Industrial Organization through Toy Models

Rebuilding Key Papers Step by Step

Felix Muñoz-Garcia

Washington State University

Preface

This book collects the handouts I prepared for EconS 594, the Ph.D. field course in Industrial Organization at Washington State University. Each chapter takes one published paper, or a small group of related papers, and rebuilds its main results in a toy model: a version with specific functional forms and numbers, simple enough that every equilibrium can be computed by hand, yet rich enough to reproduce the paper's key mechanisms. The goal is to help students read the original papers, not to replace them. Once a student has seen a result emerge from a model with two firms and linear demand, the general version in the paper becomes much easier to follow.

What this book is about

Most chapters follow the same structure. A short introduction states the economic question, explains the trade-off behind it, and gives the paper's answer in plain words. A section on the toy model then lists its ingredients: players, timing, functional forms, and parameter values. The analysis proceeds in steps (first step, second step, and so on), solved by backward induction when the model has several stages. Each step ends with a result that states its conclusion, followed by a figure or table that illustrates it and an intuitive explanation of why it holds. Most chapters close with an extension that relaxes one of the toy model's assumptions, and with a short list of what the toy model keeps from the paper and what it leaves out. When a paper contains a misprint or a claim that the toy model qualifies, a footnote says so.

The chapters are grouped by topic: price discrimination; entry barriers; mergers; collusion and cartels; regulation under asymmetric information; signaling, disclosure, and certification; information design and persuasion; advertising; platforms and network effects; innovation and patents; and environmental regulation and innovation. Each part opens with a one-page introduction that states the question its chapters share and explains how each chapter relates to the others. The chapters can be read in any order, since each one is self-contained and introduces its own notation. Students who are new to a topic may want to start with the first chapter of each part, which usually presents the most basic model of that part.

Every number and closed form in the book, and every figure drawn from a toy model's equations, was verified with a computer script that re-derives it from the toy model's primitives, using symbolic algebra and, where needed, numerical methods. These scripts are available upon request, and they make a useful starting point for students who want to change a parameter value and see how a result responds.

Acknowledgements

I thank the students of EconS 594 for their questions and comments on earlier versions of these handouts.

Felix Muñoz-Garcia

Pullman, Washington, Fall 2026

Contents

Preface **i**

I Price Discrimination 1

1 Price Discrimination and Demand Curvature **2**

Aguirre, Cowan, and Vickers (2010), "Monopoly Price Discrimination and Demand Curvature," *American Economic Review* 2

2 Price Discrimination in Oligopoly **19**

Dertwinkel-Kalt and Wey (2023), "Third-Degree Price Discrimination in Oligopoly When Markets Are Covered," *Journal of Industrial Economics* 19

3 Two-Part Tariffs with Two Firms **33**

Armstrong and Vickers (2001), "Competitive Price Discrimination," *RAND Journal of Economics* 33

4 Customer Information Sharing **46**

Liu and Serfes (2006), "Customer Information Sharing among Rival Firms," *European Economic Review* 46

II Entry Barriers 56

5 Discounts as a Barrier to Entry **57**

Ide, Montero, and Figueroa (2016), "Discounts as a Barrier to Entry," *American Economic Review* 57

6 Vertical Practices and Exclusion **68**

Asker and Bar-Isaac (2014), "Raising Retailers' Profits: On Vertical Practices and the Exclusion of Rivals," *American Economic Review* 68

7 Dynamic Vertical Foreclosure **79**

Fumagalli and Motta (2020), "Dynamic Vertical Foreclosure," *Journal of Law and Economics* 79

8 Vertical Contracts and Downstream Entry **90**

Milliou and Petrakis (2024), "Vertical Contracts and Downstream Entry," *Journal of Industrial Economics* 90

9 Entry Deterrence by Oligopolists **102**

Harrington (1987), "Oligopolistic Entry Deterrence under Incomplete Information," *RAND Journal of Economics* 102

10 Limit Pricing with Conflicting Incumbents **113**

Schultz (1999), "Limit Pricing when Incumbents Have Conflicting Interests," *International Journal of Industrial Organization* 113

III Mergers **125**

11 The Efficiency Defense in Merger Control **126**

Lagerlöf and Heidhues (2005), "On the Desirability of an Efficiency Defense in Merger Control," *International Journal of Industrial Organization* 126

12 Pricing Pressure and Pass-Through **136**

Jaffe and Weyl (2013), "The First-Order Approach to Merger Analysis," *American Economic Journal: Microeconomics* 136

13 Dynamic Merger Review **144**

Nocke and Whinston (2010), "Dynamic Merger Review," *Journal of Political Economy* . 144

14 Merger Policy with Merger Choice **151**

Nocke and Whinston (2013), "Merger Policy with Merger Choice," *American Economic Review* 151

15 Mergers and Product Innovation **157**

Federico, Langus, and Valletti (2018), "Horizontal Mergers and Product Innovation," *International Journal of Industrial Organization* 157

16 Mergers and Cost-Reducing Investment 163

Motta and Tarantino (2021), "The Effect of Horizontal Mergers, when Firms Compete in Prices and Investments," *International Journal of Industrial Organization* 163

17 Mergers and Duplicative Research 175

Federico, Langus, and Valletti (2017), "A Simple Model of Mergers and Innovation," *Economics Letters* 175

Denicolò and Polo (2018), "Duplicative Research, Mergers and Innovation," *Economics Letters* 175

18 Market Power in Customer Markets 187

Somaini and Einav (2013), "A Model of Market Power in Customer Markets," *Journal of Industrial Economics* 187

IV Collusion and Cartels 200

19 Forward Trading and Collusion 201

Liski and Montero (2006), "Forward Trading and Collusion in Oligopoly," *Journal of Economic Theory* 201

20 Leniency and Whistleblower Programs 212

Aubert, Rey, and Kovacic (2006), "The Impact of Leniency and Whistle-Blowing Programs on Cartels," *International Journal of Industrial Organization* 212

21 Collusion with Private Costs 224

Athey and Bagwell (2001), "Optimal Collusion with Private Information," *RAND Journal of Economics* 224

V Regulation under Asymmetric Information 237

22 Regulating Firms with Unknown Costs 238

Baron and Myerson (1982), "Regulating a Monopolist with Unknown Costs," *Econometrica* 238

Wang (2000), "Regulating an Oligopoly with Unknown Costs," *International Journal of Industrial Organization* 238

23 Jump-Starting a Frozen Market **245**

Tirole (2012), "Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning," *American Economic Review* 245

VI Signaling, Disclosure, and Certification 251

24 Communicating Quality **252**

Daughety and Reinganum (2008), "Communicating Quality: A Unified Model of Disclosure and Signalling," *RAND Journal of Economics* 252

25 Signaling a Clean Product with Its Price **265**

Mahenc (2008), "Signaling the Environmental Performance of Polluting Products to Green Consumers," *International Journal of Industrial Organization* 265

26 Penetration or Skimming Pricing **278**

Mahenc (2025), "Penetration or Skimming Pricing for Credence Products?" *Journal of Economic Dynamics and Control* 278

27 Label Confusion and the Groucho Effect **290**

Harbaugh, Maxwell, and Roussillon (2011), "Label Confusion: The Groucho Effect of Uncertain Standards," *Management Science* 290

28 Eco-labeling with Noisy Certification Tests **311**

Mason (2011), "Eco-Labeling and Market Equilibria with Noisy Certification Tests," *Environmental and Resource Economics* 311

29 Fraud in Green Markets **322**

Hamilton and Zilberman (2006), "Green Markets, Eco-Certification, and Equilibrium Fraud," *Journal of Environmental Economics and Management* 322

VII Information Design and Persuasion 331

30 Bayesian Persuasion **332**

Kamenica and Gentzkow (2011), "Bayesian Persuasion," *American Economic Review* . . 333

Gentzkow and Kamenica (2014), "Costly Persuasion," *American Economic Review: Papers & Proceedings* 340

Alonso and Cámara (2016), "Bayesian Persuasion with Heterogeneous Priors," *Journal of Economic Theory* 346

Gentzkow and Kamenica (2017), "Competition in Persuasion," *Review of Economic Studies* 349

31 Funding a Review Agency under Obfuscation **360**

Figueroa and Lemus (2024), "The Benefit of the Doubt: Patent Examination under Strategic Obfuscation," *Management Science* 360

VIII Advertising **374**

32 The Targeting of Advertising **375**

Iyer, Soberman, and Villas-Boas (2005), "The Targeting of Advertising," *Marketing Science* 375

33 Hype versus Real Information **386**

Johnson and Myatt (2006), "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review* 386

34 Ad-Financed Broadcasting **399**

Anderson and Coate (2005), "Market Provision of Broadcasting: A Welfare Analysis," *Review of Economic Studies* 399

IX Platforms and Network Effects **409**

35 The Attention Economy **410**

Chen (2026), "A Model of the Attention Economy: Data, Platforms, and Discovery," *American Economic Review* 410

36 Pricing Network Effects **422**

Fainmesser and Galeotti (2016), "Pricing Network Effects," *Review of Economic Studies* 422

Fainmesser and Galeotti (2020), "Pricing Network Effects: Competition," *American Economic Journal: Microeconomics* 422

X Innovation and Patents **435**

37 Research Joint Ventures **436**

Kamien, Muller, and Zang (1992), "Research Joint Ventures and R&D Cartels," *American Economic Review* 436

38 When Cost Reductions Lower Welfare **448**

Wang and Zhao (2007), "Welfare Reductions from Small Cost Reductions in Differentiated Oligopoly," *International Journal of Industrial Organization* 448

39 Patent Length and the Timing of R&D **456**

Parra (2019), "Sequential Innovation, Patent Policy, and the Dynamics of the Replacement Effect," *RAND Journal of Economics* 456

40 Killer Acquisitions **463**

Cunningham, Ederer, and Ma (2021), "Killer Acquisitions," *Journal of Political Economy* 463

XI Environmental Regulation and Innovation **474**

41 Taxes, Permits, and Abatement Technology **475**

Requate and Unold (2003), "Environmental Policy Incentives to Adopt Advanced Abatement Technology: Will the True Ranking Please Stand Up?" *European Economic Review* 475

42 Environmental R&D Cartels **487**

Poyago-Theotoky (2007), "The Organization of R&D and Environmental Policy," *Journal of Economic Behavior & Organization* 487

43 Competition and Green R&D **499**

Lambertini, Poyago-Theotoky, and Tampieri (2017), "Cournot Competition and 'Green' Innovation: An Inverted-U Relationship," *Energy Economics* 499

44 International R&D Cooperation **509**

Iida (2020), "International R&D Formations and Strategic Environmental Policy," *Environment and Development Economics* 509

45 Greener or Cheaper **518**

Strandholm, Espinola-Arredondo, and Munoz-Garcia (2025), "Greener or Cheaper Goods: Economies of Scope in R&D Investments," *Journal of Environmental Economics and Management* 518

Part I

Price Discrimination

The four chapters in this part ask the same question: when a firm can charge different consumers different prices, who gains and who loses? Each chapter adds one feature to the previous ones. Chapter 1 starts with a monopolist that sets one price in each of several markets. Discrimination raises its profit, but the effect on welfare depends on whether total output rises, and the curvature of demand determines whether it does. Chapter 2 adds competition and removes this output channel: when every consumer buys from some firm, discrimination leaves total output and welfare unchanged, but it raises prices on average, so firms gain at the consumers' expense. Chapter 3 changes the pricing tool. Instead of charging different per-unit prices in different markets, competing firms charge a fixed fee plus a price per unit; competition pushes the per-unit price down to marginal cost, and, as in Chapter 2, firms gain while consumers lose. Chapter 4 asks where the information needed to discriminate comes from. Firms learn their customers' preferences from past purchases and can sell these data to a rival, so the information that the first three chapters take as given becomes a choice of the firms.

Taken together, the chapters separate two effects of price discrimination: its effect on total output, which drives the monopoly result of Chapter 1, and its effect on competition among firms, which drives the results of Chapters 2 through 4.

Chapter 1

Third-Degree Price Discrimination and Demand Curvature

Companion note to Aguirre, I., S. Cowan, and J. Vickers (2010), "Monopoly Price Discrimination and Demand Curvature," American Economic Review 100(4): 1601–1615.

1.1 Introduction

A cinema charges students less than other adults. A drug maker charges a lower price in poorer countries. In both cases a monopolist sorts its customers into groups by a feature it can observe, such as age or location, and charges each group its own price. Is this good or bad for society? The firm earns more, the group whose price rises loses, and the group whose price falls gains, so the answer is not obvious.

Moving from one price to several prices changes welfare, the sum of consumer surplus and profit, through two channels. The first is a misallocation effect. With one price, the last consumer served in every market values the good at that price, so no other way of dividing the same output between the markets would raise welfare. With different prices, some units leave consumers who value them more and go to consumers who value them less, which always lowers welfare. The second is an output effect: price exceeds marginal cost, so an increase in total output raises welfare. Discrimination can therefore raise welfare only if it raises total output.

Aguirre et al. (2010) show that the answer depends on the curvature of demand, that is, on how the slope of demand changes with the price. Under a mild condition, if demand in the market whose price rises is at least as convex as demand in the market whose price falls, discrimination lowers welfare. For welfare to rise, demand must be more convex in the market whose price falls.

The toy model uses the pair of demands in the paper's Example 1, one linear and one exponential, and moves a single parameter to produce every case. The first step finds the prices with and without discrimination. The second step lets the firm widen the gap between the two prices little by little and splits the welfare effect into its parts. The third to fifth steps show when welfare falls, when it

rises, and what happens to output, the sixth and seventh steps treat two families of demand curves, and the eighth step lists what the toy model keeps and drops.

1.2 The Toy Model

Firm and markets. A monopolist with constant marginal cost $c = 1$ sells in two markets whose demands are independent of each other. Utility is quasi-linear, so social welfare W is consumer surplus plus profit. Demand in a generic market is $q(p)$, which is twice differentiable and decreasing, and the profit in a market is $\pi(p) = (p - c)\, q(p)$. With third-degree price discrimination the firm sets a price in each market; without it, the firm must charge the same *uniform* price in both. Throughout, both markets are served at the uniform price, so discrimination does not open a new market.

Curvature. Four objects describe a demand function. The price elasticity is $\eta \equiv -\frac{pq'}{q}$ and the Lerner index is $L \equiv \frac{p-c}{p}$. The curvature of direct demand is $\alpha \equiv -\frac{pq''}{q'}$, the elasticity of the slope of demand, and the curvature of inverse demand $p(q)$ is

$$\sigma \equiv -\frac{q\,p''}{p'} = \frac{q\,q''}{(q')^2}.$$

The two curvatures are linked by

$$\sigma = \frac{\alpha}{\eta}. \tag{1.1}$$

Linear demand has $\alpha = \sigma = 0$, convex demand has $\alpha > 0$ and $\sigma > 0$ (inverse demand is convex too), and concave demand has $\alpha < 0$ and $\sigma < 0$. The second derivative of profit is

$$\pi''(p) = 2q' + (p - c)\, q'' = q' \left(2 - L\alpha\right), \tag{1.2}$$

so profit is strictly concave if and only if $2 - L\alpha > 0$, which is assumed throughout.

The two demands. Market L has linear demand and market E has exponential demand,

$$q_L(p) = 9 - p, \qquad q_E(p) = 8\, e^{-(p-1)/b}, \qquad b > 0. \tag{1.3}$$

The two demands are normalized to be equal at the cost price, $q_L(c) = q_E(c) = 8$, so the scale b of market E is the only free parameter. Demand E is convex, with $\alpha > 0$, while demand L is the borderline linear case. In market L, $\alpha_L = \sigma_L = 0$ and $\eta_L = \frac{p}{9-p}$. In market E, $\sigma_E = 1$ and $\alpha_E = \eta_E = \frac{p}{b}$, so

$$2 - L\alpha_E = 2 - \frac{p - 1}{b},$$

and profit is strictly concave at every price below $1 + 2b$; the examples below satisfy this condition. Writing the margin as $p - c$, demand E falls by the factor e^{-1} when the margin rises by b.

1.3 First Step: Pricing with and without Discrimination

Discriminatory prices. With discrimination the firm maximizes π in each market, so

$$\pi'(p^*) = q(p^*) + (p^* - c)\, q'(p^*) = 0 \quad \Longleftrightarrow \quad L^*\eta^* = 1, \tag{1.4}$$

the Lerner condition. Since $\alpha = \sigma\eta$, it follows that $L^*\alpha^* = \sigma^*$ and $2 - L^*\alpha^* = 2 - \sigma^* > 0$ at the discriminatory price. In the toy model,

$$p_L^* = \frac{9 + c}{2} = 5, \qquad p_E^* = b + c = b + 1. \tag{1.5}$$

The uniform price. Without discrimination the firm chooses the price $\overline{p}$ that solves $\pi'_L(\overline{p}) + \pi'_E(\overline{p}) = 0$, that is,

$$10 - 2\overline{p} + 8\, e^{-(\overline{p}-1)/b}\Big(1 - \frac{\overline{p} - 1}{b}\Big) = 0, \tag{1.6}$$

which has no closed form. In terms of the margin $m = \overline{p} - c$, it reads

$$8 - 2m + 8e^{-m/b}\Big(1 - \frac{m}{b}\Big) = 0.$$

Since the two marginal profits cancel at $\overline{p}$, one market has $\pi' < 0$ and the other has $\pi' > 0$ there. Write w for the market where marginal profit is negative at $\overline{p}$ and s for the market where it is positive,

$$\pi'_w(\overline{p}) = q_w\,[1 - L\eta_w] < 0, \qquad \pi'_s(\overline{p}) = q_s\,[1 - L\eta_s] > 0.$$

Result 1 (weak and strong markets). At the uniform price:

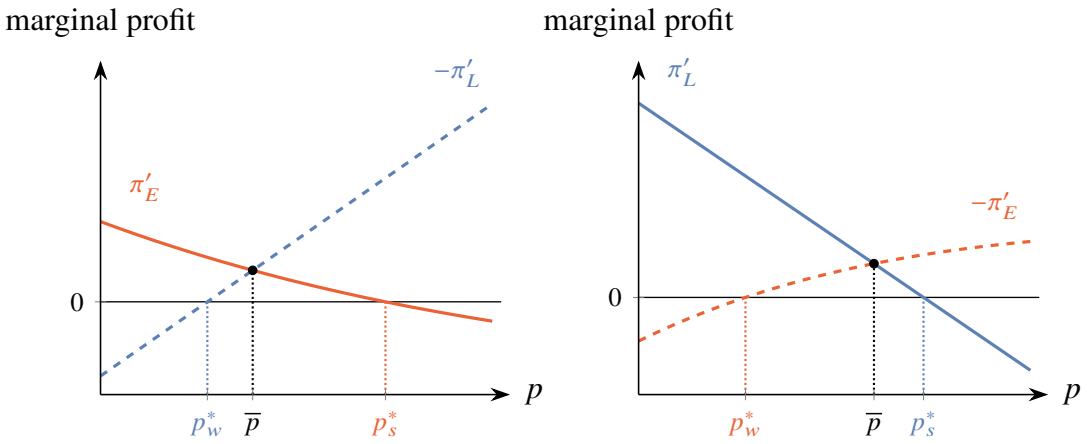
- the weak market w has the higher elasticity, $\eta_w(\overline{p}) > \eta_s(\overline{p})$;
- the uniform price lies between the discriminatory prices, $p_w^* < \overline{p} < p_s^*$, because profit is strictly concave in each market;
- with discrimination, the price falls in the weak market and rises in the strong market.

In the toy model, market E is the weak market if and only if $b + 1 < 5$, that is, $b < 4$. Table 1.1 reports the three values of b used below, $b = 5, 3, 1.8$.

b	Weak market	p_w^*	p_s^*	$\overline{p}$	$\eta_w(\overline{p})$	$\eta_s(\overline{p})$
5	L	5	6	5.2546	1.4030	1.0509
3	E	4	5	4.7217	1.5739	1.1036
1.8	E	2.8	5	4.4604	2.4780	0.9825

Table 1.1: Discriminatory and uniform prices in the toy model, with $c = 1$ and $q_E(p) = 8\, e^{-(p-1)/b}$.

Figure 1.1 depicts how the uniform price is found, with the price on the horizontal axis. The solid curve is the marginal profit of the strong market, which is positive at $\overline{p}$, and the dashed curve is minus the marginal profit of the weak market. The uniform price is where the two curves cross, so that $\pi'_s(\overline{p}) = -\pi'_w(\overline{p})$. Panel (a) has $b = 5$, where market L is weak, and panel (b) has $b = 3$, where the roles are reversed.



(a) $b = 5$: $w = L$, $s = E$ (b) $b = 3$: $w = E$, $s = L$

Figure 1.1: Marginal profit of the strong market (solid) and minus the marginal profit of the weak market (dashed), with $c = 1$ and $q_E(p) = 8\, e^{-(p-1)/b}$.

1.4 Second Step: Relaxing the Price-Difference Constraint

The paper moves from the uniform price to discrimination in small steps. The firm chooses its prices subject to the constraint $p_s - p_w \leq r$, where $r \geq 0$ is the degree of discrimination allowed. When the constraint binds, the firm maximizes $\pi_w(p_w) + \pi_s(p_w + r)$, so

$$\pi'_w(p_w) + \pi'_s(p_w + r) = 0. \tag{1.7}$$

At $r = 0$ the firm charges the uniform price $\overline{p}$, and the constraint stops binding at $r^* = p^*_s - p^*_w$, which equals 1 for $b = 5$ and $b = 3$ and equals 2.2 for $b = 1.8$. Differentiating (1.7) with respect to r,

$$p'_w(r) = \frac{-\pi''_s}{\pi''_w + \pi''_s} < 0, \qquad p'_s(r) = \frac{\pi''_w}{\pi''_w + \pi''_s} > 0, \tag{1.8}$$

so the price in the weak market falls and the price in the strong market rises as r grows, both in the direction of the discriminatory prices.

The welfare effect. The marginal effect of a price on welfare in a market is $\frac{dW}{dp} = (p - c)\, q'(p)$: a price change moves quantity, and each unit of quantity has social value equal to the price-cost margin. Hence

$$W'(r) = (p_w - c)\, q'_w\, p'_w(r) + (p_s - c)\, q'_s\, p'_s(r). \tag{1.9}$$

Adding and subtracting $\overline{p}$ and writing

$$Q'(r) \equiv q'_w\, p'_w + q'_s\, p'_s$$

for the marginal effect on total output,

$$W'(r) = \underbrace{(p_w - \overline{p})\, q'_w\, p'_w + (p_s - \overline{p})\, q'_s\, p'_s}_{\text{misallocation effect}} + \underbrace{(\overline{p} - c)\, Q'(r)}_{\text{value of the output effect}}. \tag{1.10}$$

The output term values every unit gained or lost at the same amount, $\overline{p} - c$, as if each unit were worth $\overline{p}$ to the consumer who buys it. The misallocation effect corrects this for the consumers who actually buy or stop buying, and it has two parts, one in each market. Both lower welfare.

Starting from the uniform price. At $r = 0$ the firm charges $\overline{p}$ in both markets, so the last consumer served in each market values the good at $\overline{p}$. A unit taken from one market and given to the other is then worth the same to both consumers, and moving it does not change welfare. For this reason both parts of the misallocation effect are zero at $r = 0$, and $W'(0) = (\overline{p} - c)\, Q'(0)$ is the value of the output effect alone.

A marginal increase in r. Raising r a little above zero lowers p_w below $\overline{p}$ and raises p_s above it, by (1.8). This creates the two losses:

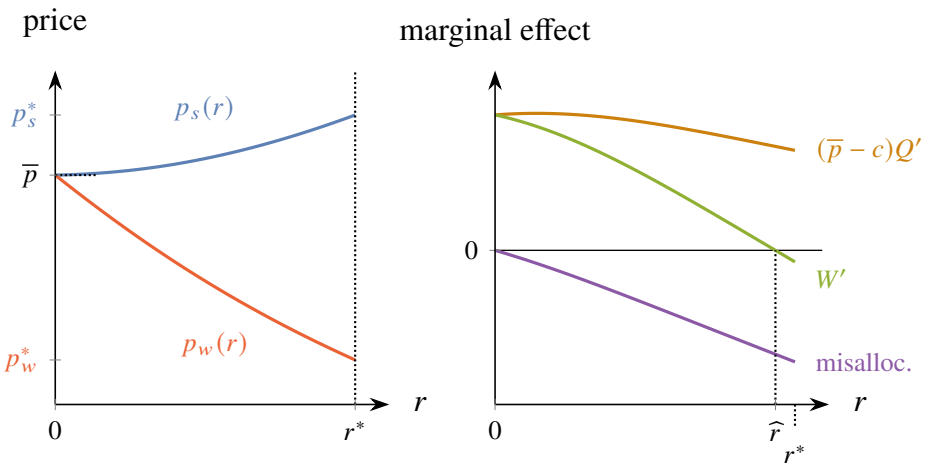
- *Weak market: units go to consumers who value them below $\overline{p}$.* The lower price p_w brings in consumers whose value lies between p_w and $\overline{p}$, and output rises by $q'_w\, p'_w > 0$. The output term counts each new unit as worth $\overline{p}$, but the consumer who buys it values it at only p_w, so the output term overstates its value by $\overline{p} - p_w$. The correction is the first term of the misallocation effect, $(p_w - \overline{p})\, q'_w\, p'_w < 0$.
- *Strong market: units leave consumers who value them above $\overline{p}$.* The higher price p_s drives out consumers whose value lies between $\overline{p}$ and p_s, and output falls by $-q'_s\, p'_s > 0$. The output term counts each lost unit as worth $\overline{p}$, but the consumer who stops buying valued it at p_s, so the output term understates the loss by $p_s - \overline{p}$. The correction is the second term, $(p_s - \overline{p})\, q'_s\, p'_s < 0$.

Each loss is a price gap times a change in quantity. At $r = 0$ the gaps $\overline{p} - p_w$ and $p_s - \overline{p}$ are zero, so the first small step away from the uniform price causes no misallocation. As r grows the gaps widen, and every further increase in r moves units from consumers who value them at p_s to consumers who value them at p_w, with $p_s - p_w = r$. When total output does not change, as with two linear demands, the units added in the weak market equal the units lost in the strong market, and the two losses add up to $-r$ times the units moved: each unit moved costs the price gap r. Aguirre (2008) depicts the two losses as two areas, one in each market, that add up over the whole move from $\overline{p}$ to the discriminatory prices.

Integrating (1.10) over $[0, r^*]$, the total welfare effect equals the two negative misallocation losses plus $(\overline{p} - c)$ times the total change in output, so an increase in output is necessary for welfare to rise.

Figure 1.2 depicts the case $b = 1.8$, with r on the horizontal axis. Panel (a) shows the two prices leaving $\overline{p} = 4.4604$ in opposite directions until they reach $p_w^* = 2.8$ and $p_s^* = 5$ at $r^* = 2.2$. Panel (b) shows the decomposition (1.10). At $r = 0$ the misallocation effect is zero, so $W'(0) = 2.109$ is the value of the output effect. As r grows the misallocation effect becomes more negative, falling to -1.74 at r^*, while the output term, after a small initial rise, falls to 1.56, and $W'(r)$ turns negative at $\widehat{r} = 2.0566$. Most of the misallocation comes from the weak market E: at $r = 1$ the weak-market loss is -0.71 and the strong-market loss is -0.03, and at r^* they are -1.49 and -0.24. The reason is that most of the price change happens in E: at $r = 1$ the price in E is 0.88 below $\overline{p}$, while the price in L is only 0.12 above it, so both the price gap and the change in quantity are larger in E.

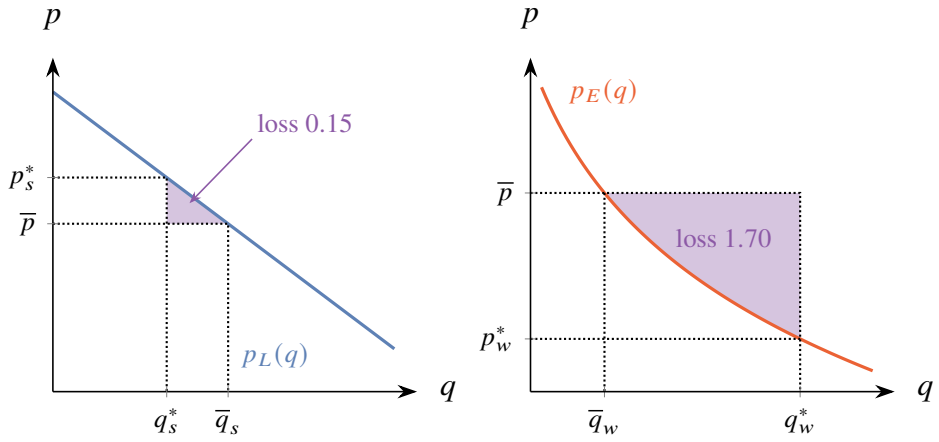
Figure 1.3 depicts the two losses over the whole move from the uniform price to the discriminatory prices, for $b = 1.8$, in the way Aguirre (2008) draws them. In panel (a), the strong market L loses



(a) Prices (b) Decomposition of $W'(r)$

Figure 1.2: Prices and the decomposition of the marginal welfare effect against the allowed price difference r, with $b = 1.8$; the labels in panel (b) name the curves that end at r^*: the output term $(\overline{p} - c)Q'(r)$, the marginal welfare effect $W'(r)$, and the misallocation effect.

the 0.54 units between $q_s^* = 4$ and $\overline{q}_s = 4.54$. These consumers value the good between $\overline{p}$ and $p_s^* = 5$, and the shaded area, 0.15, is the value they lose above $\overline{p}$. In panel (b), the weak market E gains the 1.77 units between $\overline{q}_w = 1.17$ and $q_w^* = 2.94$. These consumers value the good between $p_w^* = 2.8$ and $\overline{p}$, and the shaded area, 1.70, is the amount by which their value falls short of $\overline{p}$. The two areas add up to the misallocation loss of 1.84 reported in Table 1.3 in the fifth step.



(a) Strong market L (b) Weak market E

Figure 1.3: The two misallocation losses from full discrimination, with $b = 1.8$ and $c = 1$: units lost in the strong market by consumers who value them above $\overline{p}$ (panel a), and units gained in the weak market by consumers who value them below $\overline{p}$ (panel b).

1.5 Third Step: The Increasing Ratio Condition and When Welfare Falls

The paper rewrites (1.9) using a ratio that depends on curvature. Define

$$z(p) \equiv \frac{(p-c)\,q'(p)}{2q'(p) + (p-c)\,q''(p)} = \frac{p-c}{2 - L\alpha}, \tag{1.11}$$

the ratio of the marginal effect of the price on welfare to the second derivative of profit. Writing $(p-c)q' = z\,\pi''$ in (1.9) and using (1.8),

$$W'(r) = \underbrace{\frac{-\pi_w''\,\pi_s''}{\pi_w'' + \pi_s''}}_{>0} \Big[z_w\big(p_w(r)\big) - z_s\big(p_s(r)\big)\Big], \tag{1.12}$$

so the marginal welfare effect has the sign of $z_w - z_s$. In the toy model,

$$z_L(p) = \frac{p-1}{2}, \qquad z_E(p) = \frac{b\,(p-1)}{2b - (p-1)}. \tag{1.13}$$

The increasing ratio condition (IRC). The ratio $z(p)$ is increasing in p in each market. It holds for linear demands, for inverse demands with constant positive curvature (which include exponential and constant-elasticity demands), for direct demands with constant curvature, positive or negative, for logit and probit demands, and for demands derived from the lognormal distribution; it always holds near marginal cost, since $z'(c) = \frac{1}{2}$. It is not universal: with constant negative curvature of inverse demand it fails at high enough prices. In the toy model, both z_L and z_E are increasing on the range where profit is strictly concave.

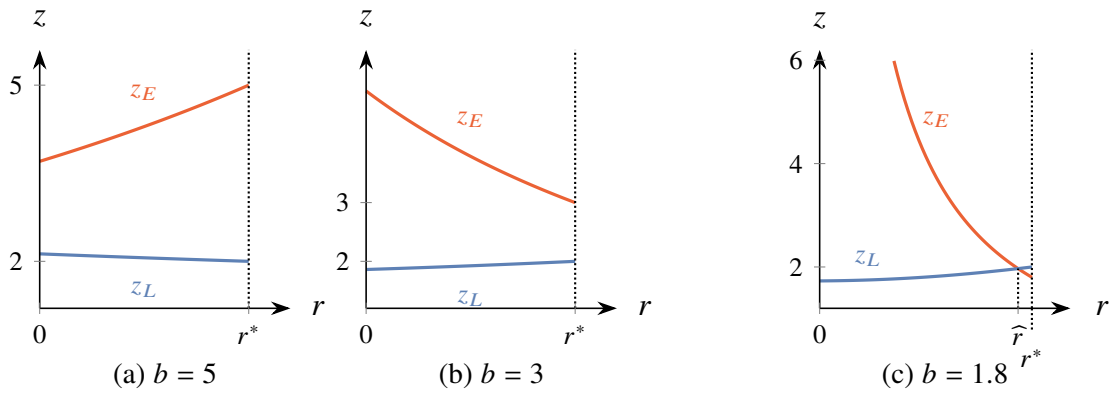
Result 2 (single peak). Under the IRC, welfare $W(r)$ is either decreasing in r, or increasing in r, or first increasing and then decreasing.

Intuitively, as r grows, p_w falls and p_s rises, so under the IRC $z_w(p_w(r))$ falls and $z_s(p_s(r))$ rises. The gap $z_w - z_s$ in (1.12) therefore decreases in r and changes sign at most once, from positive to negative. The paper states this as a Lemma: at any $\widehat{r}$ with $W'(\widehat{r}) = 0$, $W''(\widehat{r}) < 0$. Which of the three shapes arises depends on the sign of $W'(0)$ and of $W'(r^*)$.

Figure 1.4 depicts the two ratios along the constraint path, with r on the horizontal axis, for $b = 5$, $b = 3$, and $b = 1.8$. In panel (a) the strong market has the higher ratio for every r, so $W' < 0$ throughout; in panel (b) the weak market does, so $W' > 0$ throughout; and in panel (c) the curves cross at $\widehat{r} = 2.0566$, so W' changes sign.

When discrimination lowers welfare. At $r = 0$ the two markets share the price $\overline{p}$ and the Lerner index, so $z_w(\overline{p}) - z_s(\overline{p})$ has the sign of $\alpha_w(\overline{p}) - \alpha_s(\overline{p})$, the difference in the curvature of direct demand. By Result 2, if $W'(0) \le 0$ then welfare is decreasing for every $r > 0$.

Result 3 (Proposition 1 of the paper). Under the IRC, if the direct demand in the strong market is at least as convex as in the weak market at the uniform price, $\alpha_s(\overline{p}) \ge \alpha_w(\overline{p})$, then discrimination reduces welfare.



(a) $b = 5$ (b) $b = 3$ (c) $b = 1.8$

Figure 1.4: The ratios $z_w(p_w(r))$ and $z_s(p_s(r))$ against the allowed price difference r, for L (blue) and E (red); in panel (c) the curve for E is not drawn where it exceeds 6, which is the case for $r < 0.77$.

At the uniform price the misallocation effect is zero, so the welfare effect has the sign of the output effect, and the condition on α is the condition under which output does not rise locally. The IRC extends this local result to every larger r. The condition covers the case in which demand in the strong market is an affine transformation of demand in the weak market, $q_s(p) = M + N\,q_w(p)$ with $M, N > 0$ (Cowan, 2007), since the curvature of direct demand at a given price is then the same in both markets. In the toy model, $\alpha_E(\overline{p}) = \frac{\overline{p}}{b} > 0 = \alpha_L$, so Result 3 applies whenever E is the strong market, that is, $b > 4$. For $b = 5$, $\alpha_E(5.2546) = 1.0509 > 0$ and discrimination lowers welfare by 1.13. Conversely, for welfare to rise, direct demand in the weak market must be strictly more convex than in the strong market at the uniform price.

Figure 1.5 depicts why a more convex weak market gains more. The inverse demand $p_1(q) = 10 - q$ is linear, and $p_2(q) = 10 - q + 0.06\,(q - \overline{q})^2$ is more convex, with the same price and slope at the uniform quantity $\overline{q} = 3$, so the uniform price $\overline{p} = 7$ is the same for both. Their marginal revenue curves also cross at $\overline{q}$, but for $q > \overline{q}$ the price p_2 is higher, so a given increase in output lowers the price by less and $MR_2(q) > MR_1(q)$. The discriminatory quantity is therefore higher with the more convex demand, $q_2 = 5.53 > q_1 = 4.5$, and the welfare gain in the market is larger for two reasons: the output increase is bigger, and the marginal social value of output, the price minus c, is higher at every quantity above $\overline{q}$.

1.6 Fourth Step: When Welfare Rises

A sufficient condition for welfare to rise is found by looking at what happens at the discriminatory prices. Since $L^*\alpha^* = \sigma^*$, the ratio (1.11) at a discriminatory price is

$$z(p^*) = \frac{p^* - c}{2 - \sigma^*}. \tag{1.14}$$

If $z_w(p_w^*) \geq z_s(p_s^*)$, then $W'(r^*) \geq 0$ by (1.12), and by Result 2 welfare is increasing for every $r < r^*$.

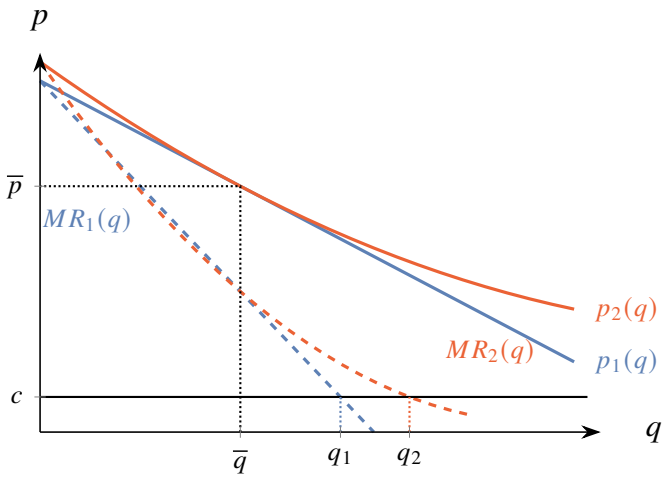


Figure 1.5: Linear inverse demand $p_1(q)$ and a more convex inverse demand $p_2(q)$ with the same price and slope at $\overline{q} = 3$, with their marginal revenue curves (dashed) and $c = 1$.

Result 4 (Proposition 2 of the paper). Under the IRC, if

$$\frac{p_w^* - c}{2 - \sigma_w^*} \geq \frac{p_s^* - c}{2 - \sigma_s^*},$$

then welfare is higher with discrimination.

Since $p_w^* < p_s^*$, the condition requires $\sigma_w^* > \sigma_s^*$: inverse demand in the weak market is more convex than in the strong market at the discriminatory prices, and the prices are close enough together that the misallocation effect is small. The condition depends only on the discriminatory prices and curvatures, so it does not depend on the multiplicative size of a market. In the toy model, with E the weak market,

$$z_E(p_E^*) = \frac{b}{2-1} = b \qquad \text{and} \qquad z_L(p_L^*) = \frac{4}{2} = 2,$$

so the condition is $b \geq 2$. For $b = 3$, discrimination raises welfare by 1.031, and by 1.527 if market E is twice as large at the cost price, $q_E(c) = 16$.

Equation (1.14) also has an empirical use. The pass-through of a cost change to the discriminatory price is

$$\frac{dp^*}{dc} = \frac{1}{2 - \sigma^*},$$

which is 1 in market E and $\frac{1}{2}$ in market L, so a regression of the monopoly price on marginal cost estimates σ^* without estimating the demand function.

The intermediate case. The opposite of the condition of Result 4 gives $W'(r^*) < 0$. Combined with $\alpha_w(\overline{p}) > \alpha_s(\overline{p})$, which gives $W'(0) > 0$, welfare first rises and then falls.

Result 5 (Proposition 3 of the paper). Under the IRC, if

- direct demand in the weak market is more convex than in the strong market at the uniform price, $\alpha_w(\overline{p}) > \alpha_s(\overline{p})$, and

- inverse demand in the strong market is at least as convex as in the weak market at the discriminatory prices, $\sigma_s^* \geq \sigma_w^*$,

then welfare rises for $r < \widehat{r}$ and falls for $r > \widehat{r}$, with a unique peak $\widehat{r} \in (0, r^*)$.

The proof uses the second condition only through $W'(r^*) < 0$, which holds in the toy model whenever E is weak and $b < 2$, even though $\sigma_E^* = 1 > \sigma_L^* = 0$. The result does not say whether full discrimination raises or lowers welfare. In the toy model the intermediate case is mild: profit in market E is strictly concave at the uniform price only for $b > 1.729$, so the case is confined to $1.729 < b < 2$. For $b = 1.8$ the peak is +2.440 at $\widehat{r} = 2.0566$ and full discrimination leaves welfare +2.427 above the uniform-price level. Figure 1.7(a) in the sixth step shows a hump that is easy to see, in a pair of demands to which Result 5 also applies, and in which full discrimination ends below the uniform-price level.

Figure 1.6 depicts $W(r) - W(0)$ for the cases of Results 3 and 4, with r on the horizontal axis: welfare falls with $b = 5$ in panel (a) and rises with $b = 3$ in panel (b).

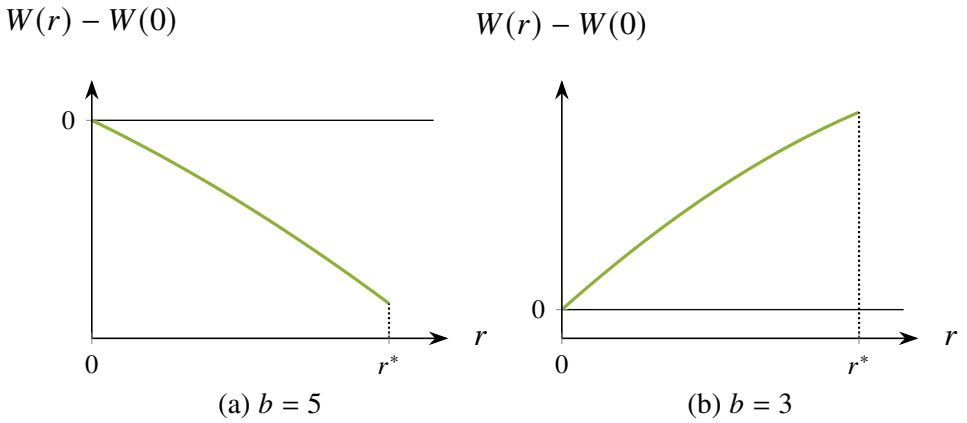


Figure 1.6: Welfare relative to the uniform price against the allowed price difference r, with $q_E(p) = 8\,e^{-(p-1)/b}$.

Table 1.2 reports the cases. The two thresholds, $b = 4$ (market E becomes the weak market) and $b = 2$ (the condition of Result 4 fails), divide the values of b into the three regimes. Both thresholds follow from the discriminatory prices alone, so they do not depend on the size of market E, and they are the same for any normalization of $q_E(c)$.

b	Result	Shape of $W(r)$	$W(r^*) - W(0)$
$b > 4$ (5)	Result 3	decreasing	−1.134
$2 \leq b < 4$ (3)	Result 4	increasing	+1.031
$1.729 < b < 2$ (1.8)	Hump, as in Result 5	rises, then falls	+2.427

Table 1.2: Welfare regimes in the toy model, with $q_E(p) = 8\,e^{-(p-1)/b}$.

1.7 Fifth Step: The Effect on Total Output

Since a rise in output is necessary for welfare to rise, the paper applies the same method to output. Using (1.8), the marginal output effect is

$$Q'(r) = \underbrace{\frac{-q'_w\, q'_s}{\pi''_w + \pi''_s}}_{>0} \left[L_w \alpha_w - L_s \alpha_s\right], \tag{1.15}$$

and, since $\alpha = \eta\sigma$,

$$L_w \alpha_w - L_s \alpha_s = L_w \eta_w \left(\sigma_w - \sigma_s\right) + \left(L_w \eta_w - L_s \eta_s\right) \sigma_s. \tag{1.16}$$

At the uniform price the Lerner indices are equal, so $Q'(0)$ has the sign of $\alpha_w(\overline{p}) - \alpha_s(\overline{p})$; at the discriminatory prices $L\eta = 1$ in each market, so $Q'(r^*)$ has the sign of $\sigma_w - \sigma_s$. In between, $L_w \leq L_s$, and the mild assumption that $L\eta$ is increasing in the price implies $L_w \eta_w > L_s \eta_s$.

Result 6 (Proposition 4 of the paper). Output changes with discrimination as follows:

- if demand is concave in the strong market and less concave, or convex, in the weak market ($\alpha_w \geq 0 \geq \alpha_s$, or $0 \geq \alpha_w \geq \alpha_s$), output increases;
- if demand is convex in the strong market and concave, or less convex, in the weak market ($\alpha_s \geq 0 \geq \alpha_w$, or $\alpha_s \geq \alpha_w \geq 0$), output decreases;
- if inverse demands are convex, and more so in the weak market ($\sigma_w \geq \sigma_s \geq 0$), output increases;
- if inverse demands are concave, and more so in the weak market ($\sigma_w \leq \sigma_s \leq 0$), output decreases.

Linear demands, where $\alpha = \sigma = 0$, sit on the boundary of all four cases, which is Pigou's (1920) result that output does not change. The paper's summary is that output rises if both direct and inverse demand are more convex in the weak market, and does not rise if both are at least as convex in the strong market. In the toy model, $\alpha_E > 0 = \alpha_L$ and $\sigma_E = 1 > 0 = \sigma_L$. If E is the weak market ($b < 4$), the first and third cases apply and output rises; if E is the strong market ($b > 4$), the second case applies and output falls. This holds regardless of the size of market E.

Table 1.3 reports the decomposition (1.10) integrated over $[0, r^*]$. In the case $b = 5$ output falls, so every term is negative and welfare falls. In the cases $b = 3$ and $b = 1.8$ output rises, and the output term exceeds the sum of the two misallocation losses in both. The loss is larger in the market whose price moves more: the strong market E when $b = 5$, and the weak market E when $b = 3$ or $b = 1.8$. The pair of demands of the sixth step with $\sigma = \frac{1}{2}$ is a case in which the misallocation term is larger: output rises by 0.211 and welfare falls by 0.491, which is the point of Varian's (1985) necessary condition: more output is necessary but not sufficient.

1.8 Sixth Step: Constant Curvature of Inverse Demand

When the two inverse demands have the same constant curvature $\sigma > 0$, Result 5 applies, because $\alpha = \sigma\eta$ is higher in the weak market, and Result 6 says that output rises. These results do not say

			Misallocation loss		
b	ΔQ	$(\overline{p}-c)\,\Delta Q$	weak market	strong market	ΔW
5	−0.218	−0.930	−0.032	−0.172	−1.134
3	+0.351	+1.306	−0.236	−0.039	+1.031
1.8	+1.233	+4.268	−1.695	−0.146	+2.427

Table 1.3: Effect of full discrimination on output, $\Delta Q = Q(r^*) - Q(0)$, and the decomposition of the welfare effect, $\Delta W = W(r^*) - W(0)$, into the output term and the two misallocation losses, with $q_E(p) = 8\,e^{-(p-1)/b}$.

whether welfare rises or falls with full discrimination. The paper answers with a second technique, which restricts quantities instead of prices. Consider inverse demands with constant curvature,

$$p_i(q) = a_i - \frac{b_i\, q^{1-\sigma}}{1-\sigma} \quad (\sigma \neq 1), \qquad p_i(q) = a_i - b_i \ln q \quad (\sigma = 1). \tag{1.17}$$

The family contains linear demand ($\sigma = 0$), exponential demand ($\sigma = 1$), and quadratic demand

$$q = \left(\frac{a-p}{2b}\right)^2 \qquad \left(\sigma = \frac{1}{2}\right),$$

and its discriminatory price is

$$p^* = \frac{(1-\sigma)a + c}{2-\sigma} \quad (\sigma \neq 1), \qquad p^* = b + c \quad (\sigma = 1).$$

Let $\overline{q}_i = q_i(\overline{p})$ and let $\Pi_i(q_i) = \pi_i(p_i(q_i))$ be profit as a function of quantity. The firm maximizes $\sum_i \Pi_i(q_i)$ subject to a limit $t \geq 0$ on how far quantities move from the uniform-price quantities,

$$-p_w'(\overline{q}_w)\,(q_w - \overline{q}_w) \leq -p_s'(\overline{q}_s)\,(q_s - \overline{q}_s) + t, \tag{1.18}$$

and the constrained-optimal quantities satisfy

$$\frac{\Pi_w'(q_w(t))}{p_w'(\overline{q}_w)} + \frac{\Pi_s'(q_s(t))}{p_s'(\overline{q}_s)} = 0. \tag{1.19}$$

At $t = 0$ the firm sells the uniform-price quantities. As t grows the firm raises quantity in the weak market and cuts it in the strong market, so prices move toward their discriminatory levels. The marginal welfare effect has the sign of

$$\delta(t) \equiv \frac{p_w - c}{2-\sigma_w}\,\frac{p_w'(\overline{q}_w)}{p_w'(q_w)} - \frac{p_s - c}{2-\sigma_s}\,\frac{p_s'(\overline{q}_s)}{p_s'(q_s)}, \tag{1.20}$$

and $\delta(0)$ has the sign of $\sigma_w - \sigma_s$. With constant σ_i, the derivative $\delta'(t)$ has the sign opposite to $1 - L\eta\sigma$ when this term has the same sign in both markets: δ keeps falling if $L\eta\sigma \leq 1$ in both markets and keeps rising if $L\eta\sigma \geq 1$ in both.

Result 7 (Proposition 5 of the paper). With constant curvature of inverse demand:

- if $1 > \sigma_s \geq \sigma_w$ and $L(\overline{p})\,\eta_w(\overline{p})\,\sigma_w \leq 1$, discrimination reduces welfare;
- if $\sigma_w \geq \sigma_s > 1$ and $L(\overline{p})\,\eta_s(\overline{p})\,\sigma_s \geq 1$, discrimination raises welfare.

If the discriminatory prices are close, $L\eta \approx 1$ and the second condition in each case follows from the first. For a common curvature σ the result is then that welfare falls if $\sigma < 1$, because the misallocation effect outweighs the positive output effect, and rises if $\sigma > 1$, because the output effect is strong enough to exceed the misallocation effect.

Figure 1.7 has two panels. Both markets have the discriminatory prices $p_w^* = 3$ and $p_s^* = 4$ and the discriminatory quantity 10, with $c = 1$, and a_i and b_i are set to match. Panel (a) depicts welfare against the allowed price difference r for $\sigma = \frac{1}{2}$: welfare rises to +0.179 at $\widehat{r} = 0.337$ and then falls to 0.491 below its uniform-price level at $r^* = 1$, the hump of Result 5 with the net loss of Result 7. Panel (b) depicts the effect of full discrimination on welfare and output against a common curvature σ, with σ on the horizontal axis. Output is unchanged at $\sigma = 0$, the linear case, falls for $\sigma < 0$, and rises for $\sigma > 0$. Welfare falls for every σ below 0.965 and rises above it; the threshold is slightly below 1 because the two discriminatory prices, 3 and 4, are not equal. At $\sigma = \frac{1}{2}$ discrimination changes welfare by −0.491 and output by +0.211, and the condition of the first case of Result 7 holds, since $L(\overline{p})\,\eta_w(\overline{p})\,\sigma_w = 0.69 \leq 1$; at $\sigma = \frac{3}{2}$ it changes welfare by +0.530 and output by +0.612, and the condition of the second case holds, since $L(\overline{p})\,\eta_s(\overline{p})\,\sigma_s = 1.38 \geq 1$.

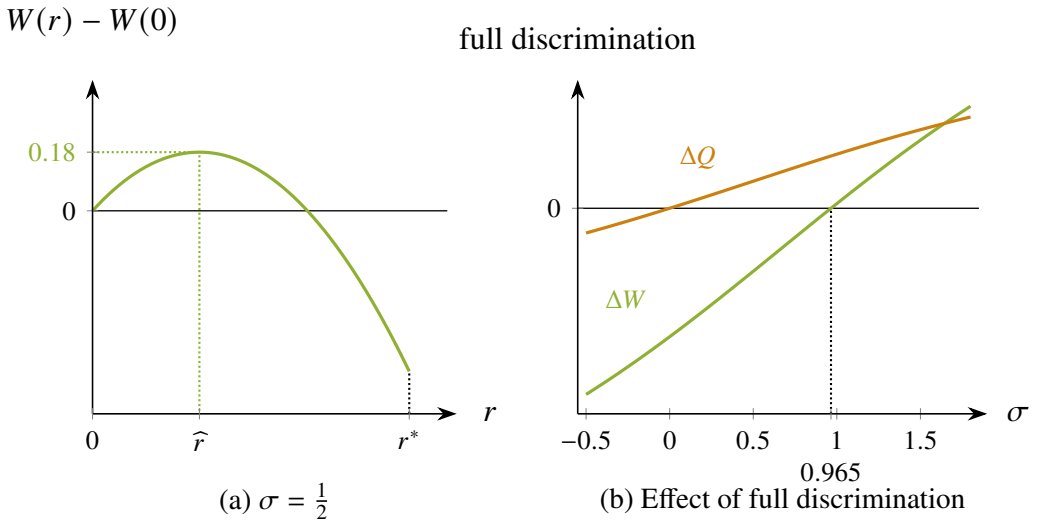


Figure 1.7: (a) Welfare against the allowed price difference r with a common curvature $\sigma = \frac{1}{2}$; (b) change in welfare (ΔW) and in output (ΔQ) from full discrimination against the common curvature σ of inverse demand; $p_w^* = 3$, $p_s^* = 4$, $q_w^* = q_s^* = 10$, and $c = 1$.

1.9 Seventh Step: Constant Elasticities

Constant-elasticity demands $q_i(p) = A_i\,p^{-\eta_i}$ with $\eta_i > 1$ are a case of constant inverse curvature, with $\sigma_i = 1 + \frac{1}{\eta_i}$, and also of constant direct curvature, with $\alpha_i = 1 + \eta_i$. The weak market has the higher elasticity $\eta_w > \eta_s$, so it has the higher α and the lower σ. Result 5 applies, so welfare first

rises and then falls with r, but Results 3, 4, 6, and 7 are silent. The discriminatory prices are

$$p_i^* = \frac{c\,\eta_i}{\eta_i - 1}. \tag{1.21}$$

Take $c = 1$, $\eta_s = 2$, $\eta_w = 3$, and A_i such that both markets sell 10 at their discriminatory prices. Then $p_w^* = 1.5$, $p_s^* = 2$, and $\overline{p} = 1.7549$.

The paper's technique is to move along a path of "harmonic mean" prices,

$$\frac{1}{\rho_i(m)} = \frac{m}{p_i^*} + \frac{1-m}{\overline{p}}, \qquad 0 \le m \le 1, \tag{1.22}$$

where $m = 0$ is the uniform price and $m = 1$ is full discrimination. Output at these prices is

$$\begin{aligned}\widetilde{Q}(m) &= \sum_i A_i\,\rho_i(m)^{-\eta_i} \\ &= \sum_i A_i \Big(\frac{m}{p_i^*} + \frac{1-m}{\overline{p}}\Big)^{\eta_i},\end{aligned}$$

a sum of strictly convex functions of m because $\eta_i > 1$. Its slope at $m = 0$ is zero by the first-order condition of the uniform price, so $\widetilde{Q}(1) > \widetilde{Q}(0)$ and output rises. Aguirre (2006) gives a short direct proof that output rises with constant-elasticity demands, for any number of markets, based on the Bernoulli inequality. Welfare $\widetilde{W}(m)$ also has zero slope at $m = 0$, and its second derivative is

$$\widetilde{W}''(m) = \sum_i -q_i'(\rho_i)\,\big[\rho_i'(m)\big]^2\,\big[L_i(\rho_i)\,(\eta_i - 1) - 1\big], \tag{1.23}$$

which is negative if $L_i(\rho_i)(\eta_i - 1) \le 1$ in both markets. This holds in the strong market because $L_s \le \frac{1}{\eta_s}$ gives

$$L_s(\eta_s - 1) \le \frac{\eta_s - 1}{\eta_s} < 1,$$

and in the weak market if $\eta_w - 1 \le \frac{1}{L(\overline{p})}$. The first-order condition of the uniform price gives

$$\frac{1}{L(\overline{p})} = s_w\eta_w + (1 - s_w)\eta_s,$$

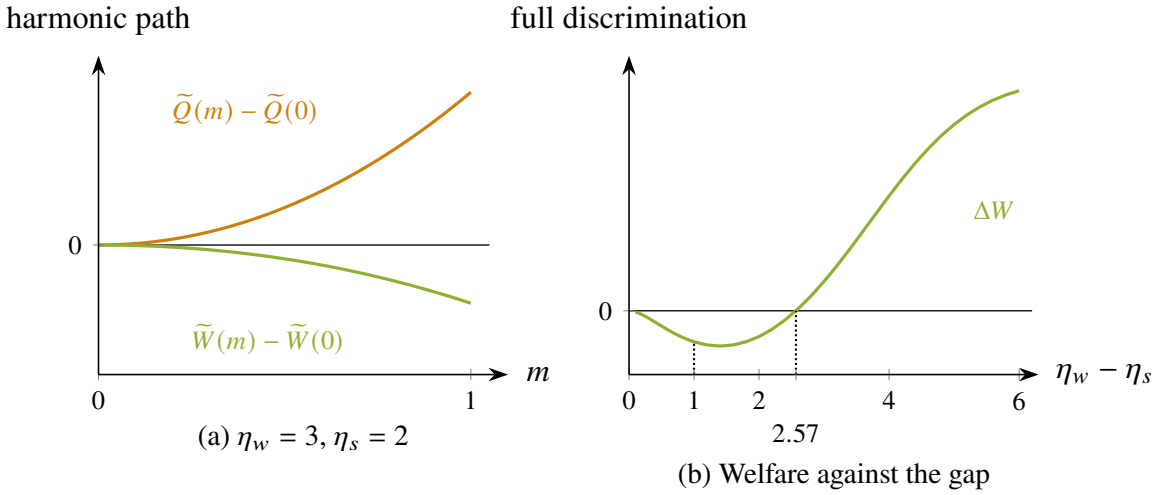
where s_w is the weak market's share of output at $\overline{p}$, so the weak-market condition is $(\eta_w - \eta_s)(1 - s_w) \le 1$.

Result 8 (Proposition 6 of the paper). With constant-elasticity demands:

- total output is higher with discrimination;
- social welfare is lower with discrimination if $(\eta_w - \eta_s)(1 - s_w) \le 1$, where $1 - s_w$ is the strong market's share of output at the uniform price, and in particular if $\eta_w - \eta_s \le 1$.

In the toy model, $s_w = 0.325$ and $\eta_w - \eta_s = 1$, so the condition holds, output rises by 0.767, and welfare falls by 0.292. Figure 1.8 depicts the argument and its limit. Panel (a) shows $\widetilde{Q}(m) - \widetilde{Q}(0)$

and $\widetilde{W}(m) - \widetilde{W}(0)$ along the harmonic path for $\eta_w = 3$: output is convex and welfare concave in m, both with zero slope at $m = 0$. Panel (b) shows the effect of full discrimination on welfare against the elasticity gap $\eta_w - \eta_s$, with $\eta_s = 2$ and both markets selling 10 at their discriminatory prices. Welfare falls for gaps up to about 2.57, well beyond the sufficient condition, and rises for larger gaps: with $\eta_w = 5$ output rises by 4.415 and welfare rises by 0.276. The condition of Result 8 is sufficient only, and it can fail when elasticities differ substantially, as Ippolito's (1980) numerical examples illustrate.



(a) $\eta_w = 3, \eta_s = 2$ (b) Welfare against the gap

Figure 1.8: Output and welfare along the harmonic-mean path of prices, and the change in welfare from full discrimination against the elasticity gap, with $c = 1$ and $q_w^* = q_s^* = 10$.

1.10 Eighth Step: What the Toy Model Keeps and Drops

The pair of linear and exponential demands, together with the constant-curvature families of the sixth and seventh steps, keeps every main conclusion of the paper:

1. Under the IRC, welfare is monotone in the degree of discrimination or has a single interior peak (the third step, Lemma).
2. Discrimination reduces welfare if direct demand in the strong market is at least as convex as in the weak market at the uniform price (the third step, Proposition 1).
3. Welfare is higher with discrimination if inverse demand in the weak market is sufficiently more convex and the discriminatory prices are close (the fourth step, Proposition 2), and welfare rises and then falls when direct demand is more convex in the weak market but inverse demand is not (the fourth step, Proposition 3).
4. Output rises if direct and inverse demand are both more convex in the weak market, and does not rise if both are at least as convex in the strong market (the fifth step, Proposition 4).
5. With constant curvature, the threshold between falling and rising welfare is $\sigma = 1$ when the discriminatory prices are close (the sixth step, Proposition 5), and with constant elasticities

discrimination raises output and lowers welfare when the elasticities are not far apart (the seventh step, Proposition 6).

The toy model also shows what the conditions of the paper leave open. The regime of Results 3 to 5 depends on the scale b of market E and not on its size, and in the intermediate case the sign of the net effect of full discrimination is not fixed by the conditions: it is positive in the linear-exponential pair and negative in the pair of quadratic demands of the sixth step.

Relative to Aguirre et al. (2010), the toy model uses two markets, and the paper notes that the results of the second to fourth steps generalize to more than two markets by constraining a weighted sum of the differences between the prices and the uniform price (Schmalensee, 1981), with weights equal to marginal profitability at the uniform price. The toy model assumes that both markets are served at the uniform price, so it leaves aside discrimination that opens new markets, and it assumes a monopolist with constant marginal cost and strictly concave profit in each market. Each of these features can be relaxed one at a time.

1.11 Connections to Other Chapters

- Chapter 2 studies third-degree price discrimination in an oligopoly where every consumer buys. Total output is then fixed, so the output effect of this chapter disappears and discrimination leaves welfare unchanged.
- Demand curvature also determines pass-through. Chapter 12 uses pass-through to turn the pricing pressure created by a merger into a predicted price change.
- Chapter 4 studies firms that use their customers' purchase histories to price discriminate, and asks when rivals trade these data.

References

Aguirre, I. (2006). "Monopolistic Price Discrimination and Output Effect under Conditions of Constant Elasticity Demand." *Economics Bulletin*, 4(23), pp. 1–6.

Aguirre, I. (2008). "Output and Misallocation Effects in Monopolistic Third-Degree Price Discrimination." *Economics Bulletin*, 4(11), pp. 1–11.

Aguirre, I., S. Cowan, and J. Vickers (2010). "Monopoly Price Discrimination and Demand Curvature." *American Economic Review*, 100(4), pp. 1601–1615. https://doi.org/10.1257/aer.100.4.1601

Cowan, S. (2007). "The Welfare Effects of Third-Degree Price Discrimination with Nonlinear Demand Functions." *RAND Journal of Economics*, 38(2), pp. 419–428.

Ippolito, R. A. (1980). "Welfare Effects of Price Discrimination When Demand Curves Are Constant Elasticity." *Atlantic Economic Journal*, 8(2), pp. 89–93.

Pigou, A. C. (1920). *The Economics of Welfare*. London: Macmillan.

Schmalensee, R. (1981). "Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination." *American Economic Review*, 71(1), pp. 242–247.

Varian, H. R. (1985). "Price Discrimination and Social Welfare." *American Economic Review*, 75(4), pp. 870–875.

Chapter 2

Third-Degree Price Discrimination in Oligopoly When Markets Are Covered

Companion note to Dertwinkel-Kalt, M., and C. Wey (2023), "Third-Degree Price Discrimination in Oligopoly When Markets Are Covered," Journal of Industrial Economics 71(2): 464–490. This chapter follows the working-paper version, DICE Discussion Paper 336 (2020).

2.1 Introduction

Firms often charge different prices in different markets. Consider two bus companies that serve both a business route and a student route, charging more on the business route. With a single seller, such price discrimination can raise or lower welfare. What happens when several firms compete in every market, and every consumer buys from one of them?

Price discrimination changes welfare through three channels. First, some consumers stop buying when their price rises. Second, buyers pay different prices in different markets. Third, with several firms, a higher price sends some consumers to a rival. Dertwinkel-Kalt and Wey (2023) shut down the first channel: markets are *covered*, meaning that every consumer buys one unit from some firm, so total demand in each market is fixed. Only switching between firms remains. Each firm then charges a high price where consumers switch reluctantly and a low price where they switch easily.

The paper shows that, with linear covered demand, discrimination changes the level of each firm's prices but not the price gaps between firms. Since only these gaps decide who buys from whom, outputs and total welfare do not change, but consumers pay more on average. Under uniform pricing, a firm's Lerner index (its markup as a share of its price) is a weighted harmonic mean of its Lerner indices under discrimination, which lies below their weighted arithmetic mean whenever its discriminatory prices differ across markets. Therefore, a ban on discrimination lowers every firm's profit and raises consumer surplus by the same amount, a gain that can be computed from observed prices and quantities alone.

The toy model has two firms and two Hotelling markets that differ only in their transport cost. The first and second steps solve the discriminatory and the uniform-pricing equilibria, the third step derives the two mean formulas, the fourth step measures the consumer surplus gain, the fifth step adds arbitrage costs, and the sixth and seventh steps draw policy implications and list what the toy model keeps and drops.

2.2 The Toy Model

Firms and markets. Two firms, 1 and 2, produce differentiated products at constant marginal costs $c_1 = 1$ and $c_2 = 4$, so firm 1 is the low-cost firm. They sell in two independent markets, A and B. In each market the consumers are spread uniformly on a Hotelling line of length one, with firm 1 at the left end and firm 2 at the right end. Every consumer buys one unit from one of the two firms, so the market is *covered*. Each market has a mass of 12 consumers, so the two markets differ only in the transport cost per unit of distance, which is 6 in market A and 3 in market B.

Demand. In market j with mass M_j and transport cost t_j, the consumer who is indifferent between the two firms sits at

$$x_0 = \frac{1}{2}\Big(1 + \frac{p_2^j - p_1^j}{t_j}\Big),$$

so firm 1 sells $M_j x_0$ and firm 2 sells the rest. Writing $a_j \equiv \frac{M_j}{2}$ and $b_j \equiv \frac{M_j}{2t_j}$,

$$D_1^j = a_j + b_j\big(p_2^j - p_1^j\big), \qquad D_2^j = a_j + b_j\big(p_1^j - p_2^j\big), \tag{2.1}$$

with $a_A = a_B = 6$, $b_A = 1$, and $b_B = 2$. This is the two-firm case of the paper's *linear covered-demand* model with m firms and n markets, in which

$$D_i^j = a_j + b_j \sum_{i' \neq i} \big(p_{i'}^j - p_i^j\big).$$

The ratio $\frac{b_j}{a_j}$ measures how intensely firms compete in market j: it is high when the transport cost is low, so that consumers switch easily. In this micro-foundation $\frac{a_j}{b_j} = t_j$, so the only difference between the markets is the ratio $\frac{a_A}{b_A} = 6$ against $\frac{a_B}{b_B} = 3$. Four properties of the demand system are used throughout, and each can be checked in (2.1):

(A1) a firm's own-price slope is $\frac{\partial D_i^j}{\partial p_i^j} = -(m-1)\, b_j$, which equals $-b_j$ for $m = 2$;

(A2) demand depends only on price *differences*: adding the same amount to every price in market j changes no one's demand;

(A3) market demand is inelastic: $\sum_i D_i^j = m\, a_j$, which is 12 in both markets;

(A4) the demand difference between two firms is determined by their price difference:

$$D_i^j - D_{i'}^j = m\, b_j\big(p_{i'}^j - p_i^j\big).$$

Figure 2.1 depicts the demand of firm 1 against the price difference $p_2^j - p_1^j$ in the two markets, as in (2.1). At a zero price difference firm 1 sells half of each market, 6 consumers, and each unit of price difference moves b_j consumers to it, so a given price difference moves more consumers in the market with the lower transport cost.

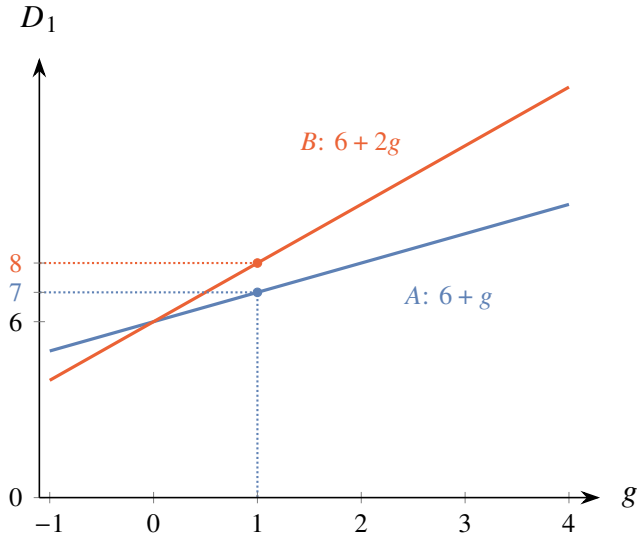


Figure 2.1: Demand of firm 1 in markets A and B against the price difference $g = p_2^j - p_1^j$; both lines equal 6 at $g = 0$, and the dotted lines mark $g = 1$.

Pricing regimes. Under *discriminatory pricing* each firm sets a price in each market, p_i^j. Under *uniform pricing* each firm must charge the same price in both markets, $\widehat{p}_i$, for example because arbitrage or a regulation forces it. Firm i's profit is $\sum_j D_i^j\,(p_i^j - c_i)$, firms choose prices simultaneously, and we look at the unique interior Nash equilibrium of each regime.

2.3 First Step: The Discriminatory Equilibrium

Under discrimination each firm solves a separate problem in each market, with first-order condition

$$-(m-1)b_j\,(p_i^j - c_i) + D_i^j = 0,$$

that is,

$$p_i^j - c_i = \frac{D_i^j}{(m-1)\,b_j}. \tag{2.2}$$

The margin of a firm equals its output divided by the slope of its demand, as for any price-setting firm. Subtract (2.2) for firm i' and for firm i in the same market and use (A4):

$$-(m-1)b_j\left[(p_{i'}^j - p_i^j) - (c_{i'} - c_i)\right] - m\,b_j\,(p_{i'}^j - p_i^j) = 0,$$

so

$$p_{i'}^j - p_i^j = \frac{m-1}{2m-1}\left(c_{i'} - c_i\right) \quad \text{in every market } j. \tag{2.3}$$

With two firms the price difference is one third of the cost difference, $\frac{3}{3} = 1$, in both markets. Using (2.3) in (A4) and (2.2), the equilibrium prices and outputs are

$$p_i^j = c_i + \frac{a_j}{(m-1)\,b_j} + \frac{1}{2m-1}\sum_{i' \neq i}\left(c_{i'} - c_i\right), \qquad D_i^j = a_j + b_j\,\frac{m-1}{2m-1}\sum_{i' \neq i}\left(c_{i'} - c_i\right). \tag{2.4}$$

A firm's price is its cost, plus a markup $\frac{a_j}{(m-1)b_j}$ that depends only on the market and falls with the intensity of competition $\frac{b_j}{a_j}$, plus a term that depends on its cost relative to its rivals. In the toy model,

$$p_1^A = 1 + 6 + 1 = 8, \quad p_2^A = 4 + 6 - 1 = 9, \quad p_1^B = 1 + 3 + 1 = 5, \quad p_2^B = 4 + 3 - 1 = 6,$$

and the outputs are

$$D_1^A = 7, \quad D_2^A = 5, \quad D_1^B = 8, \quad D_2^B = 4.$$

Market A, where consumers switch less readily, is the high-price market for both firms.

Figure 2.2 depicts the two markets at these prices. The indifferent consumer sits at $x_0 = \frac{7}{12}$ in market A and at $\frac{8}{12} = \frac{2}{3}$ in market B: the low-cost firm serves more than half of each market, and more in the market where consumers switch easily, because there the same price difference moves more consumers.

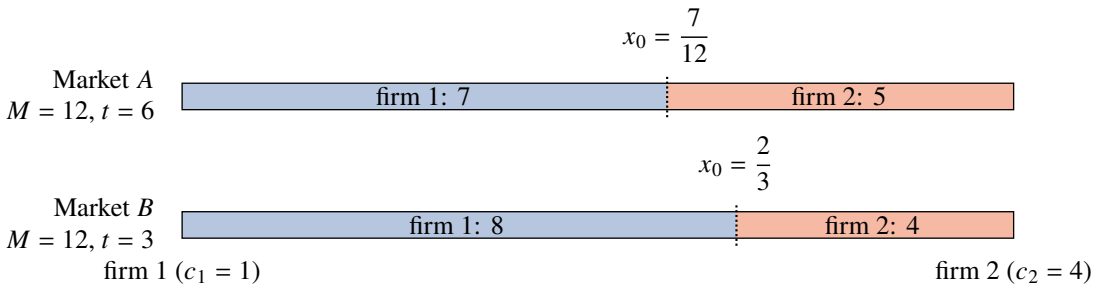


Figure 2.2: The two markets at the discriminatory prices: each firm sells to the consumers on its side of the indifferent consumer x_0.

2.4 Second Step: The Uniform-Pricing Equilibrium

Under uniform pricing firm i chooses one price $\widehat{p}_i$ for both markets. Summing its first-order conditions,

$$\sum_j \left[\, -(m-1)b_j\,(\widehat{p}_i - c_i) + D_i^j \right] = 0,$$

so

$$\widehat{p}_i - c_i = \frac{\sum_j D_i^j}{(m-1)\sum_j b_j}. \tag{2.5}$$

The same subtraction as in the first step, applied to the summed conditions, gives

$$\widehat{p}_{i'} - \widehat{p}_i = \frac{m-1}{2m-1}\,(c_{i'} - c_i),$$

so the price difference between two firms is the same under uniform pricing as under discrimination. By (A2), outputs depend only on price differences, so they are the same under both regimes.

Result 1 (Proposition 1 of the paper). In the linear covered-demand model with constant marginal costs:

- all bilateral price differences are the same under discriminatory and uniform pricing, in every market,
$$p_{i'}^j - p_i^j = \widehat{p}_{i'} - \widehat{p}_i = \frac{m-1}{2m-1}\,(c_{i'} - c_i);$$
- every firm's output in every market is the same under the two regimes,
$$D_i^j = a_j + b_j\,\frac{m-1}{2m-1}\sum_{i'\neq i}(c_{i'} - c_i).$$

Intuitively, the regime moves the *level* of each firm's prices, through its markup, but not the gaps between firms, and in a covered market only the gaps matter. The price difference $\frac{m-1}{2m-1}$ of the cost difference rises from $\frac{1}{3}$ with two firms toward $\frac{1}{2}$ as the number of firms grows. Since outputs and who serves whom do not change, total welfare does not change either: the market is covered, so no consumer is lost, and each consumer travels the same distance to the same firm.

Table 2.1 reports the two equilibria. The uniform prices are

$$\widehat{p}_1 = 1 + \frac{15}{3} = 6 \qquad \text{and} \qquad \widehat{p}_2 = 4 + \frac{9}{3} = 7,$$

whose difference is 1 as in (2.3), and each firm earns less than under discrimination.

Firm	Market	p_i^j	$\widehat{p}_i$	D_i^j	profit, discrimination	profit, uniform
1	A	8	6	7	$7 \cdot 7 = 49$	$15 \cdot 5 = 75$
	B	5		8	$8 \cdot 4 = 32$	
total for firm 1					81	75
2	A	9	7	5	$5 \cdot 5 = 25$	$9 \cdot 3 = 27$
	B	6		4	$4 \cdot 2 = 8$	
total for firm 2					33	27

Table 2.1: Discriminatory and uniform equilibria, with $c_1 = 1$ and $c_2 = 4$; outputs are the same under both regimes.

2.5 Third Step: Elasticities and Lerner Indices

Define firm i's elasticity in market j under discrimination and its aggregate elasticity under uniform pricing,

$$E_i^j \equiv -\frac{\partial D_i^j}{\partial p_i^j}\,\frac{p_i^j}{D_i^j}, \qquad \widehat{E}_i \equiv -\frac{\sum_j \frac{\partial D_i^j}{\partial p_i}\,\widehat{p}_i}{\sum_j D_i^j}, \tag{2.6}$$

and the Lerner indices

$$L_i^j \equiv \frac{p_i^j - c_i}{p_i^j} = \frac{1}{E_i^j} \qquad \text{and} \qquad \widehat{L}_i \equiv \frac{\widehat{p}_i - c_i}{\widehat{p}_i} = \frac{1}{\widehat{E}_i},$$

which the first-order conditions equate to the inverse elasticities. Let $w_i^j \equiv \frac{D_i^j}{\sum_j D_i^j}$ be the share of firm i's output sold in market j. By Result 1 the weights are the same under both regimes. Summing (2.2) over markets and comparing with the summed uniform condition,

$$\sum_j (m-1) b_j (\widehat{p}_i - c_i) = \sum_j (m-1) b_j (p_i^j - c_i),$$

which says that

$$\widehat{p}_i = \frac{\sum_j b_j p_i^j}{\sum_j b_j},$$

so the uniform price is the weighted arithmetic mean of the discriminatory prices with the weights b_j. Multiplying both sides of that equality by $(m-1)\sum_j b_j$ gives $\widehat{E}_i \sum_j D_i^j = \sum_j E_i^j D_i^j$, since $E_i^j D_i^j = (m-1)\, b_j\, p_i^j$ by (2.6); dividing by $\sum_j D_i^j$ gives the arithmetic-mean formula for elasticities, and inverting, with $L = \frac{1}{E}$, gives the harmonic-mean formula for Lerner indices.

Result 2 (Proposition 2 of the paper). For every firm i:

- (arithmetic mean formula) the aggregate elasticity under uniform pricing is the weighted arithmetic mean of the market elasticities, $\widehat{E}_i = \sum_j w_i^j E_i^j$;
- (harmonic mean formula) the Lerner index under uniform pricing is the weighted harmonic mean of the market Lerner indices,

$$\widehat{L}_i = \frac{1}{\sum_j w_i^j \frac{1}{L_i^j}}; \tag{2.7}$$

- since a harmonic mean is below the arithmetic mean unless all the numbers are equal, $\widehat{L}_i < \sum_j w_i^j L_i^j$ whenever the Lerner indices differ across markets, so uniform pricing reduces market power;
- under the same condition, the uniform price is below the quantity-weighted mean of the discriminatory prices, $\widehat{p}_i < \sum_j w_i^j p_i^j$.

The Lerner indices of firm i differ across markets if and only if its discriminatory prices differ, that is, if and only if the ratio $\frac{a_j}{b_j}$ differs across markets (given $c_i > 0$, since a firm with zero cost has a Lerner index of one in every market), by (2.4).

In the toy model, the elasticities of firm 1 are

$$E_1^A = \frac{1 \cdot 8}{7} = 1.143 \qquad \text{and} \qquad E_1^B = \frac{2 \cdot 5}{8} = 1.250,$$

and its aggregate elasticity under uniform pricing is

$$\widehat{E}_1 = \frac{3 \cdot 6}{15} = 1.2 = \frac{7}{15}\, 1.143 + \frac{8}{15}\, 1.250.$$

Its Lerner indices are $L_1^A = \frac{7}{8}$ and $L_1^B = \frac{4}{5}$, and

$$\widehat{L}_1 = \frac{1}{\frac{7}{15} \cdot \frac{8}{7} + \frac{8}{15} \cdot \frac{5}{4}} = \frac{5}{6} = 0.8333 = \frac{6-1}{6},$$

below the arithmetic mean

$$\frac{7}{15} \cdot \frac{7}{8} + \frac{8}{15} \cdot \frac{4}{5} = 0.8350.$$

For firm 2,

$$E_2^A = \frac{1 \cdot 9}{5} = 1.8, \qquad E_2^B = \frac{2 \cdot 6}{4} = 3, \qquad \widehat{E}_2 = \frac{3 \cdot 7}{9} = \frac{7}{3} = 2.333,$$

and $L_2^A = \frac{5}{9}$ and $L_2^B = \frac{1}{3}$ give

$$\widehat{L}_2 = \frac{3}{7} = 0.429 = \frac{7-4}{7},$$

below the arithmetic mean $\frac{37}{81} = 0.457$.

Figure 2.3 depicts the harmonic mean formula for firm 2, with the share w of its output sold in market A on the horizontal axis. The arithmetic mean $wL^A + (1-w)L^B$ is the straight line between the two Lerner indices, and the harmonic mean is the curve below it. At the firm's actual share $w = \frac{5}{9}$ the two means are 0.457 and 0.429, and the harmonic mean is the Lerner index $\widehat{L}_2$ that the firm earns under uniform pricing. The gap is larger the more the market Lerner indices differ.

Figure 2.4 depicts the prices of the two firms on a price axis. Each uniform price lies between the firm's two discriminatory prices, closer to the low-price market, and below the quantity-weighted mean of the discriminatory prices (6.4 for firm 1 and 7.667 for firm 2). The uniform price is the arithmetic mean with the weights b_j, $\frac{1 \cdot 8 + 2 \cdot 5}{3} = 6$, and the quantity-weighted mean gives more weight to the high-price market because the firm sells more there for each unit of b_j.

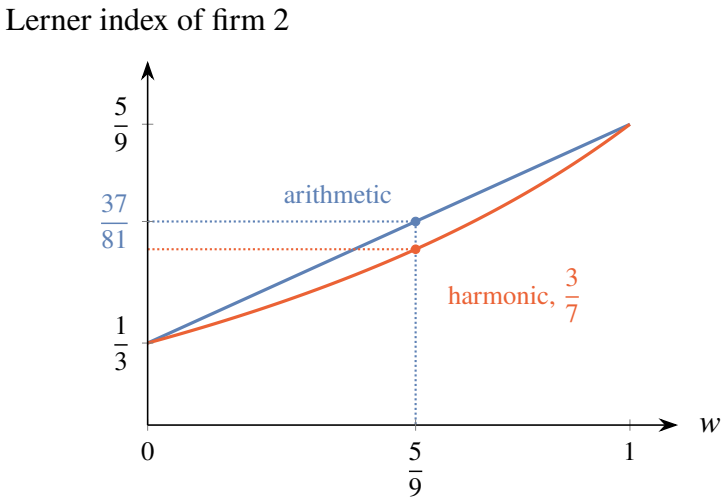


Figure 2.3: Weighted arithmetic and harmonic means of firm 2's market Lerner indices $L^B = \frac{1}{3}$ and $L^A = \frac{5}{9}$ against the share w of its output sold in market A.

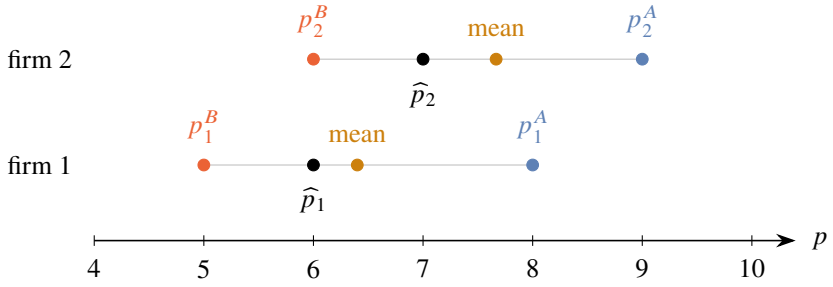


Figure 2.4: Discriminatory prices in markets B and A, uniform price $\widehat{p}_i$, and the quantity-weighted mean of the discriminatory prices (mean), for each firm.

2.6 Fourth Step: The Consumer Surplus Gain from Uniform Pricing

Since outputs and the allocation of consumers to firms do not change, total welfare is the same under the two regimes, so the gain in consumer surplus equals the loss in total profit,

$$\widehat{CS} - CS = \sum_i \pi_i - \sum_i \widehat{\pi}_i = \sum_i \sum_j \left(p_i^j - \widehat{p}_i\right) D_i^j, \tag{2.8}$$

where the last equality uses the equal outputs. In the toy model, firms 1 and 2 lose

$$81 - 75 = 6 = (8-6)\,7 + (5-6)\,8 = 14 - 8,$$
$$33 - 27 = 6 = (9-7)\,5 + (6-7)\,4,$$

respectively, so the consumer surplus gain is 12. The firms' profit loss is the same for both firms, an equality that turns out to be general (see below).

The paper's second corollary writes the gain with the discriminatory data alone. By (2.2) and (2.5), a firm's profits under discrimination and under uniform pricing are

$$\pi_i = \sum_j \frac{(D_i^j)^2}{(m-1)b_j} \qquad \text{and} \qquad \widehat{\pi}_i = \frac{\left(\sum_j D_i^j\right)^2}{(m-1)\sum_j b_j},$$

so

$$\widehat{CS} - CS = \frac{1}{m-1} \sum_i \left[\sum_j \frac{(D_i^j)^2}{b_j} - \frac{\left(\sum_j D_i^j\right)^2}{\sum_j b_j} \right], \tag{2.9}$$

where b_j can be recovered from observables by (A4): in the toy model,

$$b_A = \frac{7-5}{2\,(9-8)} = 1 \qquad \text{and} \qquad b_B = \frac{8-4}{2\,(6-5)} = 2.$$

The bracket is nonnegative by the Cauchy-Schwarz inequality, and it is zero only if $\frac{D_i^j}{b_j}$ is the same in all markets. For firm 1 it equals

$$\frac{7^2}{1} + \frac{8^2}{2} - \frac{15^2}{3} = 49 + 32 - 75 = 6.$$

Result 3 (Corollaries 1 and 2 of the paper). Relative to uniform pricing, discrimination:

- lowers consumer surplus by $\sum_i \sum_j (p_i^j - \widehat{p}_i) D_i^j > 0$, which equals the rise in total profit, and leaves total welfare unchanged;
- raises every firm's profit, so the discriminatory equilibrium is a Pareto improvement for the firms;
- produces a consumer surplus loss that can be computed from the prices and quantities of the discriminatory equilibrium alone, with (2.9).

Reading the formula. Since $D_i^j = b_j\left(\frac{a_j}{b_j} + \Delta_i\right)$, where

$$\Delta_i \equiv \frac{m-1}{2m-1}\sum_{i' \neq i}(c_{i'} - c_i)$$

is a firm-specific constant, the term Δ_i cancels from the bracket in (2.9), and the bracket becomes a variance across markets of the ratio $\frac{a_j}{b_j}$ with the weights b_j:

$$\widehat{CS} - CS = \frac{m}{m-1}\sum_j b_j\left(\frac{a_j}{b_j} - \bar{r}\right)^2, \qquad \bar{r} \equiv \frac{\sum_j a_j}{\sum_j b_j}. \tag{2.10}$$

Every firm loses the same amount,

$$\frac{1}{m-1}\sum_j b_j\left(\frac{a_j}{b_j} - \bar{r}\right)^2,$$

regardless of its cost, and the gain depends on the demand parameters only through the dispersion of $\frac{a_j}{b_j}$ across markets, that is, through how much the discriminatory prices differ across markets. With two markets,

$$\widehat{CS} - CS = \frac{m}{m-1}\,\frac{b_A b_B}{b_A + b_B}\left(\frac{a_A}{b_A} - \frac{a_B}{b_B}\right)^2 = 2 \cdot \frac{2}{3} \cdot 3^2 = 12.$$

Figure 2.5 depicts this relation against the difference

$$s = \frac{a_A}{b_A} - \frac{a_B}{b_B}$$

between the two markets, which is also the difference $p^A - p^B$ between the discriminatory prices of each firm with two firms, and the difference $t_A - t_B$ between the transport costs. The gain is zero when the markets have the same ratio and grows with the square of the difference.

Recovering the uniform price from market data. The discriminatory data also give the uniform price. From $\pi_i - \widehat{\pi}_i = \sum_j (p_i^j - \widehat{p}_i)D_i^j$ and (2.9),

$$\widehat{p}_i = \frac{\sum_j p_i^j D_i^j - (\pi_i - \widehat{\pi}_i)}{\sum_j D_i^j},$$

which for firm 1 gives $\frac{96-6}{15} = 6$ and for firm 2 gives $\frac{69-6}{9} = 7$. Neither marginal costs nor demand estimates are needed, only prices and quantities.

2.7 Fifth Step: Arbitrage Costs

Discrimination is feasible only if consumers cannot arbitrage between markets. Consider that they can, at a cost $r \geq 0$ per unit, so that each firm's price difference between the two markets cannot exceed r: $p_i^A - p_i^B \leq r$. Without the constraint, each firm's price difference is

$$r^* = p_i^A - p_i^B = \frac{1}{m-1}\left(\frac{a_A}{b_A} - \frac{a_B}{b_B}\right) = 3,$$

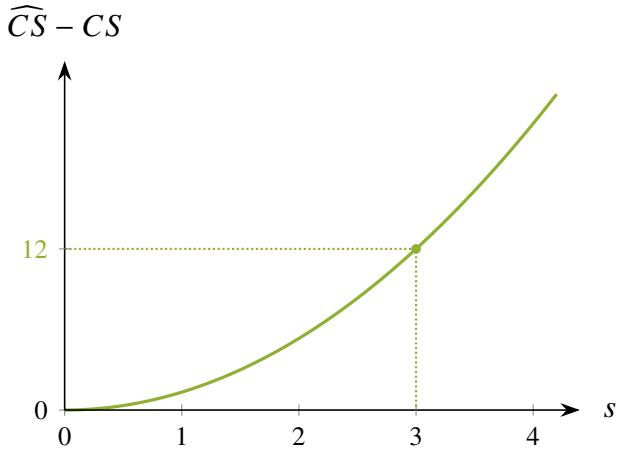


Figure 2.5: Consumer surplus gain from uniform pricing against the difference $s = \frac{a_A}{b_A} - \frac{a_B}{b_B}$, with $b_A = 1$ and $b_B = 2$.

the same for both firms, and for $r \geq r^*$ the constraint does not bind. For $r < r^*$ the constraint binds for both firms.

Result 4 (Proposition 3 of the paper). With two firms and two markets and a binding arbitrage constraint $r < r^*$, the equilibrium prices are

$$p_i^A(r) = \widehat{p}_i + \vartheta\, r, \qquad p_i^B(r) = \widehat{p}_i - (1-\vartheta)\, r, \qquad \vartheta \equiv \frac{b_B}{b_A + b_B}, \tag{2.11}$$

and all price differences between firms and all outputs remain the same as under unconstrained discrimination and uniform pricing.

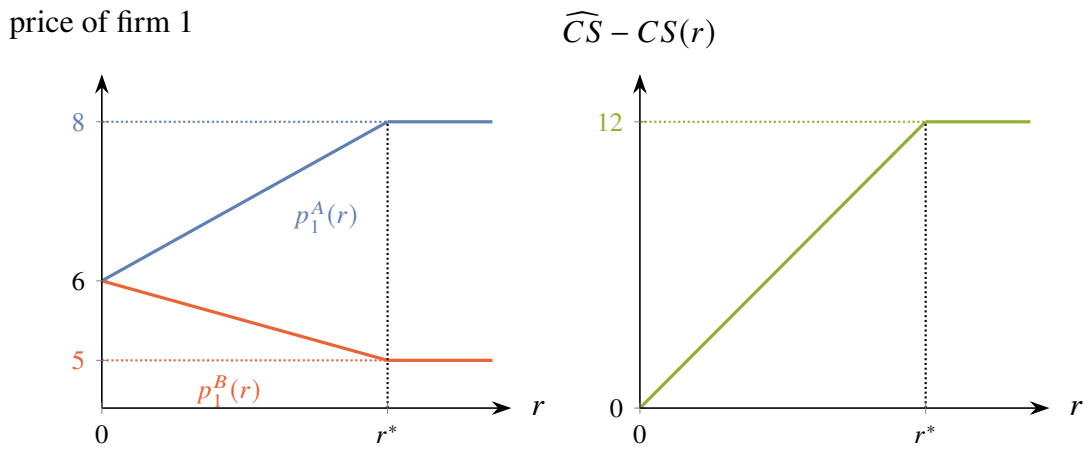
The prices (2.11) satisfy the constraint with equality, and a firm that moves both prices by the same amount leaves the price difference unchanged. The firm's first-order condition along the constraint is the uniform-pricing condition plus a term proportional to $\vartheta\, b_A - (1-\vartheta)\, b_B$, which is zero at $\vartheta = \frac{b_B}{b_A + b_B}$. In the toy model, $\vartheta = \frac{2}{3}$, so firm 1 charges $6 + \frac{2}{3}\, r$ in market A and $6 - \frac{1}{3}\, r$ in market B, and firm 2 charges $7 + \frac{2}{3}\, r$ and $7 - \frac{1}{3}\, r$; at $r = 3$ the prices are the discriminatory prices 8 and 5 (and 9 and 6). The arbitrage cost splits the price movement between the markets in proportion to b_B and b_A: the market where consumers switch more readily, B, absorbs the smaller price cut. The quantity-weighted average price of each firm rises with r, so consumer surplus falls linearly in r, and the consumer surplus gain from uniform pricing relative to the constrained equilibrium is

$$\widehat{CS} - CS(r) = \sum_i \left[\vartheta D_i^A - (1-\vartheta) D_i^B\right] r = (2+2)\, r = 4\, r \quad (r \leq 3),$$

where the bracket equals

$$\frac{2 \cdot 7 - 1 \cdot 8}{3} = 2 \text{ for firm 1} \qquad \text{and} \qquad \frac{2 \cdot 5 - 1 \cdot 4}{3} = 2 \text{ for firm 2.}$$

Figure 2.6 depicts firm 1's prices in panel (a) and the consumer surplus gain in panel (b), against the arbitrage cost r. The prices fan out from the uniform price 6 at $r = 0$ to the discriminatory prices at $r^* = 3$ and stay there. The gain from a given arbitrage cost is linear in r, so any policy that makes arbitrage easier, such as integrating markets, lowers r and raises consumer surplus.



(a) Prices of firm 1 (b) Consumer surplus gain

Figure 2.6: Prices of firm 1 and consumer surplus gain relative to uniform pricing against the arbitrage cost r.

2.8 Sixth Step: Three Policy Implications

Uniform pricing raises consumer surplus at no welfare cost. Consumers lose 12 from discrimination in the toy model, and the firms gain the same amount, so a rule that forces uniform prices moves surplus from firms to consumers. The authors stress that this matters for antitrust authorities, whose objective is typically consumer surplus. Because the market is covered, uniform pricing costs nothing in welfare, in contrast with a monopolist, where the welfare effect of discrimination depends on the curvature of demand (Aguirre et al., 2010).

Firms have a collective interest in segmenting markets. Each firm earns 6 more under discrimination, so discrimination is a Pareto improvement for the firms. Firms may therefore want to segment markets and prevent arbitrage, for example through geo-blocking. The result relies on all firms agreeing on which market is the high-price market, as they do here: the high-price market is A for both firms. The paper's survey of the literature notes that when firms disagree on where the high price belongs, discrimination intensifies competition and firms are in a prisoners' dilemma; the toy model rules this out.

Market integration is a consumer-friendly policy. By Result 4, a lower arbitrage cost raises consumer surplus linearly. Policies that make arbitrage easy, such as the Single Market in the European Union, or that ban geo-blocking, work through this channel. If markets are fully integrated, discrimination is impossible and each firm must charge one price, which is the uniform-pricing equilibrium of the second step.

2.9 Seventh Step: What the Toy Model Keeps and Drops

The toy model keeps every main conclusion of the paper:

1. Price differences between firms are the same under discrimination and uniform pricing, and outputs do not change with the regime (the second step, Proposition 1).
2. The uniform-price elasticity is the weighted arithmetic mean, and the uniform-price Lerner index the weighted harmonic mean, of the market-specific values, so uniform pricing lowers every firm's average Lerner index and average price (the third step, Proposition 2).
3. Consumers gain, firms lose, and the gain can be computed from observables (the fourth step, Corollaries 1 and 2).
4. With arbitrage costs, prices move linearly with r and outputs do not change (the fifth step, Proposition 3).

The toy model also shows two things that the general formulas leave implicit. First, all firms lose the same amount from uniform pricing, and the loss depends on the demand parameters only through the dispersion of $\frac{a_j}{b_j}$ across markets. Second, the strictness of the harmonic-mean inequality does not depend on whether the firms are symmetric: it requires only that the discriminatory prices differ across markets, that is, that $\frac{a_j}{b_j}$ differ.

Relative to Dertwinkel-Kalt and Wey (2023), the toy model has two firms and two markets instead of m firms and n markets, a Hotelling micro-foundation with one line per market, and constant marginal costs. As in the paper, market demand is completely inelastic, so the model isolates the switching channel and cannot say anything about the stay-home channel that drives the welfare effects in monopoly. Products are symmetrically substitutable, demands are linear, firms agree on which market has the higher price, and the equilibria are unique and interior. The paper notes that its results extend to market-specific delivery costs c^j that affect all firms equally, which leave the price differences unchanged, and to loyal consumers who always buy from one firm, which raise a_j but not b_j. Each of these features can be relaxed one at a time.

2.10 Connections to Other Chapters

- Chapter 1 studies the monopoly case, where discrimination also changes total output and the sign of its effect on welfare depends on the curvature of demand.
- Chapter 3 studies another pricing tool of competing firms: two-part tariffs, a simple form of nonlinear pricing.
- Chapter 4 studies competing firms that price discriminate among their past customers, using data that they can also sell to a rival.

References

Aguirre, I., S. Cowan, and J. Vickers (2010). "Monopoly Price Discrimination and Demand Curvature." *American Economic Review*, 100(4), pp. 1601–1615.

Dertwinkel-Kalt, M., and C. Wey (2023). "Third-Degree Price Discrimination in Oligopoly When Markets Are Covered." *Journal of Industrial Economics*, 71(2), pp. 464–490. Working-paper version: DICE Discussion Paper No. 336 (2020), Heinrich Heine University Düsseldorf, https://hdl.handle.net/10419/216730.

Chapter 3

Two-Part Tariffs with Two Firms

Companion note to Armstrong, M., and J. Vickers (2001), "Competitive Price Discrimination," RAND Journal of Economics 32(4): 579–605.

3.1 Introduction

Many firms charge a fixed fee plus a price for each unit. A gym charges a monthly membership plus a fee per class, and an amusement park charges an entry fee plus a price per ride. When two such firms compete for the same customers, what do their tariffs look like? And are consumers better or worse off than when firms can only charge a price per unit?

Armstrong and Vickers (2001) answer these questions by letting each firm choose the surplus it offers to consumers, rather than its prices. The per-unit price then only affects how much each customer buys, so the firm sets it at marginal cost, which makes the surplus as large as possible. The fixed fee splits this surplus and is the firm's source of profit. Under linear pricing, in contrast, a price cut also makes each customer buy more units. This *demand-expansion effect* gives firms an extra reason to cut prices, so competition is fiercer.

In a market where every consumer buys from one of the two firms, the paper shows that two-part tariffs raise both industry profit and total welfare relative to linear pricing. Profit, however, rises by more than welfare, so consumers end up worse off. Consumers therefore prefer linear pricing, even though it is less efficient. Belleflamme and Peitz (2015, §9.3.2) give a textbook treatment of the same comparison.

The toy model places two firms at the ends of a Hotelling line and keeps a general demand curve for each consumer. The first step solves the consumers' problem and the second step finds the equilibrium two-part tariff. The third to fifth steps study linear pricing and compare profit, welfare, and consumer surplus, the sixth step explains why consumers prefer linear pricing, the seventh step works out an example with linear demand, and the eighth step lists what the toy model keeps and drops.

3.2 The Toy Model

We study two firms that compete in *two-part tariffs*. A two-part tariff asks a consumer who buys q units of the good to pay a fixed fee m_i regardless of the quantity, plus a constant per-unit price p_i for each unit consumed, so that the total payment to firm i is

$$T_i(q) = m_i + p_i\, q.$$

This is a simple form of nonlinear pricing: the average price paid per unit,

$$\frac{T_i(q)}{q} = p_i + \frac{m_i}{q},$$

falls as the consumer buys more. Our question is what such a tariff looks like when firms compete, and how the outcome compares with ordinary (linear) pricing in which firms may set only the per-unit price.

The two firms sell horizontally differentiated products and sit at the two ends of a Hotelling line of unit length. Firm 1 is located at 0 and firm 2 at 1. A unit mass of consumers is spread uniformly along the line; a consumer at location $x \in [0, 1]$ incurs a "transportation" (mismatch) cost τ per unit of distance to the firm she buys from, where $\tau > 0$ measures the degree of product differentiation. Both firms produce at the same constant marginal cost c. Figure 3.1 depicts the market.

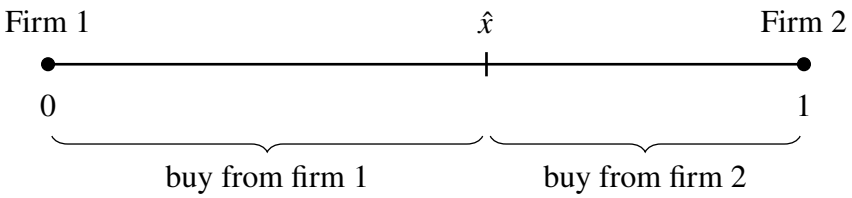


Figure 3.1: The Hotelling line. The consumer at $\hat{x}$ is indifferent between the two firms; consumers to her left buy from firm 1 and those to her right from firm 2.

The game has two stages, which we solve by backward induction:

Stage 1. Firms simultaneously choose their two-part tariffs (p_i, m_i), $i = 1, 2$.

Stage 2. Each consumer observes both tariffs, chooses the firm to patronize, and then decides how many units to buy.

We start with Stage 2 (the consumers' problem) in the first step and then move back to Stage 1 (the firms' problem) in the second step.

3.3 First Step: The Consumers in Stage 2

Consider a consumer who has decided to buy from a firm charging per-unit price p_i. She chooses her quantity q to maximize the utility she gets from consuming the good, $u(q)$, net of the per-unit

payment $p_i q$. Her *indirect utility* is the value of this problem,

$$v(p_i) \equiv \max_{q} \{\, u(q) - p_i\, q \,\}, \tag{3.1}$$

where $u(q)$ is increasing and concave. The first-order condition of (3.1) is

$$u'(q) = p_i,$$

which defines the consumer's *Walrasian demand* $q(p_i)$: she keeps buying units as long as the marginal utility of the good exceeds its per-unit price, and stops where the two are equal.

A fact we use repeatedly is Roy's identity: the derivative of the indirect utility with respect to the per-unit price equals minus the quantity demanded,

$$v'(p_i) = -\,q(p_i). \tag{3.2}$$

This follows from the envelope theorem applied to (3.1): because q is chosen optimally, a marginal change in p_i affects v only through the direct term $-p_i q$, whose derivative is $-q(p_i)$. Intuitively, a one-dollar rise in the per-unit price costs the consumer one dollar on each of the $q(p_i)$ units she was already buying.

Bringing back the fixed fee m_i, the net surplus a consumer obtains from firm i's tariff is her indirect utility from the variable good minus the fixed fee,

$$w(p_i, m_i) \;=\; \underbrace{v(p_i)}_{\text{surplus from } q(p_i) \text{ units}} \;-\; \underbrace{m_i}_{\text{fixed fee}}.$$

To keep the notation compact, write $w_1 \equiv w(p_1, m_1)$ and $w_2 \equiv w(p_2, m_2)$ for the two "deals" on offer.

A consumer located at x therefore enjoys total utility

$$\begin{aligned} r - \tau x + w_1 \qquad & \text{if she buys from firm 1,} \\ r - \tau(1-x) + w_2 \qquad & \text{if she buys from firm 2,} \end{aligned}$$

where r is the (large) gross value of consuming the product category at all. We assume the market is *covered*, so every consumer buys from one of the two firms; the constant r is then common to both options and drops out of every comparison below.

The indifferent consumer. The consumer $\hat{x}$ who is just indifferent between the two firms solves

$$r - \tau\hat{x} + w_1 \;=\; r - \tau(1-\hat{x}) + w_2.$$

Cancelling r and solving for $\hat{x}$,

$$\begin{aligned} w_1 - w_2 + \tau &= 2\tau\,\hat{x} \\ \hat{x} &= \frac{\tau + w_1 - w_2}{2\tau}. \end{aligned} \tag{3.3}$$

Everyone to the left of $\hat{x}$ prefers firm 1 and everyone to the right prefers firm 2. Because a unit mass of consumers is spread uniformly on the line, firm 1's demand is $\hat{x}$ and firm 2's demand is $1 - \hat{x}$:

$$D_1 = \hat{x} = \frac{\tau + w_1 - w_2}{2\tau}, \qquad D_2 = 1 - \hat{x} = \frac{\tau + w_2 - w_1}{2\tau}.$$

A firm attracts more consumers by offering a better deal (a larger w_i), whether it improves that deal through a lower per-unit price or a lower fixed fee.

3.4 Second Step: The Firms' Two-Part Tariffs in Stage 1

Take firm 1. Each of its $\hat{x}$ customers buys $q(p_1)$ units, on which the firm earns the per-unit margin $p_1 - c$, and each pays the fixed fee m_1. Firm 1 therefore solves

$$\max_{p_1, m_1} \quad \pi_1 \;=\; \underbrace{\hat{x}}_{\text{mass of buyers}} \cdot \Big[\underbrace{(p_1 - c)\, q(p_1)}_{\text{variable margin}} + \underbrace{m_1}_{\text{fixed fee}} \Big]. \tag{3.4}$$

It is cleaner to let the firm choose the *deal* w_1 it offers rather than the fee m_1 directly. Since $w_1 = v(p_1) - m_1$, we have $m_1 = v(p_1) - w_1$; substituting this and (3.3) into (3.4) turns the problem into a choice of the per-unit price p_1 and the deal w_1,

$$\max_{p_1, w_1} \quad \pi_1 \;=\; \underbrace{\frac{\tau + w_1 - w_2}{2\tau}}_{\hat{x}} \cdot \Big[(p_1 - c)\, q(p_1) + \underbrace{v(p_1) - w_1}_{m_1} \Big]. \tag{3.5}$$

This change of variables is what makes the toy model tractable, because the market share $\hat{x}$ now depends on w_1 but *not* on the per-unit price p_1. The two instruments split their jobs: p_1 shapes the size of the surplus firm 1 creates with each of its customers, while w_1 determines how many customers it wins.

First-order condition in p_1. Because $\hat{x}$ does not depend on p_1, differentiating (3.5) with respect to p_1 and dividing by $\hat{x} > 0$ leaves only the bracketed term,

$$q(p_1) + q'(p_1)\,(p_1 - c) + v'(p_1) = 0.$$

Now use Roy's identity (3.2), $v'(p_1) = -q(p_1)$. The first and last terms cancel,

$$\cancel{q(p_1)} + q'(p_1)\,(p_1 - c) \underbrace{-\cancel{q(p_1)}}_{v'(p_1)} = 0 \quad \Longrightarrow \quad \underbrace{q'(p_1)}_{<0 \text{ (law of demand)}} \quad (p_1 - c) = 0.$$

Since $q'(p_1) < 0$, the only way this can hold is $p_1 - c = 0$, that is,

$$p_1^* = c \qquad \text{(and, by the same argument, } p_2^* = c\text{)}. \tag{3.6}$$

Each firm therefore sets its per-unit price equal to marginal cost, so consumption is efficient. The reason is that the per-unit price does not help firm 1 win customers (that is the job of w_1); it only sets how much surplus the firm generates with each customer it already has. That surplus, variable profit plus the consumer's own gross surplus, is maximized when the good is priced at marginal cost, and the firm then collects its share of it through the fixed fee. Marginal-cost pricing makes the pie as large as possible; the fee decides how it is split.

First-order condition in w_1. Differentiating (3.5) with respect to w_1 gives two terms, one from the market share $\hat{x}$ (which rises at rate $\frac{1}{2\tau}$) and one from the fee inside the bracket (which falls one-for-one with w_1):

$$\frac{1}{2\tau}\Big[(p_1 - c)\, q(p_1) + v(p_1) - w_1\Big] - \frac{\tau + w_1 - w_2}{2\tau} = 0.$$

Multiplying through by 2τ (and recognizing $2\tau\,\hat{x} = \tau + w_1 - w_2$),

$$\Big[(p_1 - c)\,q(p_1) + v(p_1) - w_1\Big] - \big(\tau + w_1 - w_2\big) = 0. \tag{3.7}$$

Impose the optimal per-unit price $p_1 = c$ from (3.6), so the variable margin $(c - c)\,q(p_1) = 0$ vanishes, and invoke symmetry, $w_1 = w_2$, so that $w_1 - w_2 = 0$:

$$\Big[\underbrace{(c - c)\,q(p_1)}_{0} + v(p_1) - w_1\Big] - \big(\tau + \underbrace{w_1 - w_1}_{0}\big) = 0 \quad \Longrightarrow \quad v(p_1) - w_1 - \tau = 0.$$

Finally recall that $w_1 = v(p_1) - m_1$, i.e. $v(p_1) - w_1 = m_1$, so the last equality reads $m_1 - \tau = 0$, i.e.

$$m_1^* = \tau \qquad \text{(and, by the same argument, } m_2^* = \tau\text{)}. \tag{3.8}$$

Putting (3.6) and (3.8) together, each firm charges the equilibrium two-part tariff

$$T(q) = \underbrace{\tau}_{\text{fixed fee}} + \underbrace{c}_{\text{price per unit}}\, q.$$

Competition drives the per-unit price down to marginal cost, and the firms make their money entirely on the fixed fee, which they can raise up to the value of the differentiation parameter τ. The fee is the reward for being differentiated: the more differentiated the two products (the larger τ), the softer the competition and the higher the fee each firm can sustain.

Equilibrium profits. At the symmetric equilibrium each firm serves half the market, $\hat{x} = \frac{1}{2}$, earns a zero variable margin, and collects the fee $m^* = \tau$ from each customer:

$$\pi_1 = \underbrace{\frac{\tau + w_1 - w_2}{2\tau}}_{=\frac{1}{2}} \Big[\underbrace{(p_1 - c)\,q_1}_{0} + \underbrace{m_1}_{\tau}\Big] = \frac{\tau}{2}. \tag{3.9}$$

Adding the two firms, industry profit under the two-part (nonlinear) tariff is

$$\pi_{NL} = \tau.$$

We keep this number in mind: it is the benchmark against which linear pricing is measured.

3.5 Third Step: Linear Pricing and Industry Profit

Consider instead a setting of ordinary *linear pricing*, in which firms may charge only a per-unit price, with no fixed fee ($m_i = 0$). The consumer's deal is simply $w_i = v(p_i)$, and firm i's profit per customer is the variable margin $\pi_i(p_i) \equiv (p_i - c)\,q(p_i)$. We want to compare the two regimes for a general demand curve $q(\cdot)$, so a small change of variables helps.

Let firm j charge the candidate symmetric price p, and write the price p_i of its rival, firm i, as a point between marginal cost and p,

$$p_i = c + k\,(p - c), \qquad k \in [0, 1].$$

The number k measures how far the price sits above marginal cost: $k = 0$ means $p_i = c$ and $k = 1$ means $p_i = p$. Define the consumer's surplus as a function of this index,

$$s(k) \equiv v\big(c+k(p-c)\big), \qquad \text{so that} \qquad \begin{cases} s(0) = v(c) & \text{“low price” (all surplus to the consumer),} \\ s(1) = v(p) & \text{“high price.”} \end{cases} \tag{3.10}$$

Differentiating s and using Roy's identity (3.2),

$$s_k(k) \equiv \frac{\partial s(k)}{\partial k} = (p - c)\, v'\big(c + k(p - c)\big) = -(p - c)\, q\big(c + k(p - c)\big) \;< 0. \tag{3.11}$$

The consumer is worse off as the price rises, so $s_k < 0$.

The point of this parametrization is that the firm's per-customer profit can be written in terms of s_k. At price $p_i = c + k(p - c)$,

$$\pi_i\big(c+k(p-c)\big) = \underbrace{\big[c + k(p - c) - c\big]}_{=k(p-c)}\, q\big(c+k(p-c)\big) = k(p-c)\, q\big(c+k(p-c)\big) = -\,k\, s_k(k), \tag{3.12}$$

where the last equality just reads off (3.11). In particular, at $k = 1$ the per-customer (equivalently, unit-mass industry) profit under price p is

$$\pi_L \equiv \pi(p) = -\,s_k(1). \tag{3.13}$$

The firm's problem. Facing the rival's price p, firm i chooses its own price, i.e. chooses k, to solve

$$\max_k \quad \pi_i = \frac{\tau + s(k) - s(1)}{2\tau} \cdot \big(-\,k\, s_k(k)\big), \tag{3.14}$$

where we used $v(p_i) = s(k)$ and $v(p) = s(1)$ in the market share (3.3), and the per-customer profit (3.12). At a symmetric equilibrium the rival's price is a best response to itself, so $p_i = p$, i.e. $k = 1$.

First-order condition in k. Differentiating the product in (3.14) and setting it to zero,

$$-\,k\,\big(s_k(k)\big)^2 - \big[\tau + s(k) - s(1)\big]\big[s_k(k) + k\, s_{kk}(k)\big] = 0. \tag{3.15}$$

Evaluate this at the symmetric equilibrium $k = 1$, where $s(1) - s(1) = 0$ and the first bracket collapses to τ,

$$\begin{aligned} -\big[s_k(1)\big]^2 - \tau\big[s_k(1) + s_{kk}(1)\big] &= 0 \\ \Longrightarrow \quad -\,s_k(1) &= s_{kk}(1) + \frac{1}{\tau}\big[s_k(1)\big]^2. \end{aligned} \tag{3.16}$$

Using $\pi_L = -s_k(1)$ from (3.13), and hence $[s_k(1)]^2 = \pi_L^2$, equation (3.16) says that industry profit under linear pricing satisfies

$$\pi_L = s_{kk}(1) + \frac{1}{\tau}\,\pi_L^2. \tag{3.17}$$

The sign of s_{kk}. Differentiate (3.11) once more,

$$s_{kk}(k) = -(p-c)\,q'\big(c + k(p-c)\big)\cdot(p-c) = -\underbrace{(p-c)^2}_{>0}\,\underbrace{q'\big(c+k(p-c)\big)}_{<0 \text{ (law of demand)}} > 0. \tag{3.18}$$

So $s_{kk}(1) > 0$: the consumer's surplus falls at an increasing rate as the price climbs. Feeding this into (3.17),

$$\pi_L = \underbrace{s_{kk}(1)}_{>0} + \frac{1}{\tau}\,\pi_L^2 \;>\; \frac{1}{\tau}\,\pi_L^2 \quad\Longrightarrow\quad \pi_L > \frac{1}{\tau}\,\pi_L^2 \quad\Longrightarrow\quad \tau > \pi_L,$$

where the last step divides by $\pi_L > 0$. Since industry profit under the two-part tariff is $\pi_{NL} = \tau$, we have established the first comparison,

$$\pi_{NL} = \tau \;>\; \pi_L. \tag{3.19}$$

Result 1 (industry profit). Industry profit is higher under two-part tariffs than under linear pricing, $\pi_{NL} = \tau > \pi_L$.

Figure 3.2 depicts industry profit under the two regimes as a function of τ, with the linear demand $q(p) = a - p$ of the seventh step, $a = 2$, and $c = 1$. Under linear pricing, the symmetric first-order condition is

$$\tau\,(a + c - 2p) = (a-p)^2\,(p-c),$$

which gives one markup $p - c$ for each τ; the dotted lines mark $\tau = \frac{9}{32}$, the value that the example of the seventh step uses.

Intuitively, even though competition forces the per-unit price down to marginal cost, the fixed fee lets firms recover more than they could through a single linear price, so nonlinear pricing softens competition. A larger τ softens competition in both regimes. Under two-part tariffs the fee rises one-for-one with τ, while under linear pricing the price rises with τ but each price increase also lowers the units every customer buys, so industry profit grows more slowly than τ and stays below it.

3.6 Fourth Step: Welfare

Turn to total welfare, the sum of producer surplus and consumer surplus. At price $p_i = c + k(p-c)$, welfare per consumer is the firm's variable profit plus the consumer's own surplus (the fixed fee is a pure transfer and cancels),

$$W(k) = \underbrace{\pi\big(c+k(p-c)\big)}_{=-k\,s_k(k)\ \text{(PS)}} + \underbrace{v\big(c+k(p-c)\big)}_{=s(k)\ \text{(CS)}} = -\,k\,s_k(k) + s(k). \tag{3.20}$$

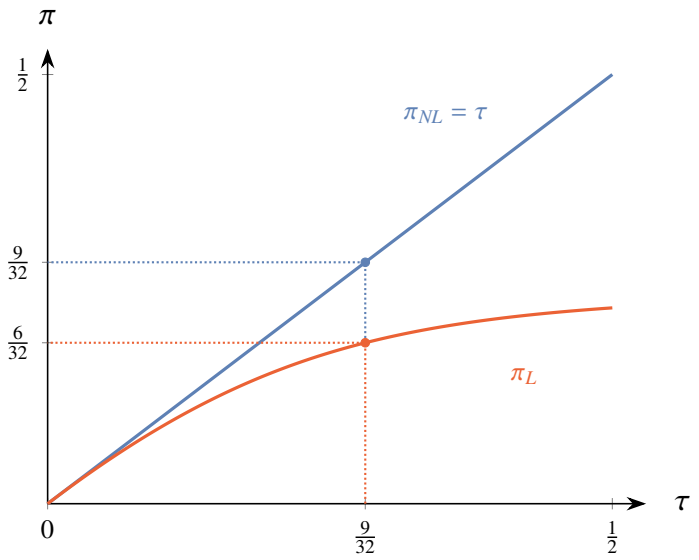


Figure 3.2: Industry profit under two-part tariffs (π_{NL}) and linear pricing (π_L) against τ, for $q(p) = a - p$, $a = 2$, and $c = 1$.

Evaluate W in the two regimes.

- *Under the two-part tariff (NL):* the per-unit price is c, so $k = 0$, and

$$W(0) = \underbrace{-0 \cdot s_k(0)}_{0} + s(0) = s(0) = v(c) \equiv W_{NL}.$$

  Welfare is maximized here, since pricing at marginal cost yields the efficient quantity.

- *Under linear pricing (L):* the price is p, so $k = 1$, and

$$W(1) = -1 \cdot s_k(1) + s(1) = s(1) - s_k(1) \equiv W_L.$$

To rank the two, note from (3.20) that

$$W'(k) = -s_k(k) - k\, s_{kk}(k) + s_k(k) = -k\, s_{kk}(k),$$

so $W'(0) = 0$ and $W'(k) < 0$ for $k > 0$ (because $s_{kk} > 0$): welfare peaks at $k = 0$ and falls thereafter, so $W_{NL} = W(0) > W(1) = W_L$.

Result 2 (welfare). Welfare is higher under two-part tariffs than under linear pricing, $W_{NL} > W_L$.

Figure 3.3 depicts the welfare curve $W(k)$, assuming, as is standard, that welfare is also concave in the price (hence in k).

Intuitively, two-part tariffs price the marginal unit at cost and so avoid the quantity distortion that a linear price above cost creates.

How large is the welfare gap? Concavity lets us bound $\Delta W \equiv W_{NL} - W_L$. For a concave function the tangent at $k = 1$ lies above the curve, so evaluating that tangent back at $k = 0$ gives

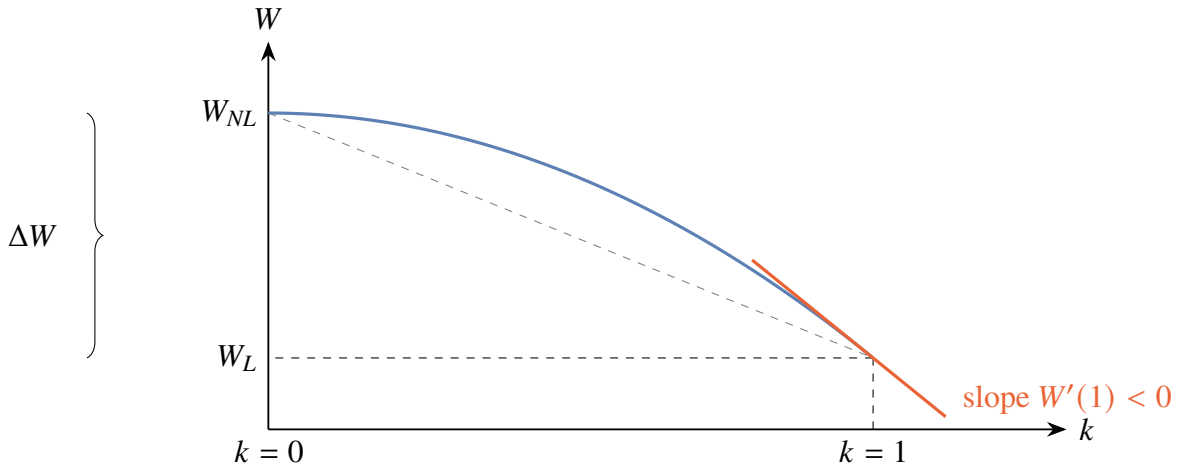


Figure 3.3: Welfare $W(k)$ is concave and decreasing in k. It is maximized at $k = 0$ (the two-part tariff, marginal-cost pricing) and falls to its linear-pricing value at $k = 1$. The gap is $\Delta W = W_{NL} - W_L$. The figure is schematic: only the shape of the curve and the ordering $W_{NL} > W_L$ matter.

$W(0) \leq W(1) - W'(1)$, i.e.

$$\Delta W = W(0) - W(1) \;\leq\; -W'(1) = s_{kk}(1), \tag{3.21}$$

using $W'(1) = -s_{kk}(1)$. Now rewrite $s_{kk}(1)$ with the equilibrium condition (3.17),

$$s_{kk}(1) = \pi_L - \frac{1}{\tau}\pi_L^2 = \frac{\pi_L}{\tau}\big(\tau - \pi_L\big). \tag{3.22}$$

Because $\pi_L < \tau$ from (3.19), the factor $\frac{\pi_L}{\tau} < 1$, while $\tau - \pi_L > 0$. Hence

$$s_{kk}(1) = \underbrace{\frac{\pi_L}{\tau}}_{<1}\,\underbrace{\big(\tau - \pi_L\big)}_{>0} < \tau - \pi_L.$$

Chaining this with the bound (3.21) and recalling $\tau - \pi_L = \pi_{NL} - \pi_L$,

$$\Delta W \;\leq\; s_{kk}(1) \;<\; \tau - \pi_L \;=\; \pi_{NL} - \pi_L \;\equiv\; \Delta\pi. \tag{3.23}$$

Result 3 (welfare gain and profit gain). The welfare gain from two-part tariffs is smaller than the profit gain, $\Delta W < \Delta\pi$.

Intuitively, firms gain more than society as a whole when moving from linear to two-part pricing. That extra gain has to come from somewhere, and the fifth step shows that it comes out of consumers' pockets.

3.7 Fifth Step: Consumer Surplus and the Overall Comparison

Total welfare is consumer surplus plus industry profit, $W = CS + \pi$, so $CS = W - \pi$. Rearranging the welfare-versus-profit comparison (3.23),

$$\begin{aligned} & W_{NL} - W_L < \pi_{NL} - \pi_L \\ \iff \quad & \underbrace{W_{NL} - \pi_{NL}}_{CS_{NL}} < \underbrace{W_L - \pi_L}_{CS_L} \\ \iff \quad & CS_{NL} < CS_L. \end{aligned}$$

Result 4 (consumer surplus). Consumers are worse off under two-part tariffs than under linear pricing, $CS_{NL} < CS_L$, even though total welfare is higher.

Figure 3.4 depicts the decomposition of welfare into industry profit and consumer surplus in the two regimes. Collecting the three comparisons, moving from linear pricing to two-part tariffs:

- total welfare rises;
- industry profit rises, by more than total welfare;
- consumer surplus falls.

In symbols,

$$W_{NL} > W_L, \qquad \pi_{NL} > \pi_L, \qquad \text{but} \qquad CS_{NL} < CS_L.$$

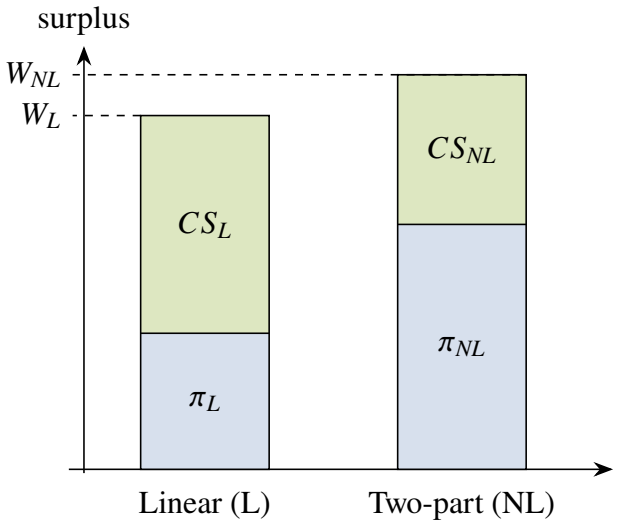


Figure 3.4: Decomposition of welfare into industry profit and consumer surplus. Moving from linear (L) to two-part (NL) pricing raises total welfare ($W_{NL} > W_L$) and profit ($\pi_{NL} > \pi_L$), but lowers consumer surplus ($CS_{NL} < CS_L$). The figure is schematic: only the orderings of the bars matter, not their heights.

Intuitively, two-part tariffs enlarge the pie *and* hand firms a larger slice, so much larger that consumers are left with less than they had under linear pricing.

3.8 Sixth Step: Why Consumers Prefer Linear Pricing

The comparisons above are driven by how hard firms compete under each pricing regime, which in turn depends on what a firm gains by giving consumers a better deal.

Under *two-part tariffs*, the per-unit price is already at marginal cost, so the only way firm i improves its deal is by cutting the fixed fee. A one-dollar cut in the fee has two effects:

(+) *Business stealing.* A better deal pulls marginal consumers away from the rival, raising the firm's market share $\hat{x}$.

(−) *Lost revenue.* The firm collects one dollar less from every customer it already had.

Under *linear pricing*, the firm improves its deal by cutting the per-unit price, which has the same two effects *and* a third one:

(+) *Demand expansion.* A lower price makes each of the firm's own customers buy more units, and the firm earns its margin on the extra volume.

This third effect gives firms a stronger reason to cut prices under linear pricing than to cut fees under two-part tariffs. Competition is therefore fiercer under linear pricing: prices are driven closer to cost, and more surplus is left in consumers' hands. Under two-part tariffs the marginal price is held at cost and the fixed fee moves no quantity, so this extra pro-competitive force is absent, firms keep fees high, and consumers end up with less, which is the ranking $CS_{NL} < CS_L$ found above. The efficiency of two-part tariffs, which price each unit at cost, is itself what removes the demand-expansion channel that would otherwise have disciplined the firms.

3.9 Seventh Step: A Numerical Example with Linear Demand

To put numbers on the four results, take the quadratic gross utility $u(q) = a\,q - \frac{q^2}{2}$, whose marginal utility is $u'(q) = a - q$. Setting $u'(q) = p$ gives the *linear demand*

$$q(p) = a - p \qquad (\text{for } p \le a),$$

and, evaluating $u(q) - p\,q$ at $q = a - p$, the indirect utility

$$v(p) = \max_q \left\{ a\,q - \frac{q^2}{2} - p\,q \right\} = \frac{1}{2}\,(a-p)^2.$$

As a check, $v'(p) = -(a-p) = -q(p)$, so Roy's identity (3.2) holds. Take $a = 2$, $c = 1$, and $\tau = \frac{9}{32}$.

Two-part tariff (NL). From the second step, the per-unit price is $p^* = c = 1$ and the fee is $m^* = \tau = \frac{9}{32}$, so each firm posts $T(q) = \frac{9}{32} + q$. Every consumer buys $q(c) = a - c = 1$ unit. Industry profit is the fee collected from the unit mass of consumers, total welfare is the efficient surplus, and consumer surplus is that surplus net of the fee:

$$\pi_{NL} = \tau = \frac{9}{32}, \qquad W_{NL} = v(c) = \frac{1}{2}\,(a-c)^2 = \frac{1}{2}, \qquad CS_{NL} = v(c) - \tau = \frac{1}{2} - \frac{9}{32} = \frac{7}{32}.$$

Linear pricing (L). With no fee, the symmetric first-order condition from the third step is

$$\tau\,(a+c-2p) = (a-p)^2\,(p-c).$$

Writing the markup as $z \equiv p-c$ and inserting $a=2$, $c=1$, $\tau=\frac{9}{32}$,

$$\frac{9}{32}\,(1-2z) = (1-z)^2\,z \quad\Longleftrightarrow\quad 32z^3-64z^2+50z-9=0 \quad\Longleftrightarrow\quad (4z-1)\big(8z^2-14z+9\big)=0.$$

The quadratic factor has no real root (its discriminant is $14^2 - 4\cdot 8\cdot 9 = -92 < 0$), so the unique symmetric equilibrium markup is $z=\frac{1}{4}$, i.e. $p_L = c+\frac{1}{4} = \frac{5}{4}$. Each consumer then buys $q_L = a - p_L = \frac{3}{4}$, and

$$\pi_L = (p_L-c)\,q_L = \frac{1}{4}\cdot\frac{3}{4} = \frac{6}{32}, \qquad CS_L = v(p_L) = \frac{1}{2}\Big(\frac{3}{4}\Big)^2 = \frac{9}{32}, \qquad W_L = \pi_L + CS_L = \frac{15}{32}.$$

Comparison. Table 3.1 reports the two regimes.

Outcome	Two-part tariff (NL)	Linear pricing (L)
Per-unit price p	1	$\frac{5}{4}$
Units per consumer q	1	$\frac{3}{4}$
Industry profit π	$\frac{9}{32} \approx 0.281$	$\frac{6}{32} \approx 0.188$
Consumer surplus CS	$\frac{7}{32} \approx 0.219$	$\frac{9}{32} \approx 0.281$
Total welfare W	$\frac{16}{32} = 0.500$	$\frac{15}{32} \approx 0.469$

Table 3.1: Equilibrium outcomes under two-part and linear pricing, for $a=2$, $c=1$, and $\tau=\frac{9}{32}$. The profit, consumer-surplus, and welfare rows are written over the common denominator 32 so the numerators can be compared directly.

All four results hold:

$$\pi_{NL} = \frac{9}{32} > \frac{6}{32} = \pi_L, \qquad W_{NL} = \frac{16}{32} > \frac{15}{32} = W_L, \qquad CS_{NL} = \frac{7}{32} < \frac{9}{32} = CS_L,$$

and the welfare gain $\Delta W = \frac{1}{32}$ is smaller than the profit gain $\Delta\pi = \frac{3}{32}$.

Moving to the two-part tariff pushes the per-unit price down from $\frac{5}{4}$ to marginal cost 1, so each consumer's purchase rises from $\frac{3}{4}$ to a full unit: consumption becomes efficient and total surplus grows by $\frac{1}{32}$. The fixed fee, however, lets firms collect $\frac{3}{32}$ more than under linear pricing, which is $\frac{2}{32}$ more than the entire welfare gain. That extra $\frac{2}{32}$ equals the fall in consumer surplus ($\frac{9}{32} - \frac{7}{32}$). Consumption becomes efficient, industry profit rises, and consumers are left worse off.

3.10 Eighth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The change of variables in which each firm chooses the surplus it offers to consumers instead of its fixed fee, so that the market share depends on the offered surplus and not on the per-unit price.
- The equilibrium two-part tariff: the per-unit price equals marginal cost and the fixed fee equals the differentiation parameter.
- The ranking of the two regimes for a general demand curve: two-part tariffs raise industry profit and total welfare, the welfare gain is smaller than the profit gain, and consumer surplus falls.
- The demand-expansion effect, which makes competition fiercer under linear pricing than under two-part tariffs.

What it drops.

- Consumer heterogeneity in the quantity demanded. All consumers share the same demand and differ only in their location; when consumers also differ in their demand and firms cannot observe it, the tariff must screen consumers.
- Price discrimination across observable groups of consumers.
- Markets that are not fully covered: every consumer buys from one of the two firms, so the fixed fee never drives a consumer out of the market.
- Asymmetric firms, with different costs or locations, and more than two firms.

3.11 Connections to Other Chapters

- Chapter 2 studies third-degree price discrimination by competing firms when every consumer buys from some firm.
- Chapter 1 studies a monopolist that charges different prices in different markets, and shows when this discrimination raises welfare.
- Chapter 8 compares wholesale prices with two-part tariffs as contracts between an input supplier and downstream firms, and shows that wholesale prices attract more entrants.

References

Armstrong, M., and Vickers, J. (2001). Competitive price discrimination. *RAND Journal of Economics*, 32(4), 579–605. `https://www.jstor.org/stable/2696383`

Belleflamme, P., and Peitz, M. (2015). *Industrial Organization: Markets and Strategies*, 2nd ed. Cambridge: Cambridge University Press.

Chapter 4

Customer Information Sharing: When a Small Firm Sells Its Data

Companion note to Liu, Q., and K. Serfes (2006), "Customer Information Sharing among Rival Firms," European Economic Review 50(6): 1571–1600.

4.1 Introduction

Consider a large supermarket and a small grocery store in the same town. Each store learns from past purchases which customers like it, and it can use these data to tailor its prices. It can also sell the data, even to its rival. When do rivals trade customer data, and what happens if a regulator bans the trade? Liu and Serfes (2006) answer these questions.

Two firms compete over two periods. In the first period they charge uniform prices; in the second, each firm knows the preferences of its own past customers and can offer each of them a personal price. A firm with a much larger loyal base lets some of its loyal consumers buy from its small rival in the first period, because it expects to buy their names later. The small rival then sells these names, and the large firm uses them to win the consumers back. Anticipating the sale, the small firm cuts its first-period price to collect more names to sell.

Banning the trade makes both firms worse off and consumers better off. Without the option to buy names back, the large firm fights harder for its loyal consumers in the first period, and prices fall further. The ban raises total welfare only when the firms are highly asymmetric and patient. In that case, data sharing leads many consumers to buy from the firm they like less in the first period, and the ban reduces this mismatch by more than the mismatch it creates in the second period.

The toy model has one parameter for the asymmetry between the firms' customer bases. The first step solves the second period without data sharing, the second step finds who sells data to whom, and the third step solves the first period. The fourth step bans the trade, the fifth step compares total welfare, the sixth step lets firms and consumers discount the future, and the seventh step lists what the toy model keeps and drops.

4.2 The Toy Model

Consumers. Two firms, A and B, produce at zero marginal cost. Consumers are indexed by their loyalty ℓ, uniformly distributed on $[-1, L]$ with density one, where $L > 1$. In each period, a consumer buys one unit, from firm A at price p_A or from firm B at price p_B, and obtains utility

$$V + \ell - p_A \quad \text{from firm } A, \qquad V - p_B \quad \text{from firm } B, \tag{4.1}$$

where V is large enough that every consumer buys. Consumers with $\ell > 0$ prefer firm A and consumers with $\ell < 0$ prefer firm B. At equal prices, firm A sells to a mass L of consumers and firm B to a mass 1, so a larger L means a larger asymmetry between the firms' customer bases. With $L = 1$ the model is the standard Hotelling line.

Information. In period 1, firms know only the distribution of ℓ and set uniform prices p_A^1 and p_B^1. Call ℓ^* the period-1 marginal consumer: those with $\ell \geq \ell^*$ buy from A and the others from B. At the end of period 1, each firm learns the ℓ of each of its own customers, its *database*, and knows nothing about the rival's customers beyond the fact that they bought from the rival.

Timing in period 2. First, each firm decides whether to sell its database to the rival; a sale occurs only if it raises the two firms' joint profit, and when firm B sells, it obtains a share $\sigma \in [0, 1]$ of the gain and firm A the rest. Second, each firm sends a *blanket coupon*, one price, to the consumers in the rival's database. Third, each firm sends *targeted coupons*, one price per consumer, to the consumers whose ℓ it knows. Firms and consumers weigh both periods equally ($\delta = 1$); the sixth step lets them discount period 2.

4.3 First Step: Period 2 without Data Sharing

Consider that $0 < \ell^* < L$, so firm A's database is $[\ell^*, L]$ and firm B's database is $[-1, \ell^*)$.

Firm A's database. These consumers all prefer firm A. For any blanket coupon that firm B sends, firm A undercuts it consumer by consumer, so firm B sends a price of zero and firm A charges each consumer ℓ. Firm A earns

$$\int_{\ell^*}^{L} \ell \, d\ell = \frac{L^2 - (\ell^*)^2}{2}.$$

Firm B's database. Firm A sends a blanket coupon p, and firm B answers with the targeted price $p - \ell$, as long as it is positive. Consumers with $\ell > p$ buy from A, so firm A earns $p(\ell^* - p)$, which is maximized at $p = \frac{\ell^*}{2}$. Firm A earns $\frac{(\ell^*)^2}{4}$, and firm B earns

$$\pi_B^{NE} = \int_{-1}^{\ell^*/2} \left(\frac{\ell^*}{2} - \ell \right) d\ell = \frac{(\ell^* + 2)^2}{8}. \tag{4.2}$$

Hence, consumers in $\left(0, \frac{\ell^*}{2}\right)$ prefer firm A but stay with firm B, while those in $\left(\frac{\ell^*}{2}, \ell^*\right)$ switch to firm A. Adding up, the two firms' joint profit is

$$\Pi^{NE} = \frac{4L^2 - (\ell^*)^2 + 4\ell^* + 4}{8}. \tag{4.3}$$

4.4 Second Step: Who Sells Data to Whom

Firm A sells to firm B. Nothing changes: firm B already charges zero to firm A's customers, and knowing their ℓ does not let it do better. There are no gains from trade.

Firm B sells to firm A. Firm A now knows every consumer. Each consumer buys from the firm she prefers: firm A charges ℓ to consumers with $\ell \geq 0$, and firm B charges $-\ell$ to consumers with $\ell < 0$, with the rival pricing at zero. Joint profit becomes $\frac{L^2+1}{2}$, so the gains from trade are

$$GFT = \frac{L^2+1}{2} - \Pi^{NE} = \frac{\ell^*(\ell^*-4)}{8}. \tag{4.4}$$

Result 1 (who sells data, Proposition 1 of the paper). In period 2:

- firm A, the firm with the larger customer base, never sells its database;
- firm B sells its database to firm A if and only if $\ell^* > 4$;
- with $L = 1$, $\ell^* < 4$ for any first-period prices, so no data are traded.

Figure 4.1 depicts the two parts of (4.4). When firm A buys the data, firm B loses $\frac{\ell^*}{2}$ on its own loyal consumers in $[-1, 0)$, who now receive a price of zero from firm A instead of the blanket coupon $\frac{\ell^*}{2}$. Firm A gains $\frac{(\ell^*)^2}{8}$ on the consumers in $\left(\frac{\ell^*}{2}, \ell^*\right)$, whom it now charges ℓ instead of $\frac{\ell^*}{2}$. In $\left(0, \frac{\ell^*}{2}\right)$ the two firms' changes cancel. The gain exceeds the loss when $\ell^* > 4$; at $\ell^* = 10$, for instance, the gains from trade are $\frac{15}{2}$.

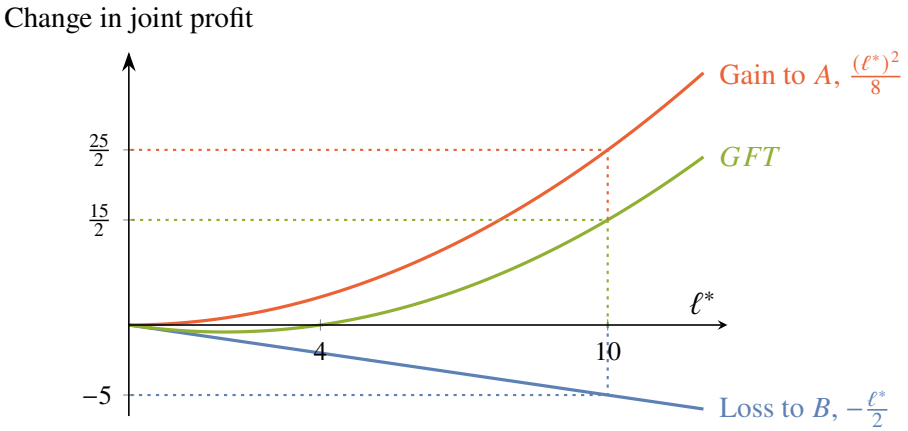


Figure 4.1: Gains from trade when firm B sells its database to firm A (green), and their two parts: the gain on consumers who prefer firm A (red) and the loss on firm B's loyal consumers (blue).

Intuitively, data are worth trading only if firm B's database contains many consumers who prefer firm A. Firm A values these names because it can extract their surplus, while firm B loses because firm A then competes harder for B's own loyal consumers. The first effect wins only when firm A has given up a large group of its loyal consumers in period 1.

4.5 Third Step: Period 1 When Data Can Be Sold

Consumers. If $\ell^* > 4$, firm B sells its data, and in period 2 firm A charges ℓ^* to the marginal consumer whichever firm she bought from in period 1. Her period-2 utility does not depend on her period-1 choice, so

$$\ell^* = p_A^1 - p_B^1. \tag{4.5}$$

Data price. Firm B's profit in period 2 is π_B^{NE} without the sale and $\frac{1}{2}$ with it. Firm A pays the information price

$$IP = \left[\frac{(\ell^*+2)^2}{8} - \frac{1}{2}\right] + \sigma\,\frac{\ell^*(\ell^*-4)}{8}, \tag{4.6}$$

which compensates firm B for its lost profit and gives it a share σ of the gains from trade. The firms' profits over the two periods are

$$\Pi_A = p_A^1(L-\ell^*) + \frac{L^2}{2} - IP, \qquad \Pi_B = p_B^1(\ell^*+1) + \frac{1}{2} + IP. \tag{4.7}$$

Result 2 (equilibrium with data sharing, Proposition 2 of the paper). If

$$L > \bar{L} \equiv 13 + 12\sqrt{\frac{5}{7-\sigma}},$$

the unique subgame perfect equilibrium has:

- first-period prices

  $$p_A^1 = \frac{L(7-\sigma)+7\sigma-1}{12} \qquad \text{and} \qquad p_B^1 = \frac{L(3-\sigma)+3+7\sigma}{12},$$

  both below the prices without a second period, $\frac{2L+1}{3}$ and $\frac{L+2}{3}$, and both decreasing in σ;
- a marginal consumer $\ell^* = \frac{L-1}{3}$, so consumers in $[0, \ell^*)$ buy from firm B although they prefer firm A;
- in period 2, firm B sells its database to firm A at the price

  $$IP = \frac{(L-1)\,[L(1+\sigma)+11-13\sigma]}{72},$$

  and every consumer buys from the firm she prefers.

With $\sigma = \frac{1}{2}$,

$$\bar{L} = 13 + 12\sqrt{\frac{10}{13}} \approx 23.52.$$

At $L = 31$, firm A charges 17 and firm B charges 7, so $\ell^* = 10$; firm B sells its database for $\frac{85}{4}$; and the firms earn $\Pi_A = \frac{3265}{4}$ and $\Pi_B = \frac{395}{4}$.

Why L must exceed $\bar{L}$. The sale requires $\ell^* > 4$, that is, $L > 13$, but firm B must also not prefer to stop the sale. If consumers expect no sale, the marginal consumer who buys from firm B receives firm A's blanket coupon $\frac{\ell^*}{2}$ in period 2 instead of paying ℓ^*, which tilts her toward firm B; the marginal consumer then satisfies $\ell^* = 2(p_A^1 - p_B^1)$ whenever this is at most 4. For $2 \le p_A^1 - p_B^1 \le 4$, neither expectation is self-fulfilling except $\ell^* = 4$, at which the gains from trade are zero. Firm B's best deviation thus raises its price to $p_B^1 = p_A^1 - 2$, which keeps $\ell^* = 4$ and yields $5(p_A^1 - 2) + \frac{9}{2}$. This deviation is unprofitable if and only if $L > \bar{L}$.

Figure 4.2 depicts firm B's profit as a function of its own first-period price, holding firm A's price at its equilibrium level. At $L = 31$ (blue), the equilibrium price 7 earns $\frac{395}{4} = 98.75$, while the deviation to 15 earns only 79.5. At $L = 20$ (red), the candidate equilibrium price earns 44.13, while raising the price to the kink earns 49.71, so data sharing cannot be an equilibrium.

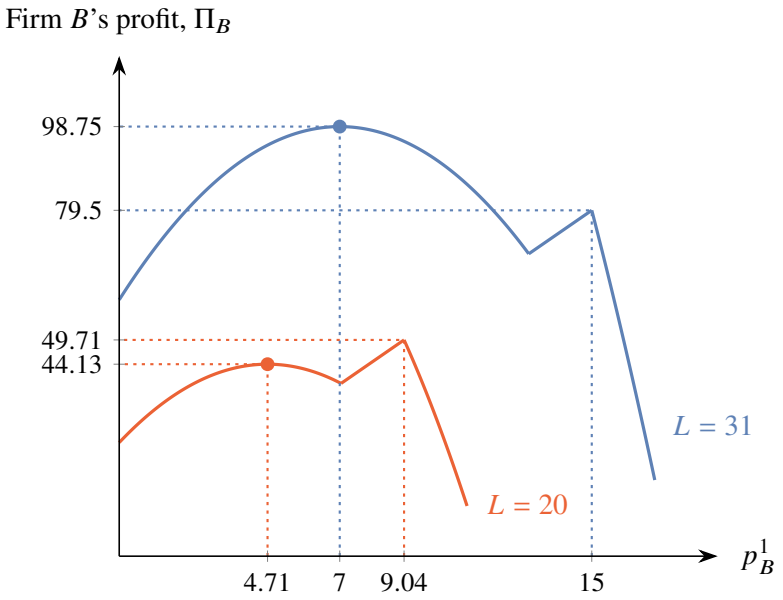


Figure 4.2: Firm B's profit over both periods as a function of its first-period price, given firm A's equilibrium price ($\sigma = \frac{1}{2}$), for $L = 31$ (blue) and $L = 20$ (red). Dots mark the candidate equilibrium prices.

Intuitively, firm A sets a high first-period price because it expects to buy back the names of the loyal consumers it gives up. Both firms price below their static levels. Firm A lowers its price to keep more names in its own database and pay less for firm B's, and firm B lowers its price to collect more names to sell. When σ rises, the sale is worth more to firm B and costs more to firm A, so both incentives grow and both prices fall. When the asymmetry is moderate, however, the data are not worth much, and firm B prefers to raise its price, sell to fewer consumers, and keep its data.

4.6 Fourth Step: Banning the Trade of Data

Consider that a regulator bans the sale of customer data, and let $L > \bar{L}$. In period 2 the firms play the first step with the marginal consumer ℓ^{**}. A consumer who bought from firm B receives firm

A's blanket coupon, so $\ell^{**} = 2(p_A^1 - p_B^1)$, and the firms' profits are

$$\Pi_A = p_A^1(L - \ell^{**}) + \frac{L^2 - (\ell^{**})^2}{2} + \frac{(\ell^{**})^2}{4}, \qquad \Pi_B = p_B^1(\ell^{**} + 1) + \frac{(\ell^{**} + 2)^2}{8}. \tag{4.8}$$

Result 3 (banning data sharing, Proposition 3 and Section 4 of the paper). With a ban:

- first-period prices are $p_A^1 = \frac{3L}{14}$ and $p_B^1 = \frac{L}{14}$, lower than with data sharing, and the marginal consumer is $\ell^{**} = \frac{2L}{7} < \ell^*$;
- in period 2, consumers in $\left(0, \frac{\ell^{**}}{2}\right)$ stay with firm B although they prefer firm A;
- profits fall to
$$\Pi_A = \frac{31L^2}{49} \qquad \text{and} \qquad \Pi_B = \frac{3L^2 + 21L + 49}{98},$$
and consumer surplus rises.

Table 4.1 compares the two regimes at $L = 31$ and $\sigma = \frac{1}{2}$. The ban cuts firm A's first-period price from 17 to 6.64 and firm B's from 7 to 2.21. Both firms lose, $\frac{40821}{196} \approx 208.3$ for firm A and $\frac{12189}{196} \approx 62.2$ for firm B, while consumers gain $\frac{1900}{7} \approx 271.4$.

Table 4.1: Data sharing allowed and banned, with $L = 31$ and $\sigma = \frac{1}{2}$.

Outcome	Sharing allowed	Sharing banned
First-period prices, (p_A^1, p_B^1)	(17, 7)	$\left(\frac{93}{14}, \frac{31}{14}\right) \approx (6.64, 2.21)$
First-period marginal consumer	$\ell^* = 10$	$\ell^{**} = \frac{62}{7} \approx 8.86$
Second-period marginal consumer	0	$\frac{31}{7} \approx 4.43$
Price of firm B's database	$\frac{85}{4}$	–
Profit of firm A	$\frac{3265}{4} \approx 816.3$	$\frac{29791}{49} \approx 608.0$
Profit of firm B	$\frac{395}{4} \approx 98.8$	$\frac{3583}{98} \approx 36.6$
Lost surplus from mismatched purchases	50	$\frac{1922}{49} + \frac{961}{98} \approx 49.03$

Intuitively, under the ban firm A cannot buy back the names it gives up, so it competes harder for them in period 1, and firm B responds by lowering its own price. Consumers gain from these lower prices more than they lose from the higher second-period prices that come with less information, and both firms lose.

4.7 Fifth Step: Total Welfare

Every consumer buys, so total welfare changes only through consumers who buy from the firm they like less. A consumer at $\ell > 0$ who buys from firm B loses ℓ. With data sharing, this happens

in $(0, \ell^*)$ in period 1 and never in period 2. With a ban, it happens in $(0, \ell^{**})$ in period 1 and in $\left(0, \frac{\ell^{**}}{2}\right)$ in period 2. Hence, the ban changes welfare by

$$\Delta W = \frac{(\ell^*)^2 - (\ell^{**})^2}{2} - \frac{(\ell^{**})^2}{8} = \frac{4L^2 - 98L + 49}{882}. \tag{4.9}$$

Result 4 (welfare, Proposition 4 of the paper with $\delta = 1$). Banning data sharing:

- raises total welfare if
$$L > \frac{49 + 21\sqrt{5}}{4} \approx 23.99;$$
- lowers total welfare if $\bar{L} < L < \frac{49+21\sqrt{5}}{4}$, an interval that exists for every σ, since
$$\bar{L} \leq 13 + 12\sqrt{\frac{5}{6}} \approx 23.95.$$

Figure 4.3 depicts the market split in both periods at $L = 31$. With data sharing, the mismatch is concentrated in period 1, where the 10 consumers in $(0, 10)$ buy from firm B, a loss of 50; period 2 is efficient. With a ban, the period-1 mismatch shrinks to $\left(0, \frac{62}{7}\right)$, a loss of $\frac{1922}{49}$, but a period-2 mismatch appears in $\left(0, \frac{31}{7}\right)$, a loss of $\frac{961}{98}$. The net change is $\Delta W = \frac{95}{98} > 0$.

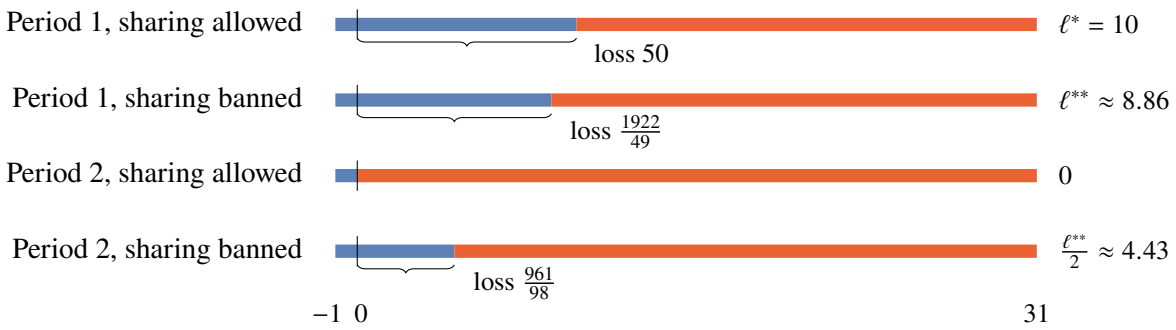


Figure 4.3: Who buys from whom at $L = 31$: consumers in blue buy from firm B and consumers in red buy from firm A; the right column gives the marginal consumer. Braces mark consumers with $\ell > 0$ who buy from firm B.

Intuitively, data sharing moves the inefficiency into period 1. Firm A is willing to give up many loyal consumers in period 1 because it expects to buy back their names, but in period 2 everyone buys the right product. A ban makes firm A fight for these consumers in period 1 and leaves some of them with firm B in period 2. The more asymmetric the firms, the more consumers firm A gives up in period 1 when sharing is allowed, and the more the ban helps.

4.8 Sixth Step: Extension to Discounting

Let firms and consumers discount period 2 by $\delta \in (0, 1]$, as in the paper. The marginal consumer with data sharing is still $\ell^* = \frac{L-1}{3}$, since her period-2 price does not depend on her period-1 choice.

With a ban, the marginal consumer satisfies

$$\ell^{**} = \frac{2(p_A^1 - p_B^1)}{2 - \delta},$$

which in equilibrium gives

$$\ell^{**} = \frac{2\left[L(2-\delta) - 2(1-\delta)\right]}{12 - 5\delta}. \tag{4.10}$$

The ban now changes welfare by

$$\Delta W = \frac{(\ell^*)^2 - (\ell^{**})^2}{2} - \delta\,\frac{(\ell^{**})^2}{8},$$

which is zero at the discount factor

$$\hat{\delta}(L) = \frac{25L^2 - 14L - 47 - (L-7)\sqrt{193L^2 - 590L + 433}}{18(L-2)^2}. \tag{4.11}$$

Result 5 (discounting, Propositions 2 and 4 of the paper). With $\delta \in (0, 1]$ and $L > \bar{L}(\delta)$:

- data sharing is part of the unique subgame perfect equilibrium, where

$$\bar{L}(\delta) \equiv 13 + 12\sqrt{\frac{5\delta}{8 - \delta(1+\sigma)}}$$

  increases in δ;
- if $L < \frac{49+21\sqrt{5}}{4}$, a ban lowers total welfare for every δ;
- if $L > \frac{49+21\sqrt{5}}{4}$, a ban raises total welfare if and only if $\delta > \hat{\delta}(L)$;
- at any δ, a ban lowers both firms' profits and raises consumer surplus.

Figure 4.4 depicts the three regions in the (L, δ) plane with $\sigma = \frac{1}{2}$. To the left of $\bar{L}(\delta)$ (blue), data sharing is not an equilibrium. The cutoff $\hat{\delta}(L)$ (red) starts at 1 when $L \approx 23.99$ and falls toward $\frac{25-\sqrt{193}}{18} \approx 0.62$ as L grows; at $L = 31$, $\hat{\delta} \approx 0.91$.

Intuitively, with data sharing the period-1 marginal consumer, $\ell^* = \frac{L-1}{3}$, does not depend on δ. Under a ban, the more firms value period 2, the harder firm A fights in period 1 for loyal consumers whose names it can no longer buy, so ℓ^{**} in (4.10) falls with δ. The ban then removes more of the period-1 mismatch, at the cost of a period-2 mismatch weighted by δ, so it raises welfare when firms are patient and asymmetric.

4.9 Seventh Step: What the Toy Model Keeps and Drops

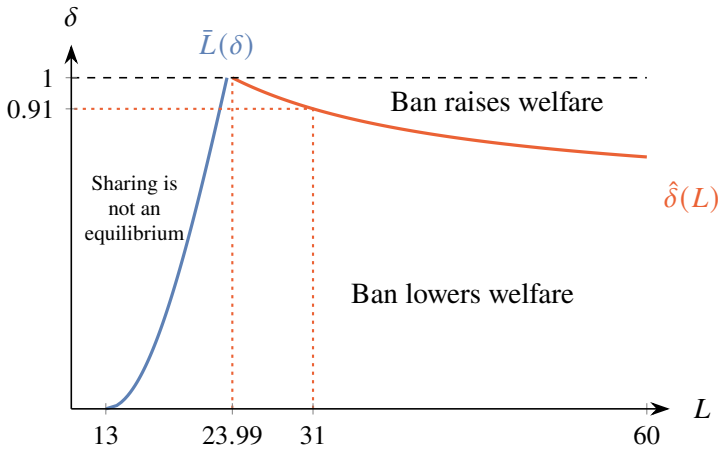


Figure 4.4: Effect of a ban on total welfare in the (L, δ) plane, with $\sigma = \frac{1}{2}$.

What the toy model keeps.

- The one-way trade of data, from the firm with the smaller customer base to the firm with the larger one, and only when the asymmetry is large.
- The first-period strategy of the large firm, which gives up loyal consumers and buys back their names.
- The effects of a ban: lower prices, lower profits, higher consumer surplus, and a welfare effect that depends on the asymmetry and the discount factor.

What it drops.

- General values of ℓ_B, the length of firm B's side of the loyalty interval in the paper: the toy model sets $\ell_B = 1$, so the paper's thresholds, which are multiples of ℓ_B, become numbers.
- Imperfect price discrimination, where data sort consumers into groups rather than identifying each one (Section 5.2 of the paper).
- Two-way trade, which arises when each firm is large in a different group of consumers, and elastic demand (Section 5.2).
- Consumers' privacy concerns, which the paper assumes away.

Fudenberg and Tirole (2000) study the same two-period setting with symmetric firms that know only whether a consumer bought from them, and Shaffer and Zhang (2002) show that a firm with a larger loyal base can gain from targeted pricing; the paper discusses both.

4.10 Connections to Other Chapters

- Chapter 1 studies the welfare effects of third-degree price discrimination by a monopolist.

- Chapter 36 studies firms that buy data on how consumers influence one another and use them to price. With competition, firms race to acquire data that lower their profits.
- In Chapter 35, platforms sell ads using data on consumer preferences, and more informative data raise product consumption but tend to lower ad revenue.

References

Fudenberg, D., and J. Tirole (2000). "Customer Poaching and Brand Switching." *RAND Journal of Economics*, 31(4), pp. 634–657.

Liu, Q., and K. Serfes (2006). "Customer Information Sharing among Rival Firms." *European Economic Review*, 50(6), pp. 1571–1600.

Shaffer, G., and Z. J. Zhang (2002). "Competitive One-to-One Promotions." *Management Science*, 48(9), pp. 1143–1160.

Part II

Entry Barriers

The six chapters in this part study entry into a market. The first three ask whether an incumbent can exclude a rival that is more efficient than itself. The Chicago School argued that such exclusion is not profitable, and each of these chapters tests this argument against a different practice. Chapter 5 shows that the argument holds for rebates that a dominant supplier offers to buyers who deal with it exclusively: a rebate does not commit the buyer before the entrant makes its offer, so a more efficient entrant can always outbid it. Chapters 6 and 7 find two practices that do exclude. In Chapter 6, the incumbent shares its profit with the retailers that the entrant needs, and no retailer carries the entrant when the entrant's profit is too small to compensate it. In Chapter 7, a vertically integrated incumbent refuses to supply the input to a downstream rival, which then cannot earn the profits it needs to cover its entry cost.

In Chapter 8, the owner of the input wants downstream firms to enter, and its contracts decide how many do. A wholesale price adds a markup that lowers the surplus of each entrant, but without a fixed fee the entrant keeps more of this surplus than its bargaining power gives it. The second effect dominates, so more firms enter under wholesale prices than under two-part tariffs, and when the supplier's bargaining power is high, consumers and the supplier itself gain from wholesale prices.

The last two chapters study limit pricing: a potential entrant that does not know the industry's cost or demand infers it from the incumbents' prices. In Chapter 9, several incumbents share private information about a common cost. Because the entrant's cost moves with theirs, a high price signals high costs and keeps the entrant out, the opposite of the single-incumbent model of Milgrom and Roberts (1982), and deterrence works only if every incumbent restricts its output. In Chapter 10, two incumbents disagree about entry, because one sells a substitute for the entrant's good and the other a complement, and they separate only when entry matters little to one of them.

In the first four chapters, entry depends on the terms on which an incumbent deals with buyers, retailers, or downstream firms. In the last two, the incumbents commit to nothing, and what keeps the entrant out is what their prices reveal.

Chapter 5

Discounts as a Barrier to Entry

Companion note to Ide, E., J.-P. Montero, and N. Figueroa (2016), "Discounts as a Barrier to Entry," American Economic Review 106(7): 1849–1877.

5.1 Introduction

Dominant suppliers often reward loyal retailers. Consider a large brewer that tells a bar: "as long as you buy only my beer, you get 10% off the list price on every keg; otherwise, you pay the list price." Such a rebate looks like an exclusive dealing contract, and the post-Chicago literature shows that exclusive dealing can keep a more efficient entrant out of the market. Should antitrust authorities treat loyalty rebates in the same way?

Ide et al. (2016) show that the answer is no, and the reason is timing. An exclusive dealing contract commits the retailer before the entrant shows up: if she later buys from the entrant, she pays a penalty fixed in advance. A rebate, in contrast, leaves the retailer free to decide whether to stay exclusive after she sees the entrant's offer. A more efficient entrant can then always offer a price that beats the rebate, since its lower cost creates surplus that it can share with the retailer.

The paper finds that off-list-price rebates, which give a discount on every unit, cannot profitably block a more efficient entrant in any of the three post-Chicago settings: an entrant with uncertain costs, two buyers and scale economies, and retailers that compete downstream. In the last setting, however, a lump-sum rebate, a fixed payment for staying exclusive, can block entry. Rebates do block entrants whose entry would destroy surplus, which raises welfare. These results matter for antitrust authorities, who treat rebates with suspicion because they look like exclusive dealing.

The toy model fixes the consumers' valuation, the incumbent's cost, and the share of consumers who can buy from either firm throughout. The first step solves the benchmark without contracts. The second to fourth steps compare exclusive dealing and rebates in the three post-Chicago settings, the fifth step shows that rebates block inefficient entrants, and the sixth step lists what the toy model keeps and drops.

5.2 The Toy Model

Consumers and manufacturers. A unit mass of final consumers values a good at $v = \frac{3}{2}$. An incumbent manufacturer, I, produces at marginal cost c_I, which we normalize to $c_I = 1$. A potential entrant, E, produces at marginal cost c_E after paying an entry cost $F \geq 0$. Whenever c_E is known, we normalize it to $c_E = 0$, so the entrant's cost advantage is $c_I - c_E = 1$ and the examples below differ only in F. Only a share $\lambda = \frac{1}{2}$ of consumers, the contestable demand, regards both products as identical; the remaining consumers buy only from I. Entry is efficient if the entrant's average cost on the contestable units is below the incumbent's cost,

$$c_E + \frac{F}{\lambda} < c_I. \tag{5.1}$$

Retailers and timing. Manufacturers sell through retailers, who have no other costs. On date 1, I offers a contract to the retailer(s); on date 2, E offers a take-or-pay contract for λ units, which is cancelled if E does not enter; on date 3, E decides whether to enter; on date 4, retailers buy according to the contracts they hold, or in the spot market, where I and E compete with nonlinear prices.

Contracts. The paper compares three contracts.

- An exclusive dealing contract (t, w, D): a payment t from I on date 1, a wholesale price w, and a penalty D that the retailer pays I if she buys from E.
- An off-list-price rebate (r, R): a list price r per unit, and a discount R on every unit if the retailer buys only from I.
- A lump-sum rebate (r, L): the list price r and a fixed payment L if the retailer buys only from I, regardless of the quantity.

A retailer that switches the contestable λ units to E loses the discount on all units, so the price that E must undercut on the contestable units, the effective price of a rebate, is

$$x \equiv r - \frac{R}{\lambda}. \tag{5.2}$$

The key difference is timing: the penalty D is committed on date 1, while the rebate is lost only if the retailer, after seeing E's offer on date 2, decides to buy from E.

5.3 First Step: The No-Contract Benchmark

Consider that entry is efficient and I offers no contract. In the spot market, E sells the λ contestable units at a price slightly below c_I, and I sells the remaining $1 - \lambda$ units with a two-part tariff, a unit price c_I and a fixed fee $(1 - \lambda)(v - c_I)$, which extracts the retailer's surplus on these units. The payoffs are

$$\pi_I^{NC} = (1 - \lambda)(v - c_I) = \frac{1}{4}, \qquad \pi_E^{NC} = \lambda(c_I - c_E) - F, \qquad \pi_B^{NC} = \lambda(v - c_I) = \frac{1}{4}. \tag{5.3}$$

The retailer keeps $\lambda(v - c_I)$ because she buys the contestable units at c_I instead of v.

Result 1 (no-contract benchmark). Without contracts, the efficient entrant enters, the incumbent earns $\frac{1}{4}$ from the non-contestable consumers, and the retailer earns $\frac{1}{4}$ from the competition between manufacturers on the contestable units.

Every contract that I offers is profitable only if it yields more than $\pi_I^{NC} = \frac{1}{4}$.

5.4 Second Step: Rent Shifting with an Uncertain Entrant (One Buyer)

Setting. Following Aghion and Bolton (1987), there is one retailer, $F = 0$, and on date 1 neither I nor the retailer knows c_E, which is uniformly distributed on $[0, c_I] = [0, 1]$, so $G(c) = c$ and every entrant is more efficient than I.

Exclusive dealing. A retailer who signed $(0, w, D)$ switches to E's offer w_E if paying w on $1 - \lambda$ units, D, and w_E on λ units costs her no more than paying w on all units, that is, if $w_E \leq w - \frac{D}{\lambda}$. The entrant therefore enters if and only if $c_E \leq y$, where

$$y \equiv w - \frac{D}{\lambda}.$$

Since the retailer pays w per unit in every case, she signs if $v - w \geq \pi_B^{NC}$, so I sets

$$w = v - \lambda(v - c_I) = \frac{5}{4}.$$

Writing $D = \lambda(w - y)$, the incumbent's expected profit is

$$\pi_I^{ED}(y) = (w - c_I) + \lambda\,(c_I - y)\,G(y) = \frac{1}{4} + \frac{1}{2}\,(1 - y)\,y. \tag{5.4}$$

The second term is the expected penalty collected from an entrant with cost below y, net of the sales lost. Maximizing over y gives

$$y^* = c_I - \frac{G(y^*)}{g(y^*)} = \frac{1}{2}, \qquad D^* = \lambda(w - y^*) = \frac{3}{8},$$

and

$$\pi_I^{ED} = \frac{1}{4} + \frac{1}{8} = \frac{3}{8} > \frac{1}{4}. \tag{5.5}$$

Result 2 (Aghion and Bolton). With an exclusive contract, the incumbent sets the effective price $y^* = \frac{1}{2}$ and earns $\frac{3}{8}$. Entrants with $c_E \in (\frac{1}{2}, 1)$, half of all entrants, are more efficient than I but stay out, and the expected welfare loss is

$$\lambda \int_{1/2}^{1} (c_I - c)\, dG(c) = \frac{1}{16}.$$

The penalty works as an entry fee: the incumbent extracts part of the entrant's cost advantage, at the price of excluding the entrants whose advantage is small.

Off-list-price rebates. Under a rebate (r, R), the retailer switches to E if

$$v - (1-\lambda)r - \lambda w_E \geq v - (r - R),$$

that is, if $w_E \leq x$, so E enters if and only if $c_E \leq x$. The incumbent sells $1-\lambda$ units at r after entry and all units at $r - R$ otherwise, and substituting $R = \lambda(r - x)$,

$$\pi_I^R(r, x) = (1-\lambda)(r - c_I) + \lambda\,(x - c_I)\,[1 - G(x)]. \tag{5.6}$$

A retailer who buys from E pays the list price r for the remaining units, so she would rather buy nothing from I if $r > v$; therefore $r \leq v$.

Result 3 (Proposition 1 of the paper). No off-list-price rebate with $x < c_I$ is profitable: for every such rebate, $\pi_I^R < (1-\lambda)(v - c_I) = \frac{1}{4}$, since $r \leq v$ and the second term of (5.6) is negative.

Figure 5.1 depicts the incumbent's expected profit as a function of the effective price it sets, under exclusive dealing (5.4) and under the best rebate (5.6) with $r = v$, which becomes

$$\pi_I^R(v, x) = \frac{1}{4} - \frac{1}{2}(1 - x)^2.$$

The exclusive contract earns more than the benchmark for every $y \in (0, c_I)$ and peaks at $y^* = \frac{1}{2}$, while the rebate earns less than the benchmark for every $x < c_I$.

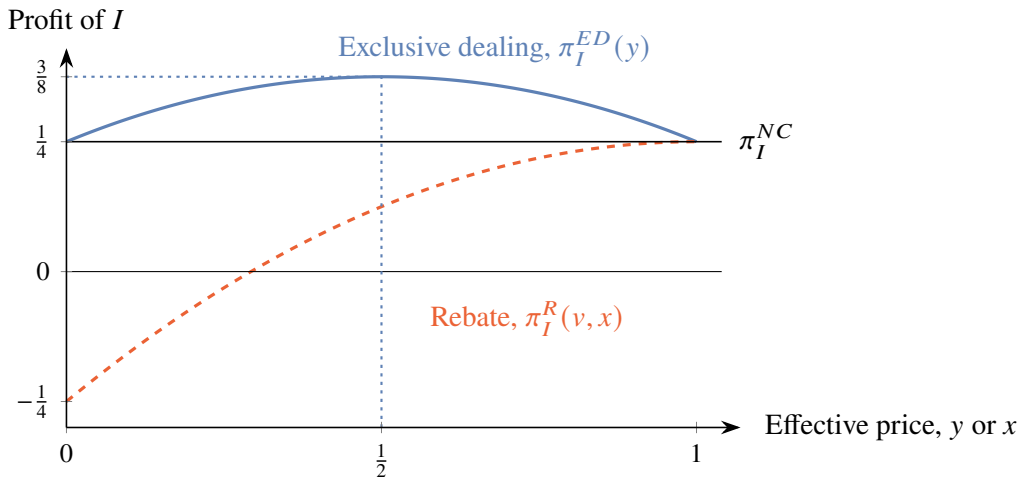


Figure 5.1: Incumbent's expected profit as a function of the effective price, with an exclusive contract and with an off-list-price rebate ($v = \frac{3}{2}$, $c_I = 1$, $\lambda = \frac{1}{2}$, c_E uniform on $[0, 1]$).

Intuitively, the two contracts look alike: relabeling $w = r - R$ and $D = (1-\lambda)R$ turns (5.6) into the exclusive's profit, since giving up the rebate on the $1-\lambda$ non-contestable units works as a penalty. The difference is the constraint on the list price. With an exclusive, the penalty is committed on date 1, so it does not enter the retailer's later decision, and the only constraint is $w \leq v$. With a rebate, nothing is committed, so the list price the retailer would pay after switching must satisfy

$$r = w + \frac{D}{1-\lambda} \leq v. \tag{5.7}$$

The exclusive optimum, $w = \frac{5}{4}$ and $D = \frac{3}{8}$, would require

$$r = \frac{5}{4} + \frac{3/8}{1/2} = 2 > \frac{3}{2},$$

a list price the retailer would never pay.

Figure 5.2 depicts this comparison in the (D, w) plane. The retailer accepts any exclusive with $w \le \frac{5}{4}$ (below the horizontal line), while a rebate must also satisfy (5.7), the region below the blue line. The dashed red lines are effective prices $w - \frac{D}{\lambda} = y$: the exclusive optimum lies on $y = \frac{1}{2}$, far outside the rebate region, and every rebate that lowers the effective price below c_I forces the incumbent to lower w as well, which costs it more than the rent it extracts.

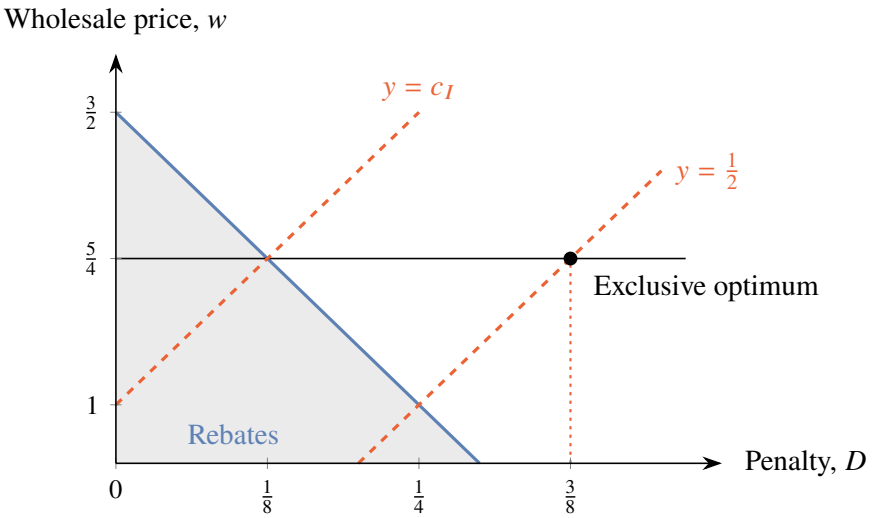


Figure 5.2: Exclusive contracts and rebates in the (D, w) plane. Rebates must lie in the shaded region, $w + \frac{D}{1-\lambda} \le v$; dashed lines are effective prices $y = w - \frac{D}{\lambda}$.

5.5 Third Step: Scale Economies and Two Buyers

Setting. Following Rasmusen et al. (1991) and Segal and Whinston (2000), there are two retailers, B_1 and B_2, each serving half of the consumers in a separate market. The entrant's cost is known, $c_E = 0$, and its entry cost is $F = \frac{7}{16}$. Entry is efficient, since

$$\lambda(c_I - c_E) - F = \frac{1}{16} > 0,$$

but the entrant needs both retailers, since serving one yields at most

$$\frac{\lambda(v - c_E)}{2} = \frac{3}{8} < F.$$

The two conditions hold together because $v < 2c_I$.

Exclusive dealing. The incumbent offers B_1 slightly more than her benchmark payoff, $\frac{\pi_B^{NC}}{2} = \frac{1}{8}$, and offers B_2 almost nothing. If B_1 signs, E cannot reach the scale it needs, so B_2 expects to pay v anyway and also signs. The incumbent earns

$$v - c_I - \frac{1}{8} = \left(1 - \frac{\lambda}{2}\right)(v - c_I) = \frac{3}{8} > \frac{1}{4},$$

since it pays one buyer to exclude and exploits the other.

Rebates. To replicate this, the incumbent offers B_1 a large rebate (v, R_1) and B_2 almost none. But B_2 has signed nothing, so E can sell to her at a price close to v, and use the profit to outbid B_1's rebate. B_1 switches if E offers her a price w_{E1} such that

$$\frac{\lambda(v - w_{E1})}{2} \geq \frac{R_1}{2},$$

and E breaks even when

$$\frac{\lambda(w_{E1} - c_E)}{2} + \frac{\lambda(v - c_E)}{2} = F,$$

so to block entry the incumbent needs

$$\frac{R_1}{2} > \lambda(v - c_I) + [\lambda(c_I - c_E) - F] = \frac{1}{4} + \frac{1}{16} = \frac{5}{16}, \tag{5.8}$$

which leaves it with at most

$$v - c_I - \frac{5}{16} = \frac{3}{16} < \frac{1}{4}.$$

Result 4 (Propositions 2 and 3 of the paper). With scale economies and two buyers:

- exclusive contracts block the efficient entrant and raise the incumbent's profit from $\frac{1}{4}$ to $\frac{3}{8}$;
- rebates that block the entrant leave the incumbent with at most

  $$(1 - \lambda)(v - c_I) - [\lambda(c_I - c_E) - F],$$

  which is below $\frac{1}{4}$ whenever entry is efficient.

Figure 5.3 depicts the incumbent's profit from a blocking rebate, $F - \frac{1}{4}$, as a function of the entry cost F. It lies below the benchmark $\frac{1}{4}$ for every F at which entry is efficient, $F < \lambda(c_I - c_E) = \frac{1}{2}$, and above it only when entry is inefficient.

Intuitively, to keep B_1 exclusive, the incumbent must give her more than the entrant can offer, and the entrant can pass on to B_1 the surplus it extracts from the unprotected B_2. When entry is efficient, this payment exceeds what the incumbent gains by exploiting the buyers.

5.6 Fourth Step: Retailers Who Compete Downstream

Setting. Following Simpson and Wickelgren (2007) and Asker and Bar-Isaac (2014), two retailers now compete in prices for the same consumers. Since E needs only one retailer to reach the

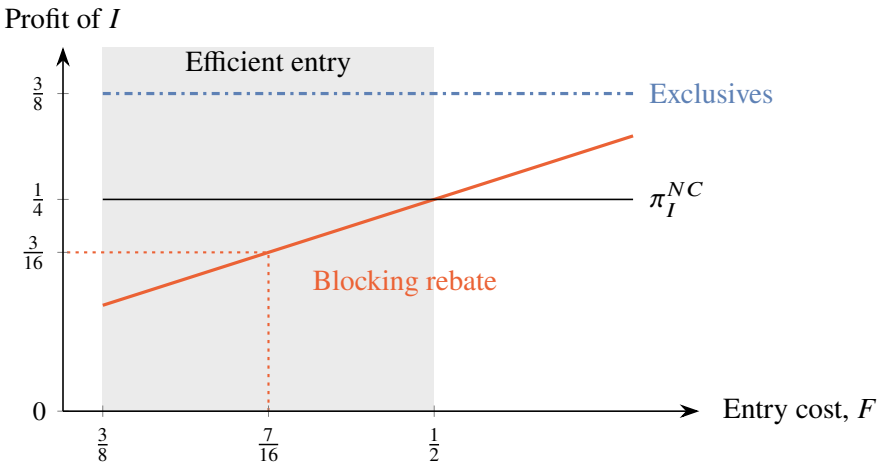


Figure 5.3: Incumbent's profit from an exclusive contract and from the smallest rebate that blocks entry, as a function of the entry cost F ($c_E = 0$). The shaded interval has scale economies and efficient entry.

contestable consumers, and price competition passes any cost saving to consumers, retailers earn zero in the benchmark.

Exclusive dealing. A retailer earns zero whether or not she signs, so the incumbent gets both retailers to sign exclusives for an arbitrarily small payment and earns $v - c_I = \frac{1}{2}$ (Proposition 4 of the paper).

Off-list-price rebates. A discount per unit lowers the retailer's marginal cost, so competition passes it to consumers: exclusive retailers sell at $r - R$. To keep E out, the incumbent must set $r - R < c_E + \frac{F}{\lambda}$, below its own cost, so blocking entry with these rebates is a form of limit pricing that loses money (Proposition 5 of the paper).

Lump-sum rebates. A fixed payment L for exclusivity does not change the retailer's marginal cost, so competition does not pass it to consumers. A retailer that switches earns $\lambda(c_I - w_E)$ on the contestable consumers and loses L, so the entrant can win her only if $L < \lambda(c_I - c_E) - F$. Paying each retailer $L = \lambda(c_I - c_E) - F$ blocks entry and yields $v - c_I - 2L$, which exceeds π_I^{NC} if and only if

$$\frac{\lambda(v - c_I)}{2} > \lambda(c_I - c_E) - F. \tag{5.9}$$

With $c_E = 0$ and $F = \frac{7}{16}$, $L = \frac{1}{16}$ and the incumbent earns

$$\frac{1}{2} - \frac{1}{8} = \frac{3}{8} > \frac{1}{4},$$

so lump-sum rebates exclude an efficient entrant. With $F = 0$, instead, $L = \frac{1}{2}$ and blocking entry costs more than it yields.

Result 5 (Propositions 4–6 of the paper). With retailers who compete downstream:

- exclusive contracts block entry for an arbitrarily small payment;
- off-list-price rebates never block an efficient entrant profitably;

- lump-sum rebates block an efficient entrant profitably if and only if (5.9) holds, that is, when the entrant's surplus from entry is below half of the incumbent's margin on the contestable consumers.

Figure 5.4 depicts condition (5.9) in the (c_E, F) plane. Entry is efficient below the solid line, $F < \lambda(c_I - c_E)$, and lump-sum rebates exclude the entrant in the shaded region, where the entrant's surplus is positive but below

$$\frac{\lambda(v - c_I)}{2} = \frac{1}{8}.$$

The dot marks the example $c_E = 0$, $F = \frac{7}{16}$.

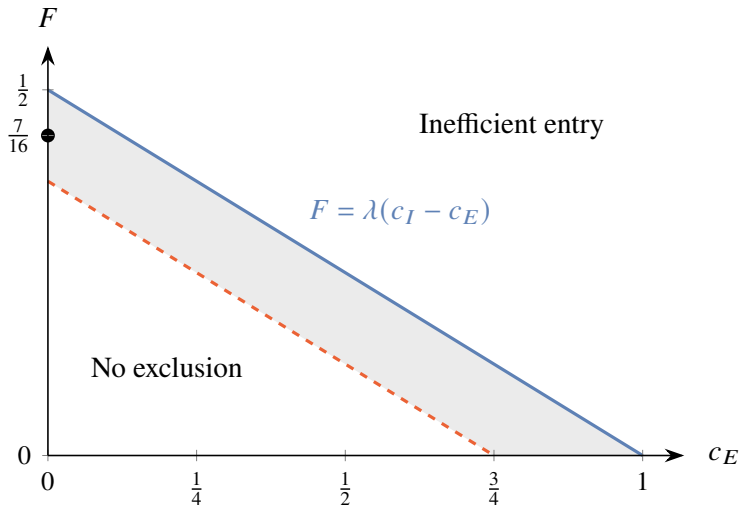


Figure 5.4: Downstream competition: parameters for which lump-sum rebates profitably exclude an efficient entrant (shaded), with $v = \frac{3}{2}$, $c_I = 1$, $\lambda = \frac{1}{2}$. The dashed line is $F = \lambda(c_I - c_E) - \frac{1}{8}$.

Intuitively, lump-sum rebates work because they reward retailers without lowering retail prices, so retailers internalize part of the incumbent's loss from entry. An off-list-price rebate, in contrast, is competed away in lower retail prices, and the retailers have nothing to lose by switching. The size of the contestable demand matters here: with $c_E = 0$ and $F = \frac{7}{16}$, entry is efficient both for $\lambda = \frac{1}{2}$ and for $\lambda = 1$, but condition (5.9) holds only for $\lambda = \frac{1}{2}$, since a smaller contestable demand raises the entrant's entry cost per contestable unit. The paper's remark that the size of the contestable demand does not matter applies to off-list-price rebates and to settings without entry costs.

5.7 Fifth Step: A Generalized Chicago Critique and Inefficient Entry

The common thread of the second to fourth steps is what the paper calls a generalized Chicago critique. Excluding an efficient rival destroys surplus, so if every party could bargain at the same time, exclusion could not occur. Post-Chicago models obtain exclusion because one party, the entrant or final consumers, is absent when the buyer signs, and the buyer is locked in before that

party shows up. Exclusive dealing provides this lock-in; off-list-price rebates do not, because the buyer decides whether to stay exclusive only after seeing the entrant's offer. The paper notes that adding an unconditional transfer from the buyer to the incumbent, paid whether or not she later stays exclusive, would restore the commitment, but such transfers have not been documented in rebate cases.

Inefficient entry. Rebates without commitment do block entrants whose entry destroys surplus. Consider that $c_E = 0$ and $F = \frac{5}{8}$ with one retailer, so E's average cost on the contestable units is $c_E + \frac{F}{\lambda} = \frac{5}{4}$, above c_I but below v. Without a contract, E would enter, since the surplus from entry that E and the retailer can share, $\lambda(v - c_E) - F = \frac{1}{8}$, is positive. The incumbent offers the list price $r = v$ and a discount $R = \frac{1}{8}$, so the retailer pays $\frac{11}{8}$ per unit if exclusive, and the effective price on the contestable units is $x = c_E + \frac{F}{\lambda} = \frac{5}{4}$. The entrant cannot undercut this price, and the incumbent earns

$$\pi_I = (1 - \lambda)(v - c_I) + \lambda \left(c_E + \frac{F}{\lambda} - c_I \right) = \frac{1}{4} + \frac{1}{8} = \frac{3}{8}. \tag{5.10}$$

Result 6 (inefficient entry). With one retailer and a known entrant's cost, the incumbent profitably blocks entry with an off-list-price rebate if and only if the entrant's average cost on the contestable units exceeds the incumbent's cost, $c_E + \frac{F}{\lambda} > c_I$. The incumbent's gain from blocking, $\lambda(c_E + \frac{F}{\lambda} - c_I)$, equals the welfare loss that entry would cause.

Figure 5.5 depicts, as a function of the entry cost F with $c_E = 0$, the welfare effect of entry,

$$\lambda(c_I - c_E) - F = \frac{1}{2} - F,$$

and the incumbent's gain from a blocking rebate. The gain is zero while entry is efficient, $F < \frac{1}{2}$, and equals the welfare loss from entry above that cost, up to $F = \lambda v = \frac{3}{4}$, where the entrant would not enter even without a contract.

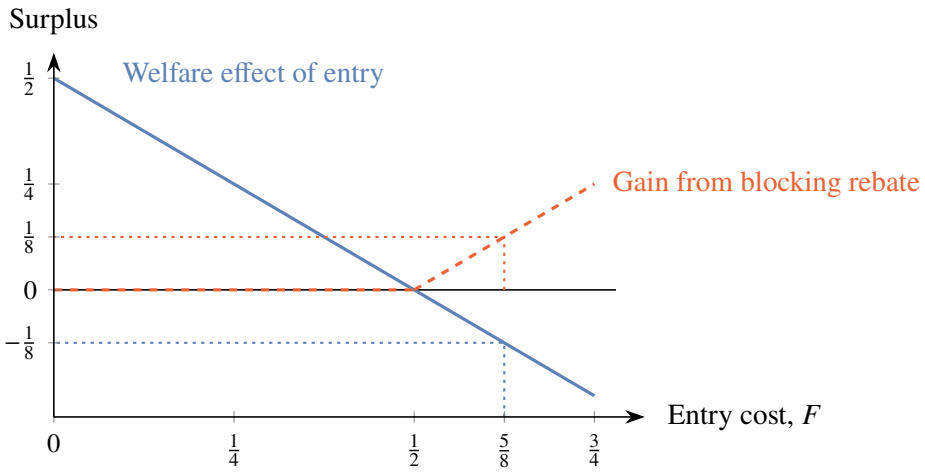


Figure 5.5: Welfare effect of entry and incumbent's gain from a blocking off-list-price rebate, as a function of the entry cost ($c_E = 0$, one retailer).

Intuitively, the incumbent is the residual claimant of the surplus created by removing an inefficiency, so it can always outbid the entrant for the retailer's loyalty, without any commitment.

5.8 Sixth Step: What the Toy Model Keeps and Drops

Table 5.1 reports the incumbent's profits in the examples of the second to fifth steps.

Table 5.1: Incumbent's profit with each contract ($v = \frac{3}{2}$, $c_I = 1$, $\lambda = \frac{1}{2}$; benchmark $\pi_I^{NC} = \frac{1}{4}$).

Setting	Exclusive dealing	Blocking rebate	Entry blocked?
Uncertain entrant, $c_E \sim U[0,1]$	$\frac{3}{8}$	below $\frac{1}{4}$	Exclusive only
Two buyers, $c_E = 0$, $F = \frac{7}{16}$	$\frac{3}{8}$	$\frac{3}{16}$	Exclusive only
Downstream competition, off-list-price	$\frac{1}{2}$	below 0	Exclusive only
Downstream competition, lump-sum, $F = \frac{7}{16}$	$\frac{1}{2}$	$\frac{3}{8}$	Both
Inefficient entrant, $c_E = 0$, $F = \frac{5}{8}$	–	$\frac{3}{8}$	Rebate (efficiently)

What the toy model keeps.

- The role of commitment: exclusive dealing locks the buyer in before the entrant appears, while an off-list-price rebate leaves her free to switch after seeing the entrant's offer.
- The equivalence that fails: relabeling a rebate as an exclusive works for profits but not for the constraint $r \leq v$ on the list price.
- The difference between off-list-price and lump-sum rebates when retailers compete.
- Rebates that block inefficient entry, which raise welfare.

What it drops.

- Downward-sloping demand and incomplete surplus extraction in the non-contestable segment, which underlie bundling theories of exclusion.
- Asymmetric information about the buyer, contracts that combine rebates with unconditional transfers, and repeated interaction.
- Antitrust implications: the paper argues that the off-list-price rebates in *Michelin II* and *British Airways* could not have blocked a more efficient rival, while the apparently lump-sum rebates in *AMD v. Intel* were potentially exclusionary.

5.9 Connections to Other Chapters

- Chapter 9 studies entry deterrence when the entrant does not know the industry's cost, so the incumbents' prices signal it.
- Chapter 19 studies another contract signed before competition takes place: forward sales, which make the spot market more competitive when firms meet only once.

- Chapter 38 shows that a cost reduction by an inefficient firm can lower welfare, because it moves production away from more efficient rivals. In this chapter, an off-list-price rebate cannot keep a more efficient entrant out.
- Chapter 6 shows that vertical practices such as resale price maintenance can exclude an entrant when they let the incumbent share its profit with the retailers that the entrant needs.
- Chapter 7 studies an integrated incumbent that refuses to supply a more efficient downstream rival, so that the rival cannot cover its entry cost.

References

Aghion, P., and Bolton, P. (1987). Contracts as a barrier to entry. *American Economic Review*, 77(3), 388–401.

Asker, J., and Bar-Isaac, H. (2014). Raising retailers' profits: On vertical practices and the exclusion of rivals. *American Economic Review*, 104(2), 672–686.

Ide, E., Montero, J.-P., and Figueroa, N. (2016). Discounts as a barrier to entry. *American Economic Review*, 106(7), 1849–1877.

Rasmusen, E. B., Ramseyer, J. M., and Wiley, J. S., Jr. (1991). Naked exclusion. *American Economic Review*, 81(5), 1137–1145.

Segal, I. R., and Whinston, M. D. (2000). Naked exclusion: Comment. *American Economic Review*, 90(1), 296–309.

Simpson, J., and Wickelgren, A. L. (2007). Naked exclusion, efficient breach, and downstream competition. *American Economic Review*, 97(4), 1305–1320.

Chapter 6

Raising Retailers' Profits to Exclude a Rival

Companion note to Asker, J., and H. Bar-Isaac (2014), "Raising Retailers' Profits: On Vertical Practices and the Exclusion of Rivals," American Economic Review 104(2): 672–686.

6.1 Introduction

Many manufacturers share part of their profits with the retailers that carry their products. Resale price maintenance, in which the manufacturer sets the retail price, guarantees retailers a margin, while slotting fees and loyalty rebates pay them directly. When the U.S. Supreme Court ended the per se ban on resale price maintenance in its 2007 *Leegin* decision, it warned that a manufacturer with market power could use the practice to give retailers a reason not to sell the products of new entrants. Can sharing profits with retailers keep a more efficient rival out of the market, even if no contract forbids retailers from carrying it?

Asker and Bar-Isaac (2014) show that it can, when a new manufacturer needs at least one retailer to reach consumers and its entry would lower industry profits. A retailer that carries the entrant then loses the stream of profits that the incumbent shares with it. The entrant can offer this retailer at most what it expects to earn after entry. When this amount falls short of what each retailer receives from the incumbent, the entrant cannot compensate any retailer for its loss.

Retailers then protect the incumbent's market, and the more efficient entrant stays out. Each additional retailer must also be paid, so exclusion is harder with more retailers. Ide et al. (2016) show that rebates that do not commit buyers cannot exclude a more efficient entrant; here, retailers are not committed either, yet the loss of future profits keeps them loyal. For policy, banning a single practice is not enough when other practices can transfer the same profits.

The toy model uses linear demand, an incumbent with a higher cost than the entrant, and a small number of retailers. The first step solves the benchmark without vertical practices, the second step shows how the incumbent shares its profits with retailers, the third step finds when retailers refuse to carry the entrant, the fourth step studies welfare and a ban on these practices, and the fifth step

lists what the toy model keeps and drops.

6.2 The Toy Model

Manufacturers and demand. An incumbent manufacturer, I, and a potential entrant, E, produce a homogeneous good. The incumbent's marginal cost is $c = \frac{1}{10}$ and the entrant's is zero, so the entrant is the more efficient producer. At retail price p, consumers buy

$$q = 1 - p \tag{6.1}$$

units. To enter, the entrant pays a one-time cost $F \geq 0$. The paper works with general profit functions and allows for differentiated products; the toy model uses the homogeneous-good setting of the paper's application to resale price maintenance, where the paper also sets $F = 0$, and keeps the entry cost of the paper's baseline model.

Retailers. Consumers buy only through $n \geq 2$ retailers, which are perfect substitutes for each other. A retailer's only cost is the wholesale price it pays, retailers compete in prices, and retailers that charge the same price split demand equally. The entrant can reach consumers only if at least one retailer agrees to carry its good, and a retailer that agrees gives the entrant access to the market forever.

States and timing. Time is infinite, and all firms discount future profits with the discount factor $\delta \in (0, 1)$. The market is in state M, where the incumbent is the only manufacturer, or in state C, where both manufacturers are active. The game starts in state M, and a period in state M proceeds as follows:

(1) the incumbent sets its wholesale price and, if allowed, a vertical practice that gives each retailer an amount $T \geq 0$ at the end of the period, provided that the incumbent is still the only manufacturer;

(2) the entrant offers each retailer r a payment $T_r \geq 0$, paid if retailer r carries its good;

(3) retailers simultaneously accept or reject the entrant's offers;

(4) if at least one retailer accepts, the entrant pays F and enters, and the market is in state C from this period on; otherwise, retailers sell the incumbent's good, profits and transfers are realized, and the next period starts in state M.

State C is absorbing: in every period, the two manufacturers simultaneously set their wholesale prices (and, if allowed, their vertical practices), and then retailers compete in prices. We study stationary Markov perfect equilibria, in which strategies depend only on the state, as in the paper. In the paper, profits in state M are realized before the entrant's offer, and the entrant becomes active one period after a retailer accepts. Starting competition at once simplifies the discounting, and the condition for exclusion in the third step coincides with the paper's.

6.3 First Step: The Benchmark without Vertical Practices

Consider first that manufacturers can only set wholesale prices.

Retail competition. Retailers are perfect substitutes and face the same cost, the wholesale price w. Price competition between them drives the retail price down to w, so every retailer earns zero, in both states.

State M. The incumbent chooses w to maximize $(w - c)(1 - w)$. The first-order condition, $1 - 2w + c = 0$, gives the monopoly outcome

$$p_M = \frac{1+c}{2} = \frac{11}{20}, \qquad q_M = \frac{1-c}{2} = \frac{9}{20}, \qquad \pi_M = \frac{(1-c)^2}{4} = \frac{81}{400}. \tag{6.2}$$

Since retailers earn nothing, π_M is also the industry profit before entry.

State C. The two manufacturers compete in wholesale prices with costs $\frac{1}{10}$ and 0. As in Bertrand competition with asymmetric costs, the entrant sets a wholesale price equal to the incumbent's cost and serves the whole market. This price lies below the entrant's own monopoly price, $\frac{1}{2}$, so

$$p_C = c = \frac{1}{10}, \qquad q_C = \frac{9}{10}, \qquad \pi_{EC} = c\,(1-c) = \frac{9}{100}, \qquad \pi_{IC} = 0, \tag{6.3}$$

where π_{EC} and π_{IC} are the entrant's and the incumbent's per-period profits after entry. Entry lowers industry profit from $\frac{81}{400}$ to $\frac{9}{100} = \frac{36}{400}$.

Accommodation. The entrant's value from entering, net of the entry cost but before any payment to retailers, is

$$V_E = \frac{\pi_{EC}}{1-\delta} - F = \frac{9}{100(1-\delta)} - F, \tag{6.4}$$

which is positive if $F < \frac{9}{100(1-\delta)}$. A retailer earns zero in both states, so, when every other retailer rejects, a retailer that receives a positive offer strictly prefers to accept it. Hence, the entrant can always enter by offering a small payment to one retailer. When one retailer accepts, entry occurs regardless of what the others do, so accepting even a zero offer is a best response, and the entrant pays nothing.

Result 1 (no vertical practices). Consider that $F < \frac{9}{100(1-\delta)}$ and that manufacturers set only wholesale prices. Then:

- retailers earn zero in both states;
- in every equilibrium, the entrant enters in the first period, and there is an equilibrium in which it pays retailers nothing;
- the retail price falls from $p_M = \frac{11}{20}$ to $p_C = \frac{1}{10}$, and per-period industry profit falls from $\pi_M = \frac{81}{400}$ to $\pi_{EC} = \frac{9}{100}$.

Figure 6.1 depicts the two outcomes. Before entry, the incumbent sells $\frac{9}{20}$ units at $\frac{11}{20}$ and earns the blue rectangle; after entry, the entrant sells $\frac{9}{10}$ units at $\frac{1}{10}$ and earns the green rectangle, which is smaller.

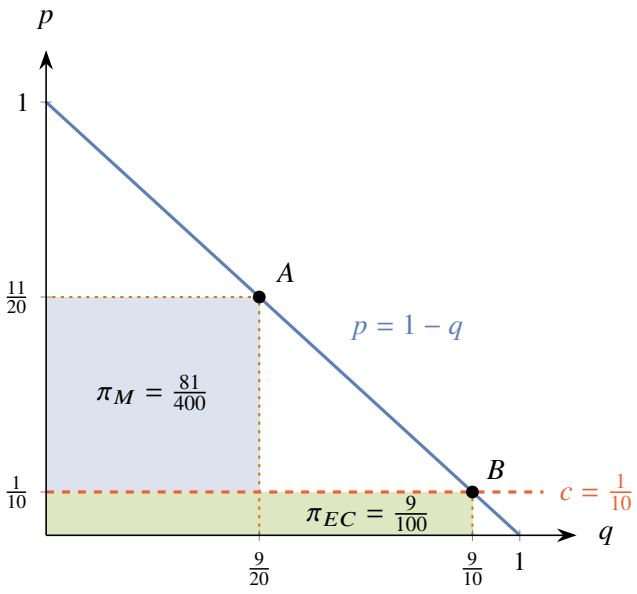


Figure 6.1: Monopoly outcome before entry (point A) and Bertrand outcome after entry (point B). The entrant's marginal cost is zero; the dashed line is the incumbent's marginal cost.

Intuitively, price competition between retailers leaves them nothing to protect. They pass the wholesale price on to consumers in both states, so carrying the entrant costs them nothing, and the more efficient entrant is accommodated, as in standard Bertrand competition.

6.4 Second Step: Sharing Profits with Retailers

Resale price maintenance. Consider now that the incumbent sets both its wholesale price w and the retail price, and that it sets the retail price at $p_M = \frac{11}{20}$. Retailers can no longer undercut each other, so they split the monopoly output $q_M = \frac{9}{20}$, and each retailer earns

$$T = \frac{(p_M - w)\, q_M}{n} \tag{6.5}$$

per period, while the incumbent earns $(w - c)\, q_M$. The two add up to π_M for every w. By lowering w from p_M toward its cost c, the incumbent shifts profit to retailers without changing the retail price. A lump-sum payment of T per retailer, such as a slotting fee or a loyalty payment, together with $w = p_M$, delivers the same profits. The paper shows that exclusive territories, which turn each retailer into a local monopolist, also work this way.

After entry. In state C, no manufacturer shares profits with retailers. A lump-sum payment in the absorbing state changes neither sales nor the state, so no manufacturer makes it. A manufacturer that imposes a retail price above its rival's loses all sales, so the two manufacturers' retail prices fall to the incumbent's cost, and the outcome is the one in (6.3), with zero profit for retailers.

Result 2 (profit sharing). With resale price maintenance:

- in state M, a wholesale price $w \in [c, p_M]$ and the retail price p_M give each retailer T in (6.5) per period and leave the retail price and industry profit unchanged;
- the largest transfer per retailer, reached at $w = c$, is

$$\bar{T} = \frac{\pi_M}{n} = \frac{81}{400\,n}; \tag{6.6}$$

- in state C, retailers earn zero.

Figure 6.2 depicts the largest transfer $\bar{T}$ in (6.6) as a function of the number of retailers, together with the entrant's per-period profit after entry, $\pi_{EC} = \frac{9}{100}$, which the third step uses. Treating n as continuous, the two cross at $n = \frac{9}{4}$; with three retailers, $\bar{T} = \frac{27}{400}$ lies below $\frac{9}{100}$.

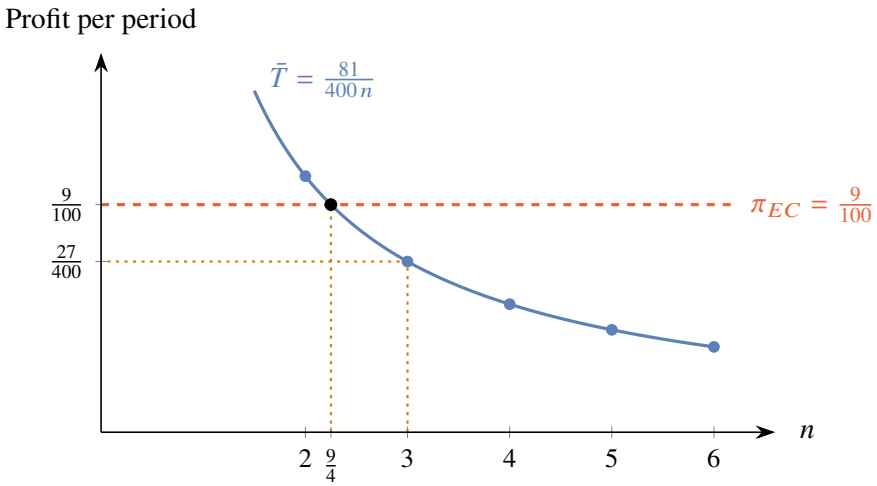


Figure 6.2: Largest per-period transfer to each retailer under resale price maintenance and the entrant's per-period profit after entry, as functions of the number of retailers.

Intuitively, resale price maintenance removes price competition between retailers, which creates a retail margin. This margin is a share of the incumbent's monopoly profit, and it lasts only while the incumbent is the only manufacturer: after entry, competition between manufacturers drives every margin down to zero.

6.5 Third Step: When Retailers Refuse to Carry the Entrant

A retailer's choice. Consider a stationary equilibrium in which, in every period in state M, the incumbent gives each retailer T, and consider that all other retailers reject the entrant. If retailer r also rejects, it receives T in this period and in every later period. If it accepts, it receives T_r from the entrant, but competition starts at once, so by Result 2 it receives nothing from then on. Retailer r therefore rejects if and only if

$$\frac{T}{1-\delta} \;\geq\; T_r. \tag{6.7}$$

The entrant's best offer. The entrant needs only one retailer, so it makes an offer to a single retailer, and the most it can offer is its value from entering, V_E in (6.4). Every retailer rejects this offer if (6.7) holds with $T_r = V_E$, that is, if T is at least

$$T^* = (1-\delta)\,V_E = \pi_{EC} - (1-\delta)F = \frac{9}{100} - (1-\delta)F. \tag{6.8}$$

The incumbent's choice. Paying T^* to every retailer in every period gives the incumbent

$$\frac{\pi_M - nT^*}{1-\delta},$$

while letting the entrant in gives it $\pi_{IC} = 0$ forever. The incumbent pays T^* if and only if $nT^* \le \pi_M$, that is, if and only if $T^* \le \bar{T}$ in (6.6), which gives the condition for exclusion

$$\frac{\pi_M}{n} \;\ge\; \pi_{EC} - (1-\delta)F \qquad \Longleftrightarrow \qquad \frac{81}{400\,n} \;\ge\; \frac{9}{100} - (1-\delta)F. \tag{6.9}$$

Dividing by $1-\delta$ and using $\pi_{IC} = 0$, condition (6.9) reads

$$\frac{\pi_M - \pi_{IC}}{n(1-\delta)} \;\ge\; \frac{\pi_{EC}}{1-\delta} - F,$$

which is the paper's condition for an exclusionary equilibrium.

Two examples. With $n = 2$ and $F = 0$, the incumbent pays each retailer $T^* = \frac{9}{100}$ and keeps

$$\frac{81}{400} - \frac{72}{400} = \frac{9}{400}$$

per period. Under resale price maintenance, it does so with $w = \frac{3}{20}$, since

$$\left(\frac{11}{20} - \frac{3}{20}\right)\frac{9}{20} = \frac{18}{100} = 2 \cdot \frac{9}{100}.$$

With $n = 3$, $\delta = \frac{1}{2}$, and $F = \frac{1}{10}$, the entrant's value is

$$V_E = \frac{18}{100} - \frac{1}{10} = \frac{2}{25},$$

so $T^* = \frac{1}{25}$; the incumbent pays $\frac{3}{25}$ in total, through $w = \frac{17}{60}$, and keeps $\frac{33}{400}$ per period.

Result 3 (exclusion). An equilibrium in which the entrant never enters exists if and only if condition (6.9) holds. In this equilibrium, the incumbent gives each retailer T^* in (6.8) per period, and the retail price stays at $\frac{11}{20}$ forever. Moreover:

- with $F = 0$, exclusion occurs if and only if $n \le \frac{9}{4}$, that is, with $n = 2$ retailers;
- with $n = 3$, exclusion occurs if and only if $F \ge \frac{9}{400(1-\delta)}$, while entry is profitable if and only if $F < \frac{9}{100(1-\delta)}$;

- condition (6.9) becomes harder to satisfy as n increases, as F decreases, and, when $F > 0$, as δ increases;
- if $F < \frac{9}{100(1-\delta)}$, an equilibrium in which the entrant enters in the first period also exists.

Figure 6.3 depicts, for $n = 3$, the pairs (δ, F) at which the entrant is excluded (shaded): above the lower blue curve, $F = \frac{9}{400(1-\delta)}$, and below the red curve, $F = \frac{9}{100(1-\delta)}$, where entry stops being profitable. At $\delta = \frac{1}{2}$, exclusion occurs for F between $\frac{9}{200}$ and $\frac{9}{50}$, and the example $F = \frac{1}{10}$ (point X) lies inside this interval. With a fourth retailer, the lower boundary rises to the green curve, $F = \frac{63}{1600(1-\delta)}$.

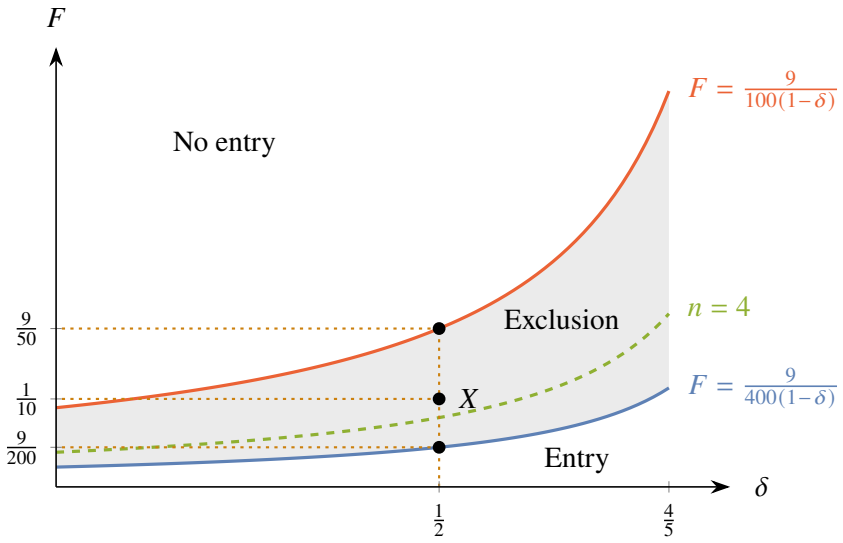


Figure 6.3: Exclusion region with $n = 3$ retailers (shaded). Below the blue curve, a retailer accepts the entrant's best offer; above the red curve, entry is not profitable. The dashed green curve is the lower boundary with $n = 4$ retailers.

Intuitively, the condition for exclusion resembles the condition for sustaining collusion with grim-trigger strategies. The stream of transfers plays the role of the cartel profit, and carrying the entrant is a deviation punished by the loss of that stream. The entrant can offer a retailer at most what it expects to earn, and since competition lowers industry profit, this amount can fall short of the share of monopoly profit that the incumbent hands to each retailer. Each additional retailer must also be paid, which makes exclusion harder. The entry cost lowers what the entrant can offer; since it is paid once while profits accrue every period, more patient firms value the entrant's profits more relative to F, and exclusion becomes harder. With $F = 0$, patience does not matter, because the retailer's loss and the entrant's gain are both flows.

Which markets? Keeping $F = 0$, let the incumbent's cost c vary in $(0, \frac{1}{2}]$, so that the entrant still sets a wholesale price equal to c after entry. Using (6.2) and (6.3), condition (6.9) becomes

$$\frac{(1-c)^2}{4n} \;\geq\; c\,(1-c),$$

that is,

$$n \;\leq\; \bar{n}(c) \equiv \frac{1-c}{4c}. \tag{6.10}$$

Result 4 (cost advantage and exclusion). With $F = 0$ and $c \in (0, \frac{1}{2}]$:

- exclusion occurs if and only if $n \leq \bar{n}(c)$ in (6.10), and $\bar{n}(c)$ decreases in c, with $\bar{n}(\frac{1}{10}) = \frac{9}{4}$;
- with $n \geq 2$ retailers, exclusion requires $c \leq \frac{1}{9}$;
- for $c > \frac{1}{5}$, entry raises industry profit, and $\bar{n}(c) < 1$.

Figure 6.4 depicts $\bar{n}(c)$. Exclusion occurs in the shaded area below the curve. The curve crosses $n = 2$ at $c = \frac{1}{9}$ and $n = 1$ at $c = \frac{1}{5}$, where $\pi_M = \pi_{EC} = \frac{4}{25}$.

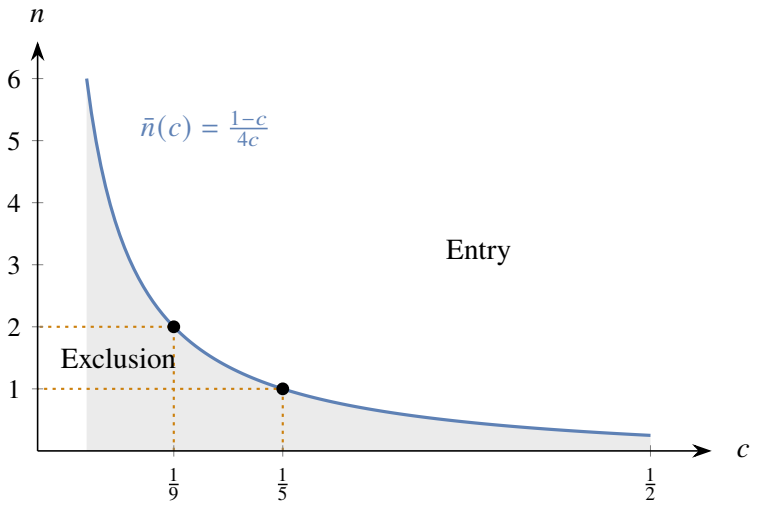


Figure 6.4: Largest number of retailers at which exclusion occurs ($F = 0$), as a function of the incumbent's marginal cost; the entrant's marginal cost is zero.

Intuitively, exclusion needs competition to destroy industry profit. When the entrant's cost advantage is small, it can only charge a price close to its cost after entry, and its profit is a small fraction of the monopoly profit that the incumbent can share. When its cost advantage is large, the entrant earns almost as much as a monopolist, and it can outbid the incumbent for a retailer. Exclusion therefore keeps out entrants that are more efficient than the incumbent, but not much more.

6.6 Fourth Step: Welfare and a Ban on Vertical Practices

Surplus in each state. Consumer surplus is $\frac{q^2}{2}$ under the linear demand (6.1). Before entry, consumer surplus and total surplus per period are

$$\frac{(9/20)^2}{2} = \frac{81}{800} \qquad \text{and} \qquad \frac{81}{800} + \frac{81}{400} = \frac{243}{800},$$

however the incumbent splits its profit with retailers. After entry, they are

$$\frac{(9/10)^2}{2} = \frac{324}{800} \qquad \text{and} \qquad \frac{324}{800} + \frac{72}{800} = \frac{396}{800},$$

with the entry cost F paid once.

A ban. Consider that an antitrust authority bans every practice that lets a manufacturer share profits with retailers: resale price maintenance, slotting fees, loyalty payments, and exclusive territories. Then only wholesale prices remain, and Result 1 applies: the entrant enters in the first period. Discounted total surplus is

$$\frac{396}{800(1-\delta)} - F \text{ with the ban} \qquad \text{and} \qquad \frac{243}{800(1-\delta)} \text{ under exclusion.}$$

Result 5 (a ban on profit sharing). Consider that condition (6.9) holds and that $F < \frac{9}{100(1-\delta)}$. A ban on practices that share profits with retailers:

- raises consumer surplus by $\frac{243}{800}$ per period, from $\frac{81}{800}$ to $\frac{324}{800}$;
- lowers the per-period profits of the incumbent and the retailers to zero;
- raises discounted total surplus if and only if $F \leq \frac{153}{800(1-\delta)}$, which holds whenever entry is profitable, since

$$\frac{9}{100(1-\delta)} = \frac{72}{800(1-\delta)}.$$

Table 6.1 reports per-period surplus under the two examples of the third step and under the ban. Under exclusion, the split between the incumbent and the retailers depends on T^*, but total surplus does not.

Table 6.1: Per-period surplus under exclusion and under a ban ($q = 1 - p$, $c = \frac{1}{10}$). Under the ban, the entrant also pays the entry cost F once.

Surplus	Exclusion, $n = 2$, $F = 0$	Exclusion, $n = 3$, $\delta = \frac{1}{2}$, $F = \frac{1}{10}$	Ban (entry)
Consumers	$\frac{81}{800}$	$\frac{81}{800}$	$\frac{324}{800}$
Incumbent	$\frac{18}{800}$	$\frac{66}{800}$	0
Retailers (all)	$\frac{144}{800}$	$\frac{96}{800}$	0
Entrant	0	0	$\frac{72}{800}$
Total	$\frac{243}{800}$	$\frac{243}{800}$	$\frac{396}{800}$

Intuitively, exclusion causes two losses. Consumers face the monopoly price instead of the competitive one, and the market is served by the high-cost producer instead of the efficient one. Exclusion survives even though entry raises total surplus by $\frac{153}{800}$ per period, because this gain goes to consumers, who take no part in the retailers' decision: consumer surplus rises by $\frac{243}{800}$, while industry profit falls from $\frac{162}{800}$ to $\frac{72}{800}$.

Policy screens. The toy model suggests the screens that the paper proposes for assessing a vertical practice. First, entry must require the cooperation of a retailer that receives profits from the incumbent; if the entrant could integrate into retailing, or retail entry were free, exclusion would fail. Second, entry must lower industry profit, which holds here only if $c < \frac{1}{5}$ (Result 4). Third,

the profits that retailers receive must disappear after entry, as in Result 2; a practice whose profits for retailers would survive entry gives them no reason to refuse the entrant. Finally, more retailers make exclusion harder (Result 3), so an incumbent gains from having fewer retailers to pay. Since both the incumbent and the retailers gain under exclusion, which side first proposes the practice says little about its purpose. A ban on a single practice, such as resale price maintenance, is not enough when slotting fees or loyalty payments can transfer the same profits.

6.7 Fifth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Exclusion without any exclusivity clause: a retailer that carries the entrant breaches no contract, and refuses only because entry would end the profits that the incumbent shares with it.
- The paper's condition for exclusion, which compares each retailer's share of the profit that entry destroys with the entrant's value from entering, and its comparative statics in the number of retailers and the entry cost.
- The equivalence between resale price maintenance and lump-sum transfers, the coexistence of an exclusionary equilibrium with one in which entry occurs, and the welfare loss from excluding a more efficient producer.

What it drops.

- General demand, differentiated products, and an incumbent that keeps positive profits after entry; with $\pi_{IC} > 0$, only the profit that entry destroys, $\pi_M - \pi_{IC}$, can be shared with retailers.
- Other post-entry conduct, such as collusion between manufacturers, which the paper discusses and which changes π_{IC} and π_{EC}.
- Retail services. Klein and Murphy (1988) argue that retail margins can reward retailers for providing services; with services, a restraint can raise demand as well as exclude, and its welfare effect involves a trade-off that the toy model does not have.
- Explicit exclusive dealing contracts, as in Rasmusen et al. (1991) and Segal and Whinston (2000), and offers by the entrant that depend on what other retailers do.

6.8 Connections to Other Chapters

- Chapter 5 shows that off-list-price rebates, which do not commit buyers, cannot keep a more efficient entrant out. In this chapter, vertical practices exclude because retailers share the incumbent's profit.
- Chapter 7 studies another way in which an incumbent excludes a more efficient rival: refusing to supply it, so that it cannot cover its entry cost.

References

Asker, J., and Bar-Isaac, H. (2014). Raising retailers' profits: On vertical practices and the exclusion of rivals. *American Economic Review*, 104(2), 672–686.

Ide, E., Montero, J.-P., and Figueroa, N. (2016). Discounts as a barrier to entry. *American Economic Review*, 106(7), 1849–1877.

Klein, B., and Murphy, K. M. (1988). Vertical restraints as contract enforcement mechanisms. *Journal of Law and Economics*, 31(2), 265–297.

Rasmusen, E. B., Ramseyer, J. M., and Wiley, J. S., Jr. (1991). Naked exclusion. *American Economic Review*, 81(5), 1137–1145.

Segal, I. R., and Whinston, M. D. (2000). Naked exclusion: Comment. *American Economic Review*, 90(1), 296–309.

Chapter 7

Refusing to Supply a More Efficient Rival: Dynamic Vertical Foreclosure

Companion note to Fumagalli, C., and M. Motta (2020), "Dynamic Vertical Foreclosure," Journal of Law and Economics 63(4): 763–812.

7.1 Introduction

Some firms own an essential input and also sell the final product. Should such a firm refuse to supply the input to a downstream rival that is more efficient than its own downstream unit? The Chicago School argued that it should not: selling the input to the rival, and capturing part of the rival's efficiency gains, earns more than serving the market itself at a higher cost. Most theories of vertical foreclosure, such as Ordover et al. (1990) and Hart and Tirole (1990), therefore look for static reasons why the input owner cannot capture enough of these gains.

Fumagalli and Motta (2020) propose a dynamic reason. A downstream rival needs the profits it earns early on, its critical profits, to cover its entry cost. A refusal to supply today removes these profits and can keep the rival out of the market altogether. The input owner, the incumbent, gives up profits today, but it keeps the downstream market to itself tomorrow, when a more efficient upstream rival may enter. As the only buyer of the new input, the incumbent then captures part of the upstream entrant's efficiency gains. When the upstream entrant needs downstream rivals to sell its input, refusing to supply also keeps it out.

A test that looks only at the current period would therefore miss the reason for refusing to supply. Consumers lose from foreclosure in every period, while the effect on total surplus depends on the entry costs that it saves. The paper relates the argument to the European Commission's *Telefónica* broadband case and to the *Genzyme* case in the United Kingdom.

The toy model uses linear demand, constant costs, and two periods. The first step reports the post-entry profits and the static benchmark, the second step shows how a refusal to supply keeps the rival out, the third and fourth steps find when foreclosure is profitable if upstream entry occurs

anyway and if it is deterred, the fifth step studies welfare and a duty to supply, and the sixth step lists what the toy model keeps and drops.

7.2 The Toy Model

Firms and technologies. A vertically integrated incumbent I owns the only upstream firm, which produces an input at marginal cost $c > 0$, and a downstream unit that turns one unit of the input into one unit of a homogeneous final good at marginal cost $\gamma > 0$. A downstream rival D can enter in period 1; it transforms the input at marginal cost 0, so γ is its efficiency advantage. An upstream rival U can enter in period 2; it produces the input at marginal cost 0, so c is its efficiency advantage. Entry costs $F_D \geq 0$ and $F_U \geq 0$ are sunk. Inputs are homogeneous, and downstream firms compete in prices. As in the paper, we assume $c + \gamma < \frac{1}{2}$, so that the entrants are not so efficient that they could charge the monopoly price while the incumbent is active. The paper allows D to have a positive cost γ_E; with linear demand, setting it to zero is a normalization, since every threshold in the paper follows from the ones below after measuring c and $\gamma = \gamma_I - \gamma_E$ relative to $1 - \gamma_E$.

Demand and monopoly profits. In each period, demand is $Q = 1 - p$. A monopolist with total unit cost k charges $\frac{1+k}{2}$, sells $\frac{1-k}{2}$, and earns

$$\pi^m(k) = \frac{(1-k)^2}{4}, \tag{7.1}$$

while consumer surplus and total surplus are

$$\frac{(1-k)^2}{8} \qquad \text{and} \qquad \frac{3(1-k)^2}{8}.$$

If a downstream firm with total cost 0 competes in prices against the incumbent's own chain, with total cost $c + \gamma < \frac{1}{2}$, it charges the limit price $c + \gamma$ and earns

$$\pi^d = (c + \gamma)(1 - c - \gamma). \tag{7.2}$$

Contracting. Firms trade the input through two-part tariffs, and the incumbent can commit to close its downstream unit. When the incumbent faces a single entrant, each of them makes a take-it-or-leave-it offer with probability $\frac{1}{2}$. If they do not trade, the incumbent serves the market alone and earns $\pi^m(c + \gamma)$; if they trade, they earn the monopoly profit with the better technology, and the firm that makes the offer keeps the whole gain. In expectation, each firm obtains half of the gain from trade. When both entrants are active, the paper shows that offers that can depend on whether the partner trades exclusively sustain the largest industry profit, $\pi^m(0) = \frac{1}{4}$: D buys from U and sells at price $\frac{1}{2}$, and the incumbent is paid to keep its downstream unit out. Among these outcomes, the paper selects the one that is best for the incumbent, which is the least favorable for foreclosure. The incumbent earns its marginal contribution, $\frac{1}{4} - \pi^d$, the profit that D and U would lose if the incumbent's chain competed against them, and D and U share the remaining π^d. The toy model takes this three-firm outcome from the paper.

Timing. Figure 7.1 depicts the timing. In period 0, the incumbent decides whether to refuse to supply D; the refusal binds through period 1. Profits are not discounted. We write $\pi_I(D,U)$, $\Pi_D(D,U)$, and $\Pi_U(D,U)$ for the profits of the incumbent and the entrants, gross of entry costs, when both entrants are active, and use $(D,\emptyset)$, $(\emptyset,U)$, and $(\emptyset,\emptyset)$ when only D, only U, or neither is active.

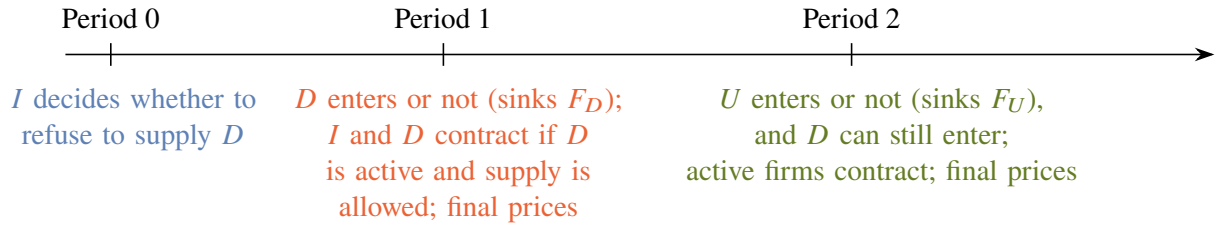


Figure 7.1: Timing of the game.

7.3 First Step: Post-Entry Profits and the Static Benchmark

Consider first each market configuration in a single period. With no entrant, the incumbent is a monopolist with cost $c+\gamma$. With only D, the incumbent sells its input to D and closes its downstream unit; D sells at $\frac{1+c}{2}$, and the two firms split equally the gain $\pi^m(c)-\pi^m(c+\gamma)$. With only U, the incumbent buys the input from U and sells at $\frac{1+\gamma}{2}$, and the two firms split equally the gain $\pi^m(\gamma)-\pi^m(c+\gamma)$. With both entrants, the price is $\frac{1}{2}$, the incumbent earns $\frac{1}{4}-\pi^d$, and D and U earn

$$\Pi_D(D,U)=\frac{\pi^d-\left[\pi^m(\gamma)-\pi^m(c)\right]}{2} \qquad \text{and} \qquad \Pi_U(D,U)=\frac{\pi^d+\left[\pi^m(\gamma)-\pi^m(c)\right]}{2}, \tag{7.3}$$

where $\pi^m(\gamma)-\pi^m(c)$ is the gain that U would create by trading alone with the incumbent minus the gain that D would create. Table 7.1 reports these profits after substituting (7.1) and (7.2).

Table 7.1: Post-entry profits in one period, gross of entry costs.

Active entrants	Incumbent	Downstream rival D	Upstream rival U
None, $(\emptyset,\emptyset)$	$\frac{(1-c-\gamma)^2}{4}$	0	0
Only D, $(D,\emptyset)$	$\frac{(1-c)^2+(1-c-\gamma)^2}{8}$	$\frac{\gamma(2-2c-\gamma)}{8}$	0
Only U, $(\emptyset,U)$	$\frac{(1-\gamma)^2+(1-c-\gamma)^2}{8}$	0	$\frac{c(2-c-2\gamma)}{8}$
Both, (D,U)	$\frac{1}{4}-(c+\gamma)(1-c-\gamma)$	$\frac{2c+6\gamma-3c^2-8c\gamma-5\gamma^2}{8}$	$\frac{6c+2\gamma-5c^2-8c\gamma-3\gamma^2}{8}$

Result 1 (post-entry profits and the Chicago benchmark). For every $c,\gamma>0$ with $c+\gamma<\frac{1}{2}$:

- in a single period, supplying D raises the incumbent's profit:

$$\pi_I(D,\emptyset) - \pi_I(\emptyset,\emptyset) = \frac{\gamma(2-2c-\gamma)}{8} > 0;$$

- the incumbent also gains when only U is active:

$$\pi_I(\emptyset,U) - \pi_I(\emptyset,\emptyset) = \frac{c(2-c-2\gamma)}{8} > 0;$$

- the incumbent loses when both entrants are active:

$$\pi_I(\emptyset,\emptyset) - \pi_I(D,U) = \frac{(c+\gamma)(2-3c-3\gamma)}{4} > 0;$$

- each entrant earns more when the other entrant is also active: $\Pi_D(D,U) > \Pi_D(D,\emptyset)$ and $\Pi_U(D,U) > \Pi_U(\emptyset,U)$.

With $c = \frac{1}{10}$ and $\gamma = \frac{1}{5}$, the example used below, the incumbent earns $\frac{49}{400}$ alone, $\frac{13}{80}$ with D, $\frac{113}{800}$ with U, and $\frac{1}{25}$ with both; D earns $\frac{1}{25}$ alone and $\frac{101}{800}$ with U; and U earns $\frac{3}{160}$ alone and $\frac{67}{800}$ with D.

Intuitively, the first bullet is the Chicago argument: in a static market, the incumbent's downstream unit only serves as its outside option, and dealing with a more efficient rival creates a gain that the incumbent shares. The same logic applies to a more efficient upstream entrant. When both entrants are active, however, they can trade with each other, and the incumbent is left with the value of keeping its own chain out of the market. Each entrant, in turn, benefits from the presence of the other, because it no longer depends on a single trading partner that owns a competing technology.

7.4 Second Step: Critical Profits and the Ability to Exclude

Entry costs. As in the paper, entry costs are such that each entrant enters when the incumbent supplies D:

$$F_U \le \Pi_U(D,U) \qquad \text{and} \qquad F_D \le \overline{F}_D \equiv \Pi_D(D,\emptyset) + \Pi_D(D,U) = \frac{2c+8\gamma-3c^2-10c\gamma-6\gamma^2}{8}. \tag{7.4}$$

Period 2. If D entered in period 1, U earns $\Pi_U(D,U) \ge F_U$ and enters. If D did not enter, D can still enter in period 2, but it then earns at most $\Pi_D(D,U)$, its profit when U is active as well.

Period 1. Without a refusal to supply, D earns

$$\Pi_D(D,\emptyset) + \Pi_D(D,U) \ge F_D$$

by entering in period 1, more than by waiting. With a refusal to supply, D earns nothing in period 1, so it enters only if its period-2 profit alone covers the entry cost.

Result 2 (ability to exclude). Under (7.4):

- without a refusal to supply, D enters in period 1 and U enters in period 2;
- with a refusal to supply, D stays out in both periods if and only if $F_D > \Pi_D(D, U)$;
- when D stays out, U enters in period 2 if $F_U \leq \Pi_U(\emptyset, U)$ and stays out if $\Pi_U(\emptyset, U) < F_U \leq \Pi_U(D, U)$.

In the example, a refusal to supply excludes D if and only if $\frac{101}{800} < F_D \leq \frac{133}{800}$. When it does, U still enters if $F_U \leq \frac{3}{160}$, and is kept out if $\frac{3}{160} < F_U \leq \frac{67}{800}$.

Intuitively, the entry cost of D is a fixed cost that it can recover only over two periods. A refusal to supply removes the first period, so the rival lacks the critical profits it needs, and it does not enter at all; a financial constraint that requires early profits to fund a later investment would work in the same way. A margin squeeze, a high input price combined with a low retail price, would have the same effect on the rival's period-1 profit. The refusal also changes period 2: without D, the upstream entrant has a single buyer, the incumbent, and earns $\Pi_U(\emptyset, U)$ instead of $\Pi_U(D, U)$, which deters it when its entry cost is high.

7.5 Third Step: Foreclosure that Protects the Downstream Monopoly

Consider that refusal to supply excludes D, $F_D \in \left(\Pi_D(D, U), \overline{F}_D\right]$, and that U enters anyway, $F_U \leq \Pi_U(\emptyset, U)$. Supplying D and refusing yield, respectively,

$$\pi_I(D, \emptyset) + \pi_I(D, U) \qquad \text{and} \qquad \pi_I(\emptyset, \emptyset) + \pi_I(\emptyset, U).$$

Define the incumbent's period-1 sacrifice and its period-2 gain from refusing,

$$\begin{aligned} S &\equiv \pi_I(D, \emptyset) - \pi_I(\emptyset, \emptyset) = \frac{\gamma(2 - 2c - \gamma)}{8}, \\ G_U &\equiv \pi_I(\emptyset, U) - \pi_I(D, U) = \frac{4\gamma + 6c - 7c^2 - 14c\gamma - 6\gamma^2}{8}. \end{aligned} \tag{7.5}$$

Refusal to supply changes the incumbent's two-period profit by

$$\Delta_U \equiv G_U - S = \frac{2\gamma + 6c - 7c^2 - 12c\gamma - 5\gamma^2}{8}, \tag{7.6}$$

and define the cutoff

$$\gamma_P(c) \equiv \frac{1 - 6c + \sqrt{c^2 + 18c + 1}}{5}. \tag{7.7}$$

Result 3 (refusal to protect the downstream monopoly). When refusal to supply excludes D and U enters anyway, the incumbent refuses to supply if and only if $\Delta_U \geq 0$. Moreover:

- refusal is profitable whenever $\gamma \leq c$;

- if $c \geq \frac{1}{12}$, refusal is profitable for every $\gamma < \frac{1}{2} - c$;
- if $c < \frac{1}{12}$, refusal is profitable if and only if $\gamma \leq \gamma_P(c)$, where $\gamma_P(c) < \frac{1}{2} - c$.

Table 7.2 reports the incumbent's profit in each period in the example, where

$$S = \frac{1}{25}, \qquad G_U = \frac{81}{800}, \qquad \text{and} \qquad \Delta_U = \frac{49}{800} > 0.$$

With $c = \frac{1}{20}$, instead, the cutoff is $\gamma_P(\frac{1}{20}) \approx 0.416$, and with $c = \frac{1}{40}$ and $\gamma = \frac{9}{20}$ refusal is not profitable.

Table 7.2: The incumbent's profit along each path, with $c = \frac{1}{10}$ and $\gamma = \frac{1}{5}$.

Incumbent's choice and outcome	Period 1	Period 2	Total
Supply D: D enters, then U enters	$\frac{13}{80}$	$\frac{1}{25}$	$\frac{81}{400}$
Refuse: D stays out, U enters anyway	$\frac{49}{400}$	$\frac{113}{800}$	$\frac{211}{800}$
Refuse: D and U stay out	$\frac{49}{400}$	$\frac{49}{400}$	$\frac{49}{200}$

Figure 7.2 depicts the sacrifice S and the gain G_U in (7.5) as functions of D's efficiency advantage γ, for $c = \frac{1}{20}$. The sacrifice rises with γ, while the gain first rises and then falls. Refusal to supply is profitable to the left of the crossing point, $\gamma_P(\frac{1}{20}) \approx 0.416$, and unprofitable between this point and the largest admissible advantage, $\frac{1}{2} - c = \frac{9}{20}$.

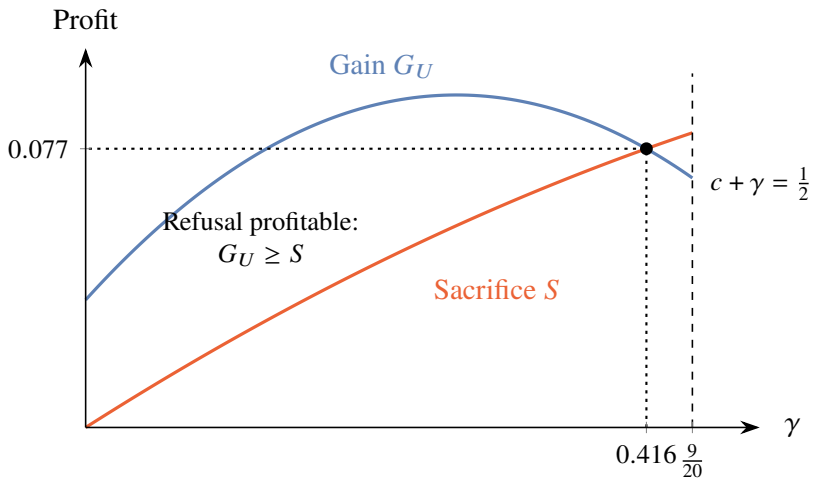


Figure 7.2: Period-1 sacrifice and period-2 gain from refusing to supply when the upstream rival enters anyway ($c = \frac{1}{20}$).

Intuitively, refusing to supply has a cost today and a benefit tomorrow. The cost is the incumbent's share of D's efficiency gain in period 1, which is small when D is only slightly more efficient than the incumbent's own unit. The benefit comes from period 2. Without D, the incumbent is the only buyer of U's input and captures half of U's efficiency gain. With D in the market, D competes for

U's input, and the incumbent is reduced to its marginal contribution, $\frac{1}{4} - \pi^d$. The incumbent's half of U's gain,

$$\frac{1}{2}\left[\pi^m(\gamma) - \pi^m(c+\gamma)\right],$$

grows with the upstream advantage c, and when $c \geq \frac{1}{12}$ the benefit exceeds the cost for every admissible γ. In a static market the incumbent never refuses (Result 1); the reason to refuse is the change in market structure that the refusal brings about in period 2.

7.6 Fourth Step: Foreclosure that Protects Both Monopolies

Consider now that refusal to supply excludes D and that U's entry cost is high, $\Pi_U(\emptyset, U) < F_U \leq \Pi_U(D, U)$, so that U enters only if D does. Refusing now yields $\pi_I(\emptyset, \emptyset)$ in both periods. The period-2 gain becomes

$$G_0 \equiv \pi_I(\emptyset, \emptyset) - \pi_I(D, U) = \frac{(c+\gamma)(2-3c-3\gamma)}{4},$$

and refusal changes the incumbent's two-period profit by

$$\Delta_0 \equiv G_0 - S = \frac{2\gamma + 4c - 6c^2 - 10c\gamma - 5\gamma^2}{8} = \Delta_U - \frac{c(2-c-2\gamma)}{8}. \tag{7.8}$$

Define the cutoffs

$$c_{PP} \equiv 1 - \frac{\sqrt{3}}{2} \approx 0.134 \qquad \text{and} \qquad \gamma_{PP}(c) \equiv \frac{1 - 5c + \sqrt{1 + 10c - 5c^2}}{5}. \tag{7.9}$$

Result 4 (refusal to protect both monopolies). When refusal to supply excludes both D and U, the incumbent refuses to supply if and only if $\Delta_0 \geq 0$. Moreover:

- $\Delta_0 < \Delta_U$, and the difference is the incumbent's gain from dealing with U alone,

$$\pi_I(\emptyset, U) - \pi_I(\emptyset, \emptyset) = \frac{c(2-c-2\gamma)}{8};$$

- if $c \geq c_{PP}$, refusal is profitable for every $\gamma < \frac{1}{2} - c$;
- if $c < c_{PP}$, refusal is profitable if and only if $\gamma \leq \gamma_{PP}(c)$, where $\gamma_{PP}(c) < \frac{1}{2} - c$, and $\gamma_{PP}(c) < \gamma_P(c)$ for $c < \frac{1}{12}$.

In the example, $\Delta_0 = \frac{17}{400}$, below $\Delta_U = \frac{49}{800}$ but still positive (Table 7.2). With $c = \frac{1}{20}$, the cutoff falls from $\gamma_P \approx 0.416$ to $\gamma_{PP} \approx 0.394$.

Figure 7.3 depicts the parameter regions of Results 3 and 4 in the (c, γ) plane. Below the red curve $\gamma_{PP}(c)$, refusal is profitable whether or not it also deters U; between the red curve and the blue curve $\gamma_P(c)$, or the boundary $c + \gamma = \frac{1}{2}$ when $c \geq \frac{1}{12}$, it is profitable only when U enters anyway; and above the blue curve, refusal is not profitable. Both curves start at $\gamma = \frac{2}{5}$ when c is close to

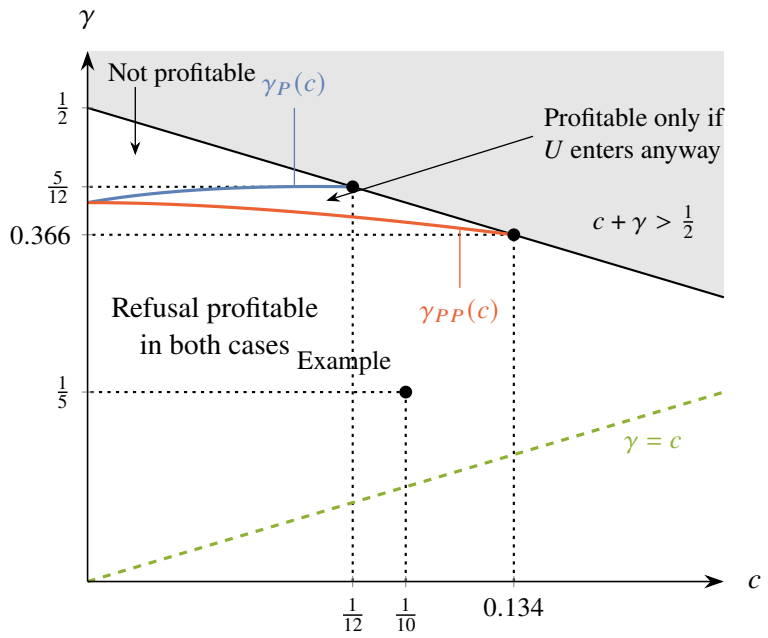


Figure 7.3: Parameter regions where refusal to supply is profitable. The shaded area violates $c + \gamma < \frac{1}{2}$.

zero, and reach the boundary $c + \gamma = \frac{1}{2}$ at $c = \frac{1}{12}$ and at $c = c_{PP}$, respectively. The green dashed line is $\gamma = c$, below which Result 3 holds without computing $\gamma_P(c)$. The dot marks the example.

Intuitively, once D is out of the market, the incumbent would rather see U enter, since it could then capture part of U's efficiency gain. Deterring U therefore gives up this gain, and the incentive to refuse is weaker than in the third step. Refusal can remain profitable because a monopoly in both markets yields more than facing efficient rivals in both, and the sacrifice in period 1 is small when γ is small. This case is close to Carlton and Waldman (2002), where tying deters entry into a complementary market and, through it, later entry into the primary market. The case of the third step has no counterpart there: foreclosure can be profitable even when upstream entry cannot be prevented, and the incentive is then stronger.

7.7 Fifth Step: Welfare and a Duty to Supply

Without a refusal to supply, the final price is $\frac{1+c}{2}$ in period 1 and $\frac{1}{2}$ in period 2. With a refusal, the incumbent's chain serves the market at $\frac{1+c+\gamma}{2}$ in period 1; in period 2 the price is $\frac{1+\gamma}{2}$ if U enters and $\frac{1+c+\gamma}{2}$ if it does not. Every market is served by a monopolist, so total surplus in a period with unit cost k is $\frac{3(1-k)^2}{8}$. Define the entry-cost thresholds

$$F_W \equiv \frac{3\gamma(2-c-\gamma)}{4} \qquad \text{and} \qquad F_{W+U} \equiv \frac{3(2c+4\gamma-c^2-4c\gamma-2\gamma^2)}{8}. \tag{7.10}$$

Result 5 (welfare). When refusal to supply occurs in equilibrium:

- consumers lose in both periods: the period-1 price rises from $\frac{1+c}{2}$ to $\frac{1+c+\gamma}{2}$, and the period-2 price rises from $\frac{1}{2}$ to $\frac{1+\gamma}{2}$ if U enters anyway and to $\frac{1+c+\gamma}{2}$ if it does not;
- if U enters anyway, refusal lowers total surplus if $F_D < F_W$ and raises it if $F_D > F_W$;
- if U is kept out, refusal lowers total surplus if $F_D + F_U < F_{W+U}$ and raises it if $F_D + F_U > F_{W+U}$.

In the example, consumer surplus falls by $\frac{17}{200}$ when U enters anyway and by $\frac{83}{800}$ when it does not. Since $F_W = \frac{51}{200}$ exceeds $\overline{F}_D = \frac{133}{800}$, and $F_{W+U} = \frac{249}{800}$ exceeds the largest admissible $F_D + F_U$, which is

$$\frac{133}{800} + \frac{67}{800} = \frac{1}{4},$$

refusal lowers total surplus for every admissible entry cost.

Figure 7.4 depicts, for $c = \frac{1}{10}$ and small values of γ, the interval of entry costs at which refusal to supply excludes D, between $\Pi_D(D, U)$ and $\overline{F}_D$ (shaded), together with the welfare threshold F_W. Since $c \geq \frac{1}{12}$, refusal is profitable at every point of the shaded interval when U enters anyway (Result 3). The threshold F_W lies below the interval, and refusal raises total surplus whenever it occurs, for

$$\gamma < 3 + c - \sqrt{9 + 4c + 4c^2} \approx 0.028;$$

it lies above the interval, and refusal lowers total surplus, for

$$\gamma > \frac{c(2 - 3c)}{4(1 + c)} = \frac{17}{440} \approx 0.039;$$

in between, the answer depends on F_D.

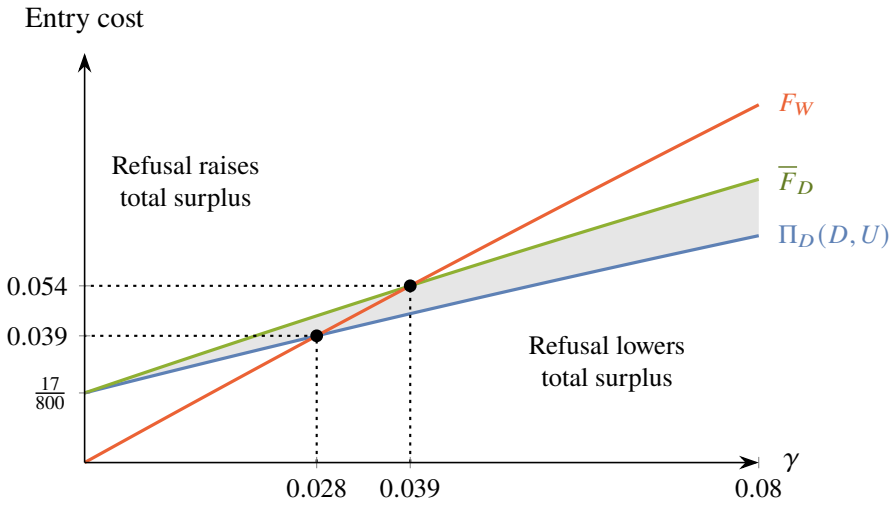


Figure 7.4: Entry costs at which refusal to supply excludes the downstream rival (shaded) and the welfare threshold F_W, for $c = \frac{1}{10}$.

Intuitively, a refusal to supply replaces the efficient technologies with the incumbent's own in both periods, so consumers always lose. Total surplus also counts the entry costs that foreclosure saves. When D's advantage is small and its entry cost is high, entry adds little surplus at a large cost,

the excess-entry problem of Mankiw and Whinston (1986), and foreclosure can raise total surplus. When D's advantage is larger, foreclosure lowers total surplus for every entry cost at which it excludes.

A duty to supply. An obligation to supply D in period 1 restores the path without refusal: D enters in period 1 and U in period 2. In the example, it raises consumer surplus by $\frac{17}{200}$ when U would have entered anyway, and raises total surplus by

$$F_W - F_D = \frac{51}{200} - F_D > 0.$$

Two features of the toy model matter for applying such a duty. First, a test that looks only at the current period would find that refusing lowers the incumbent's profit, and would conclude that the incumbent has no reason to refuse; the reason lies in how the market will evolve. As the paper argues, what matters is that upstream entry is likely in the future, whether it occurs anyway or depends on downstream entry. Second, a refusal to supply excludes only if it is a commitment: if the incumbent could resume supply once D has entered, D would anticipate it and enter, and refusal would not occur in equilibrium.

7.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The Chicago benchmark: in a static market, the incumbent prefers to supply a more efficient downstream rival.
- The dynamic rationale: refusing to supply deprives the rival of the period-1 profits it needs to cover its entry cost, and changes the market structure in period 2.
- The two motives, protecting the downstream monopoly to capture rents from an upstream entrant and protecting the monopoly in both markets, with a weaker incentive in the second; the role of the downstream efficiency gap γ and the upstream efficiency gap c; and the harm to consumers, with an ambiguous effect on total surplus.

What it drops.

- General demand, a general contracting game, and the derivation of the three-firm contracting outcome, which the toy model takes from the paper.
- Discounting: the paper and the toy model weigh both periods equally. A discount factor below one would lower the period-2 gain from refusal, and also the upper bound $\overline{F}_D$, which narrows the interval of entry costs at which refusal excludes.
- The paper's extensions: a refusal that binds forever, which makes foreclosure more likely; a lower probability that upstream firms make the offers, which strengthens the incentive to refuse; and vertically integrated entrants, which are harder to exclude.
- Alternatives to refusal: exclusive dealing between the incumbent and D, which deters U and makes refusal less likely, and long-term contracts with commitments on future input prices,

which, under strong commitment assumptions, let the incumbent capture the entrants' efficiency gains without excluding them.

7.9 Connections to Other Chapters

- Chapter 5 studies rebates that an incumbent offers to buyers, and shows that off-list-price rebates, which do not commit buyers, cannot exclude a more efficient entrant.
- In Chapter 6, an incumbent excludes an entrant by sharing its profit with the retailers that the entrant needs.
- Chapter 40 studies another strategy that removes a future competitor: buying an innovator to shut down its project.

References

Carlton, D. W., and Waldman, M. (2002). The strategic use of tying to preserve and create market power in evolving industries. *RAND Journal of Economics*, 33(2), 194–220.

Fumagalli, C., and Motta, M. (2020). Dynamic vertical foreclosure. *Journal of Law and Economics*, 63(4), 763–812.

Hart, O., and Tirole, J. (1990). Vertical integration and market foreclosure. *Brookings Papers on Economic Activity: Microeconomics*, 1990, 205–286.

Mankiw, N. G., and Whinston, M. D. (1986). Free entry and social inefficiency. *RAND Journal of Economics*, 17(1), 48–58.

Ordover, J. A., Saloner, G., and Salop, S. C. (1990). Equilibrium vertical foreclosure. *American Economic Review*, 80(1), 127–142.

Chapter 8

Wholesale Prices, Two-Part Tariffs, and Downstream Entry

Companion note to Milliou, C., and E. Petrakis (2024), "Vertical Contracts and Downstream Entry," Journal of Industrial Economics 72(1): 598–629.

8.1 Introduction

Many final products use an input that only one firm supplies. The supplier can charge a price per unit of the input, a wholesale price contract, or add a fixed fee to it, a two-part tariff. The standard view favors two-part tariffs: a unit price above cost adds a second markup to the final price, while a two-part tariff can price the input at cost and collect profit through the fee.

This view takes the number of downstream firms as given. When entry is free, the contract also changes how many firms enter. A wholesale price above cost shrinks the surplus that the supplier and an entrant create together; this *inefficiency effect* discourages entry. Without a fee, however, the supplier can collect only by raising the unit price, which shrinks that surplus further, so the entrant keeps more than its bargaining power would give it; this *distribution effect* encourages entry.

Milliou and Petrakis (2024) show that the distribution effect dominates, so more firms enter under wholesale price contracts. When the supplier's bargaining power is high, the extra entrants more than offset the lower output of each one, and aggregate output, consumer surplus, and welfare can be higher with wholesale prices. The supplier may prefer them, and it may gain by giving up part of its bargaining power. For competition policy, a ranking of contracts obtained with a fixed number of firms can reverse once entry responds.

The toy model fixes product substitutability at $\frac{1}{2}$, normalizes the demand intercept to one and costs to zero, and treats the number of entrants as continuous, as the paper does. The first step splits an entrant's gain from wholesale prices into the two effects. The second step fixes the number of entrants, the third step lets firms enter freely, the fourth to sixth steps compare output, welfare, and the supplier's profit, and the seventh step lists what the toy model keeps and drops.

8.2 The Toy Model

Firms and demand. An upstream monopolist, U, produces an input at marginal cost $c = 0$ and sells it to n downstream entrants, $D_1, \dots, D_n$. Each entrant turns one unit of the input into one unit of its own variety of the final good, at no further cost. The inverse demand for variety i is

$$p_i = 1 - q_i - \tfrac{1}{2} Q_{-i}, \tag{8.1}$$

where Q_{-i} is the total output of the other entrants. This is the linear demand of the paper, from Singh and Vives (1984), with demand intercept $a = 1$, input cost $c = 0$, and substitutability $\gamma = \frac{1}{2}$. It comes from a representative consumer with utility

$$\sum_{i=1}^{n} q_i - \frac{1}{2}\Big[\sum_{i=1}^{n} q_i^2 + \frac{1}{2}\sum_{i \neq j} q_i q_j\Big],$$

who values variety: the same total output spread over more varieties yields higher utility. When n entrants each produce q, consumer surplus, utility minus spending, is

$$CS = \frac{1}{2} n\Big[1 + \frac{n-1}{2}\Big] q^2 = \frac{n(n+1)\, q^2}{4}. \tag{8.2}$$

Contracts and bargaining. Under a *wholesale price contract* ($\mathcal{W}$), D_i pays w_i per unit of input. Under a *two-part tariff* ($\mathcal{T}$), it pays w_i per unit plus a fixed fee F_i. The terms are set by generalized Nash bargaining between U and each D_i, where β is the bargaining power of U and $1 - \beta$ that of D_i. The negotiations are simultaneous, separate, and secret: each pair takes the outcome of the other negotiations as given (Nash-in-Nash bargaining; Horn and Wolinsky, 1988; Collard-Wexler et al., 2019), and D_i never observes the terms of its rivals.

Timing. In the first stage, downstream firms decide whether to enter, paying the entry cost $K = \frac{1}{16}$. In the second stage, each entrant bargains with U over its contract. In the third stage, entrants choose quantities simultaneously.

Free entry. As in the paper, the number of entrants n is treated as a continuous variable, and firms enter until the gross profit of an entrant equals K. The paper writes $k \equiv \frac{\sqrt{K}}{a-c}$, which is $\frac{1}{4}$ in the toy model, and assumes $k \leq \frac{\sqrt{1-\beta}}{2}$ (its Assumption 1) so that at least one firm enters under either contract. In the toy model this condition reads

$$\beta \leq \frac{3}{4}, \tag{8.3}$$

and every result below considers $\beta \in [0, \frac{3}{4}]$. Under free entry, each entrant's profit net of the entry cost is zero, so welfare is consumer surplus plus the profit of U.

8.3 First Step: Bargaining between the Supplier and One Entrant

Quantities. In the third stage, D_i takes the output of its rivals as given. Define its *residual intercept* $x \equiv 1 - \frac{1}{2}Q_{-i}$, the price it would get for a first unit. Maximizing $(x - q_i - w_i)\, q_i$ gives

$$q_i = \frac{x - w_i}{2} \qquad \text{and} \qquad \pi_{D_i} = \frac{(x - w_i)^2}{4}, \tag{8.4}$$

where π_{D_i} is the operating profit of D_i, before any fee and before the entry cost. An entrant's *gross profit* is its operating profit minus any fee, before the entry cost.

The surplus of the pair. Since D_j does not observe the terms agreed with D_i, and the pair takes the other negotiations as given, the revenue of U from the other entrants is the same whether or not U and D_i agree. It cancels from the gain of U from an agreement. If the pair disagrees, D_i does not produce and its entry cost is sunk, so the pair (U, D_i) bargains over the surplus it creates, given x,

$$J(w_i) = w_i \frac{x - w_i}{2} + \frac{(x - w_i)^2}{4}, \tag{8.5}$$

the revenue of U from D_i plus the operating profit of D_i.

Two-part tariff. The pair chooses w_i to maximize J and uses the fee to divide it. Since $J'(w_i) = -\frac{w_i}{2}$, the surplus is maximized at $w_i = 0$, the marginal cost of the input, where $J = \frac{x^2}{4}$. Maximizing the Nash product $F_i^{\beta}\left(\frac{x^2}{4} - F_i\right)^{1-\beta}$ gives $F_i = \frac{\beta x^2}{4}$, so D_i keeps $(1-\beta)\frac{x^2}{4}$. The wholesale price equals marginal cost because of the opportunism problem of U under secret contracts (Hart and Tirole, 1990; O'Brien and Shaffer, 1992): it cannot commit not to offer a lower price to a rival, and it collects its share through the fee.

Wholesale price contract. Without a fee, the per-unit price must both create and divide the surplus. The pair maximizes the Nash product

$$\left[w_i \frac{x - w_i}{2}\right]^{\beta}\left[\frac{(x - w_i)^2}{4}\right]^{1-\beta},$$

whose logarithm has first-order condition $\frac{\beta}{w_i} = \frac{2-\beta}{x - w_i}$. Therefore

$$w_i = \frac{\beta x}{2}, \qquad q_i = \frac{(2-\beta)\, x}{4}, \qquad \pi_{D_i} = \frac{(2-\beta)^2 x^2}{16}, \qquad J = \frac{(4 - \beta^2)\, x^2}{16}, \tag{8.6}$$

and U earns $w_i q_i = \frac{\beta(2-\beta)x^2}{8}$ from D_i.

Two effects. Compare the gross profit of D_i under the two contracts, given the same residual intercept x. Under the two-part tariff, D_i gets the share $1 - \beta$ of the largest surplus, $\frac{x^2}{4}$. Under the wholesale price contract the surplus is smaller, which is the *inefficiency effect*, and D_i gets more than the share $1 - \beta$ of it, which is the *distribution effect*.

Result 1 (inefficiency and distribution effects). Given the residual intercept x:

- under a two-part tariff, $w_i = 0$, $F_i = \frac{\beta x^2}{4}$, and the gross profit of D_i is $\frac{(1-\beta)x^2}{4}$;
- under a wholesale price contract, $w_i = \frac{\beta x}{2}$, and the gross profit of D_i is the share $\frac{2-\beta}{2+\beta} = 1-\beta+\frac{\beta^2}{2+\beta}$ of the surplus $\frac{(4-\beta^2)x^2}{16}$;
- the gap between the two profits is

$$\frac{(2-\beta)^2x^2}{16} - \frac{(1-\beta)x^2}{4} = \underbrace{-\frac{(1-\beta)\beta^2x^2}{16}}_{\text{inefficiency effect}} + \underbrace{\frac{(2-\beta)\beta^2x^2}{16}}_{\text{distribution effect}} = \frac{\beta^2x^2}{16} \geq 0, \tag{8.7}$$

  where the inefficiency effect is $(1-\beta)$ times the change in the surplus, $\frac{(4-\beta^2)x^2}{16} - \frac{x^2}{4} = -\frac{\beta^2x^2}{16}$, and the distribution effect is the excess of the profit of D_i over the share $1-\beta$ of the surplus $\frac{(4-\beta^2)x^2}{16}$.

Figure 8.1 depicts the two effects and their sum, in units of $\frac{x^2}{16}$, with the bargaining power β on the horizontal axis. At $\beta = \frac{3}{4}$, the distribution effect is $\frac{45}{64}$, the inefficiency effect is $-\frac{9}{64}$, and the gap is $\frac{9}{16}$.

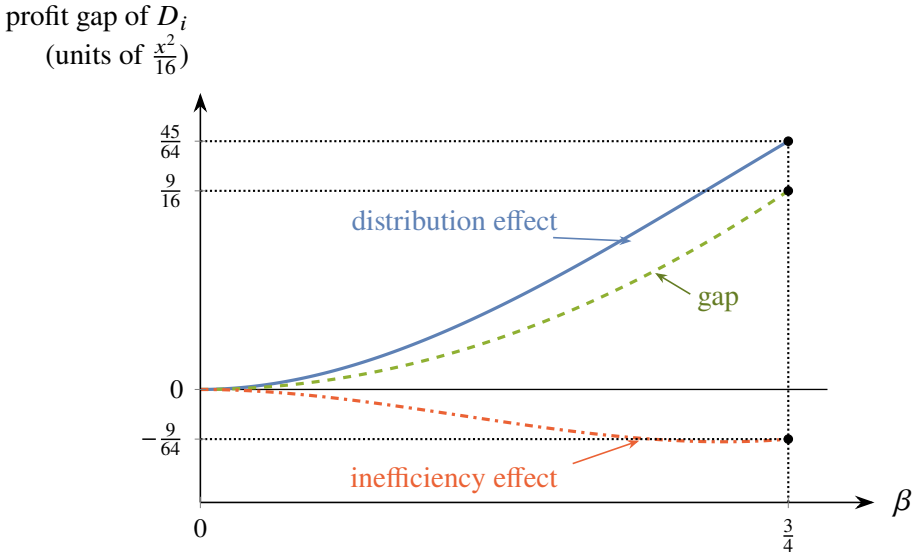


Figure 8.1: Decomposition of the gap between the gross profit of an entrant under a wholesale price contract and under a two-part tariff, for a given residual intercept x.

Intuitively, a wholesale price above zero cuts the output of D_i and with it the surplus of the pair, which lowers the profit of D_i. At the same time, U can raise its own payoff only by raising w_i, which shrinks the surplus it is trying to capture, so U stops short of taking the share β and leaves more to D_i. Both effects vanish at $\beta = 0$, where U sets $w_i = 0$ under either contract, and both grow with β. The distribution effect is $\frac{2-\beta}{1-\beta}$ times the size of the inefficiency effect, so it always dominates. The paper shows that both effects arise with a general demand system when contracts are secret.

8.4 Second Step: A Fixed Number of Entrants

Consider first that n entrants are in the market, as in most of the literature on vertical contracts. In a symmetric equilibrium each rival produces the same q, so $x = 1 - \frac{1}{2}(n-1)\,q$. Under the two-part tariff, $q = \frac{x}{2}$, and under the wholesale price contract, $q = \frac{(2-\beta)x}{4}$ from (8.6). Solving for x,

$$q^{\mathcal{T}n} = \frac{2}{n+3}, \qquad q^{\mathcal{W}n} = \frac{2(2-\beta)}{8+(2-\beta)(n-1)}, \qquad w^{\mathcal{W}n} = \frac{4\beta}{8+(2-\beta)(n-1)}, \tag{8.8}$$

which follow from equations (7) and (16) of the paper at $a = 1$, $c = 0$, and $\gamma = \frac{1}{2}$. The gross profit of an entrant is $(1-\beta)(q^{\mathcal{T}n})^2$ under the two-part tariff and $(q^{\mathcal{W}n})^2$ under the wholesale price contract, and U earns $\beta n(q^{\mathcal{T}n})^2$ and $n\,w^{\mathcal{W}n}q^{\mathcal{W}n}$, respectively. With n fixed, the entry costs are the same under both contracts, so welfare is measured as consumer surplus plus the gross profits of all firms.

Result 2 (fixed number of entrants). For $n \in [1,5]$ and $\beta \in (0, \frac{3}{4}]$:

- each entrant produces more under the two-part tariff, $q^{\mathcal{T}n} - q^{\mathcal{W}n} = \frac{8\beta}{(n+3)[8+(2-\beta)(n-1)]} > 0$, so aggregate output, consumer surplus, and welfare are higher under the two-part tariff;
- the profit of U is higher under the two-part tariff;
- the gross profit of each entrant is higher under the wholesale price contract.

Table 8.1 reports the case $n = 3$ and $\beta = \frac{3}{4}$, with every calculation.

Variable	Two-part tariff	Wholesale price contract
Residual intercept x	$1 - \frac{1}{2}\cdot 2\cdot\frac{1}{3} = \frac{2}{3}$	$\frac{8}{8+\frac{5}{4}\cdot 2} = \frac{16}{21}$
Wholesale price w	0	$\frac{3}{4}\cdot\frac{16}{21}\cdot\frac{1}{2} = \frac{2}{7}$
Output per entrant q	$\frac{1}{2}\cdot\frac{2}{3} = \frac{1}{3}$	$\frac{5}{4}\cdot\frac{16}{21}\cdot\frac{1}{4} = \frac{5}{21}$
Aggregate output	$3\cdot\frac{1}{3} = 1$	$3\cdot\frac{5}{21} = \frac{5}{7}$
Fixed fee F	$\frac{3}{4}\cdot\left(\frac{1}{3}\right)^2 = \frac{1}{12}$	none
Gross profit of an entrant	$\frac{1}{4}\cdot\left(\frac{1}{3}\right)^2 = \frac{1}{36} \approx 0.028$	$\left(\frac{5}{21}\right)^2 = \frac{25}{441} \approx 0.057$
Profit of U	$3\cdot\frac{1}{12} = \frac{1}{4}$	$3\cdot\frac{2}{7}\cdot\frac{5}{21} = \frac{10}{49} \approx 0.204$
Consumer surplus	$\frac{3\cdot 4}{4}\left(\frac{1}{3}\right)^2 = \frac{1}{3}$	$\frac{3\cdot 4}{4}\left(\frac{5}{21}\right)^2 = \frac{25}{147} \approx 0.170$
Welfare	$\frac{1}{3} + \frac{1}{4} + 3\cdot\frac{1}{36} = \frac{2}{3}$	$\frac{25}{147} + \frac{10}{49} + 3\cdot\frac{25}{441} = \frac{80}{147} \approx 0.544$

Table 8.1: Equilibrium with $n = 3$ entrants and $\beta = \frac{3}{4}$. Welfare is consumer surplus plus the gross profits of U and of the three entrants; the entry costs are the same under both contracts.

Intuitively, with n fixed the only difference between the contracts is the markup $w^{\mathcal{W}n} > 0$, which raises the marginal cost of every entrant and lowers output, so consumers and welfare lose. The

supplier loses too. The markup softens competition among entrants, so the surplus of each pair can be larger, but U keeps only the share $\frac{2\beta}{2+\beta} < \beta$ of it. In Table 8.1, the surplus of each pair is $\frac{55}{441}$ instead of $\frac{1}{9} = \frac{49}{441}$, while the share of U falls from $\frac{3}{4}$ to $\frac{6}{11}$. In the paper, the supplier prefers wholesale prices with a fixed n only when products are close substitutes and n is large; at $\gamma = \frac{1}{2}$ this requires $(n-1)\sqrt{2-\beta} > 4\sqrt{2}$, which fails for every $n \leq 5$. The entrants, in contrast, gain from wholesale prices, because the distribution effect of Result 1 dominates. This gain, $\frac{25}{441}$ instead of $\frac{1}{36}$ in Table 8.1, drives the reversal in the next steps.

8.5 Third Step: Free Entry

Entry continues until the gross profit of an entrant equals $K = \frac{1}{16}$. By Result 1, this condition determines the residual intercept under each contract:

$$\frac{(1-\beta)\,x^2}{4} = \frac{1}{16} \implies x^{\mathcal{T}} = \frac{1}{2\sqrt{1-\beta}}, \qquad \frac{(2-\beta)^2 x^2}{16} = \frac{1}{16} \implies x^{\mathcal{W}} = \frac{1}{2-\beta}. \tag{8.9}$$

Output per entrant is then $q^{\mathcal{T}} = \frac{x^{\mathcal{T}}}{2} = \frac{1}{4\sqrt{1-\beta}}$ and $q^{\mathcal{W}} = \frac{(2-\beta)x^{\mathcal{W}}}{4} = \frac{1}{4}$. The number of entrants follows from $x = 1 - \frac{1}{2}(n-1)\,q$.

Result 3 (more entry under wholesale prices). The free-entry numbers of entrants are

$$n^{\mathcal{T}} = 8\sqrt{1-\beta} - 3 \qquad \text{and} \qquad n^{\mathcal{W}} = \frac{10-9\beta}{2-\beta}, \tag{8.10}$$

and

- $n^{\mathcal{T}} = n^{\mathcal{W}} = 5$ at $\beta = 0$, and both decrease in β;
- $n^{\mathcal{W}} > n^{\mathcal{T}}$ for every $\beta \in (0, \frac{3}{4}]$, and the gap $n^{\mathcal{W}} - n^{\mathcal{T}}$ increases in β;
- at $\beta = \frac{3}{4}$, $n^{\mathcal{T}} = 1$ and $n^{\mathcal{W}} = \frac{13}{5}$.

Figure 8.2 depicts the two numbers of entrants against β. Expressions (8.10) are equations (9) and (18) of the paper at $\gamma = \frac{1}{2}$ and $k = \frac{1}{4}$, and the second bullet is its Proposition 1.

Intuitively, Result 1 says that, for any residual intercept, an entrant earns more under the wholesale price contract, with a gap of $\frac{\beta^2 x^2}{16}$. To bring its profit down to K, the residual intercept must therefore be lower under wholesale prices, $x^{\mathcal{W}} < x^{\mathcal{T}}$, since $(2-\beta)^2 - 4(1-\beta) = \beta^2 > 0$. A lower residual intercept means more output by rivals, which requires more entrants, all the more because each entrant produces less under wholesale prices. At $\beta = 0$ the two contracts coincide, and as β grows the profit gap of Result 1 grows, so the gap in entry widens. A higher β lowers entry under both contracts, but much more under the two-part tariff, where the fee takes the share β of the surplus of the pair; at $\beta = \frac{3}{4}$, a single firm enters.

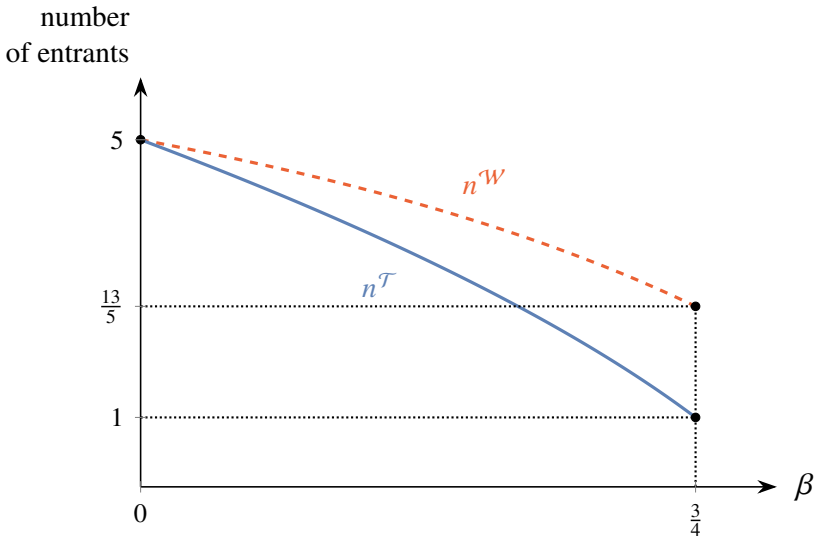


Figure 8.2: Free-entry number of entrants under a two-part tariff ($n^{\mathcal{T}}$, solid) and under a wholesale price contract ($n^{\mathcal{W}}$, dashed), against the bargaining power of U.

8.6 Fourth Step: Output per Entrant and Aggregate Output

Multiplying output per entrant by the number of entrants in (8.10), aggregate output is

$$Q^{\mathcal{T}} = 2 - \frac{3}{4\sqrt{1-\beta}} \qquad \text{and} \qquad Q^{\mathcal{W}} = \frac{10-9\beta}{4(2-\beta)}. \tag{8.11}$$

Writing $s = \sqrt{1-\beta}$, the condition $Q^{\mathcal{T}} = Q^{\mathcal{W}}$ becomes $s^3 + 3s^2 - 7s + 3 = 0$, that is, $(s-1)(s^2 + 4s - 3) = 0$. The root $s = 1$ is $\beta = 0$, and the root in $(0,1)$ is $s = \sqrt{7} - 2$, which gives the cutoff $\beta_Q = 1 - (\sqrt{7} - 2)^2$.

Result 4 (output). Under free entry:

- each entrant produces less under the wholesale price contract, $q^{\mathcal{W}} = \frac{1}{4} < q^{\mathcal{T}} = \frac{1}{4\sqrt{1-\beta}}$, for every $\beta \in (0, \frac{3}{4}]$;
- aggregate output is higher under the wholesale price contract, $Q^{\mathcal{W}} > Q^{\mathcal{T}}$, if and only if

$$\beta > \beta_Q \equiv 4\sqrt{7} - 10 \approx 0.583.$$

Figure 8.3 depicts aggregate output under both contracts against β. Both start at $\frac{5}{4}$ at $\beta = 0$ and cross at β_Q, where $Q \approx 0.839$, $n^{\mathcal{T}} \approx 2.17$, and $n^{\mathcal{W}} \approx 3.35$. At $\beta = \frac{3}{4}$, $Q^{\mathcal{T}} = \frac{1}{2}$ and $Q^{\mathcal{W}} = \frac{13}{20}$. The cutoff is the one in Proposition 2 of the paper, $\beta_Q(\gamma) = \frac{2(-1-\gamma+\gamma^2+\sqrt{1+(2-\gamma)\gamma})}{\gamma^2}$, evaluated at $\gamma = \frac{1}{2}$.

Intuitively, two forces work in opposite directions. Each entrant produces less under the wholesale price contract, because of the markup on its input, but more firms enter (Result 3). The second force dominates when β is high, because the distribution effect, and with it the gap in entry, grows

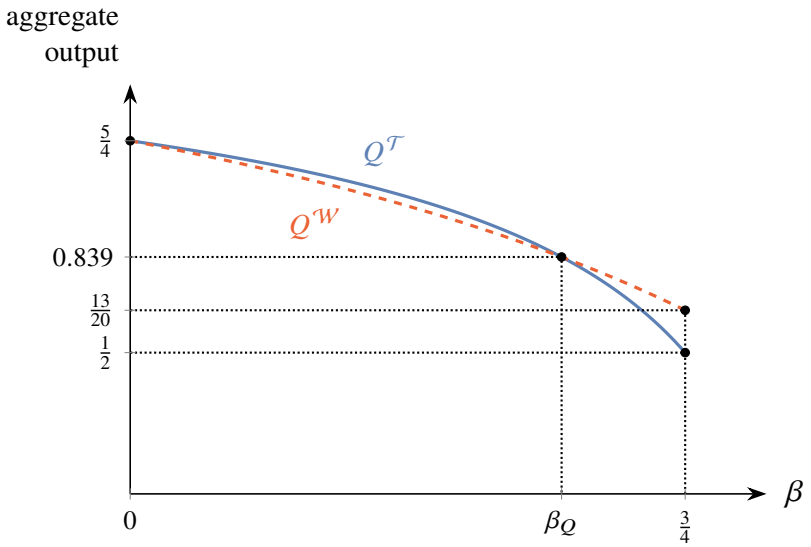


Figure 8.3: Free-entry aggregate output under a two-part tariff ($Q^{\mathcal{T}}$, solid) and under a wholesale price contract ($Q^{\mathcal{W}}$, dashed), against the bargaining power of U.

with β, while output per entrant under the wholesale price contract does not change with β. In the paper, $\beta_Q(\gamma)$ increases in γ: with closer substitutes, the entry gap is smaller and the supplier needs more bargaining power for the ranking to reverse.

8.7 Fifth Step: Consumer Surplus and Welfare

Consumer surplus (8.2) under free entry is

$$CS^{\mathcal{T}} = \frac{(8s-3)(4s-1)}{32\,s^2} \qquad \text{and} \qquad CS^{\mathcal{W}} = \frac{(10-9\beta)(6-5\beta)}{32\,(2-\beta)^2}, \tag{8.12}$$

with $s = \sqrt{1-\beta}$. Under the two-part tariff, U collects the fee $\beta(q^{\mathcal{T}})^2$ from each entrant. Under the wholesale price contract, it earns $w^{\mathcal{W}} q^{\mathcal{W}}$ from each entrant, with $w^{\mathcal{W}} = \frac{\beta x^{\mathcal{W}}}{2} = \frac{\beta}{2(2-\beta)}$. Therefore

$$\pi_U^{\mathcal{T}} = \frac{\beta(8s-3)}{16\,s^2} \qquad \text{and} \qquad \pi_U^{\mathcal{W}} = \frac{\beta(10-9\beta)}{8\,(2-\beta)^2}, \tag{8.13}$$

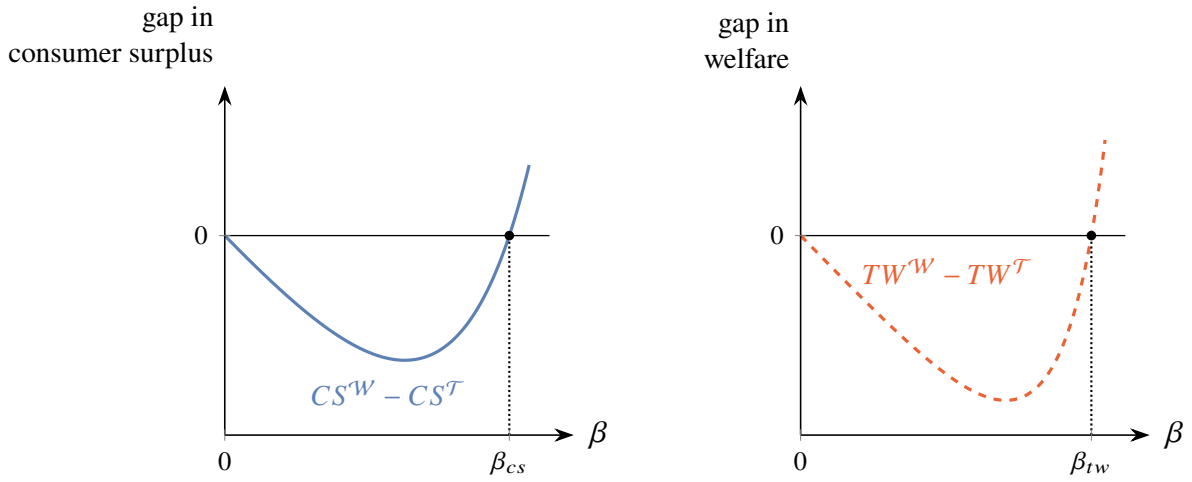
which are equations (11) and (20) of the paper at $\gamma = \frac{1}{2}$ and $k = \frac{1}{4}$. Welfare is $TW^{\rho} = CS^{\rho} + \pi_U^{\rho}$ for $\rho = \mathcal{T}, \mathcal{W}$, since entrants earn zero profit net of the entry cost.

Result 5 (consumer surplus and welfare). Under free entry:

- consumer surplus is higher under the wholesale price contract, $CS^{\mathcal{W}} > CS^{\mathcal{T}}$, if and only if $\beta > \beta_{cs} \approx 0.701$;
- welfare is higher under the wholesale price contract, $TW^{\mathcal{W}} > TW^{\mathcal{T}}$, if and only if $\beta > \beta_{tw} \approx 0.716$;

- the cutoffs are ranked $\beta_Q < \beta_{cs} < \beta_{tw} < \frac{3}{4}$.

Figure 8.4 depicts the gaps $CS^{\mathcal{W}} - CS^{\mathcal{T}}$ and $TW^{\mathcal{W}} - TW^{\mathcal{T}}$ against β. Both gaps are zero at $\beta = 0$ and negative at intermediate values, with lowest values of -0.037 at $\beta \approx 0.44$ and -0.050 at $\beta \approx 0.51$, respectively. Both turn positive near the upper end of (8.3), where they reach $\frac{17}{800} \approx 0.021$ and $\frac{23}{800} \approx 0.029$. The two results are Proposition 5 of the paper.



(a) Consumer surplus (b) Welfare

Figure 8.4: Consumer surplus and welfare under a wholesale price contract minus their values under a two-part tariff, for $\beta \in [0, \frac{3}{4}]$.

Intuitively, consumers gain from more output, so consumer surplus follows aggregate output. But by (8.2), with aggregate output $Q = nq$,

$$CS = \frac{Q\,(Q+q)}{4},$$

so, for a given Q, consumer surplus is higher when fewer firms each produce more. More varieties raise the utility from a given Q, but they raise prices by more, since $p = 1 - \frac{Q+q}{2}$ in a symmetric outcome. At β_Q the two contracts produce the same Q, but the wholesale price contract spreads it over more entrants, so $CS^{\mathcal{W}} < CS^{\mathcal{T}}$ there, and the supplier needs more bargaining power, $\beta > \beta_{cs}$, for consumers to gain. Welfare adds the profit of U, which is still higher under the two-part tariff at β_{cs} (see the sixth step), so the welfare cutoff is higher still. In the paper, both cutoffs decrease with the entry cost, so a higher entry cost works in favor of wholesale price contracts.

8.8 Sixth Step: The Supplier's Profit and Its Bargaining Power

The profit of U in (8.13) is its share of the surplus of each pair times the number of entrants. A higher β raises the share but, by Result 3, lowers entry. Under the two-part tariff, in terms of $s = \sqrt{1-\beta}$, the derivative of $\pi_U^{\mathcal{T}}$ is zero where

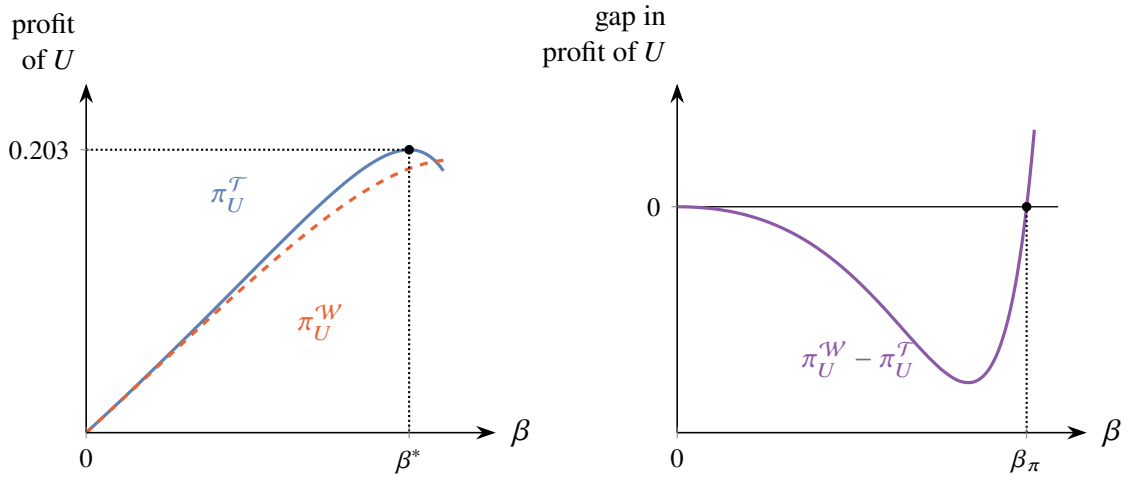
$$4s^3 + 4s = 3, \tag{8.14}$$

which has a single real root, $s^* \approx 0.567$. Under the wholesale price contract, the derivative of $\pi_U^{\mathcal{W}}$ has the sign of $10 - 13\beta$.

Result 6 (the supplier's profit). Under free entry:

- under the two-part tariff, $\pi_U^{\mathcal{T}}$ increases in β for $\beta < \beta^*$ and decreases for $\beta > \beta^*$, where $\beta^* = 1 - (s^*)^2 \approx 0.678$, and $\pi_U^{\mathcal{T}}(\beta^*) \approx 0.203 > \pi_U^{\mathcal{T}}(\frac{3}{4}) = \frac{3}{16}$;
- under the wholesale price contract, $\pi_U^{\mathcal{W}}$ increases in β on $[0, \frac{3}{4}]$;
- the supplier earns more under the wholesale price contract, $\pi_U^{\mathcal{W}} > \pi_U^{\mathcal{T}}$, if and only if $\beta > \beta_\pi \approx 0.734$.

Figure 8.5 depicts the two profits against β in panel (a) and their difference in panel (b). The first bullet is Proposition 3(i) of the paper: at $\gamma = \frac{1}{2}$, its cutoff $k_T(\gamma, \beta) = \frac{(2-\beta)\sqrt{1-\beta}}{3}$ equals $k = \frac{1}{4}$ at β^*, and $\beta^* \geq \gamma$. Under the wholesale price contract, the maximum is at $\beta = \frac{10}{13}$, where $\pi_U^{\mathcal{W}} = \frac{25}{128}$, which lies outside (8.3). With a higher entry cost, $K = \frac{9}{100}$, Assumption 1 requires $\beta \leq 0.64$ and the maximum is at $\beta = \frac{22}{35} \approx 0.629$, as Proposition 3(ii) allows. The third bullet is Proposition 4, which requires γ below about 0.601, as $\gamma = \frac{1}{2}$ is.



(a) Profit of U (b) Difference

Figure 8.5: Free-entry profit of U under a two-part tariff ($\pi_U^{\mathcal{T}}$, solid) and under a wholesale price contract ($\pi_U^{\mathcal{W}}$, dashed), and their difference, for $\beta \in [0, \frac{3}{4}]$.

Intuitively, when β is low the supplier gains from a larger share, but when β is high its share is already large and few firms enter, so the surplus to be shared is small. Under the two-part tariff, lowering β from $\frac{3}{4}$ to β^* raises entry from one firm to about 1.54, and the profit of U rises by about 0.015: the supplier gains by giving up bargaining power, for instance by delegating the negotiations to weaker agents, as the paper suggests. The same force explains the third bullet. At a high β, the wholesale price contract keeps entry, and therefore the demand for the input, much higher (Result 3), and this outweighs the lower output per entrant and the smaller share of the surplus that

U keeps. With a fixed number of entrants, the supplier preferred the two-part tariff for every $n \leq 5$ (Result 2); once entry responds, it prefers the wholesale price contract at high β.

Table 8.2 reports the free-entry equilibrium at $\beta = \frac{3}{4}$, with every calculation, for comparison with Table 8.1 at the same β. Every ranking of Table 8.1 is reversed: aggregate output, consumer surplus, welfare, and the profit of U are all higher under the wholesale price contract.

Variable	Two-part tariff	Wholesale price contract
Residual intercept x	$\frac{1}{2\sqrt{1-\frac{3}{4}}} = 1$	$\frac{1}{2-\frac{3}{4}} = \frac{4}{5}$
Number of entrants n	$8 \cdot \frac{1}{2} - 3 = 1$	$\frac{10-\frac{27}{4}}{\frac{5}{4}} = \frac{13}{5}$
Wholesale price w	0	$\frac{3}{4} \cdot \frac{4}{5} \cdot \frac{1}{2} = \frac{3}{10}$
Output per entrant q	$\frac{1}{2} \cdot 1 = \frac{1}{2}$	$\frac{5}{4} \cdot \frac{4}{5} \cdot \frac{1}{4} = \frac{1}{4}$
Aggregate output	$1 \cdot \frac{1}{2} = \frac{1}{2}$	$\frac{13}{5} \cdot \frac{1}{4} = \frac{13}{20}$
Fixed fee F	$\frac{3}{4} \cdot \left(\frac{1}{2}\right)^2 = \frac{3}{16}$	none
Gross profit of an entrant	$\frac{1}{4} \cdot \left(\frac{1}{2}\right)^2 = \frac{1}{16} = K$	$\left(\frac{1}{4}\right)^2 = \frac{1}{16} = K$
Profit of U	$1 \cdot \frac{3}{16} = \frac{3}{16} = 0.1875$	$\frac{13}{5} \cdot \frac{3}{10} \cdot \frac{1}{4} = \frac{39}{200} = 0.195$
Consumer surplus	$\frac{1 \cdot 2}{4}\left(\frac{1}{2}\right)^2 = \frac{1}{8} = 0.125$	$\frac{1}{4} \cdot \frac{13}{5} \cdot \frac{18}{5} \cdot \left(\frac{1}{4}\right)^2 = \frac{117}{800} \approx 0.146$
Welfare	$\frac{1}{8} + \frac{3}{16} = \frac{5}{16} = 0.3125$	$\frac{117}{800} + \frac{39}{200} = \frac{273}{800} \approx 0.341$

Table 8.2: Free-entry equilibrium with $\beta = \frac{3}{4}$ and $K = \frac{1}{16}$. Welfare is consumer surplus plus the profit of U, since entrants earn zero profit net of the entry cost.

8.9 Seventh Step: What the Toy Model Keeps and Drops

With $\gamma = \frac{1}{2}$, $a = 1$, $c = 0$, and $K = \frac{1}{16}$, the toy model keeps every main conclusion of the paper:

1. Wholesale price contracts lower the surplus of each pair (inefficiency effect) but leave each entrant more than its bargaining share (distribution effect), and the second effect dominates (the first step; Section IV(iii) of the paper).
2. With a fixed number of entrants, two-part tariffs give higher output, consumer surplus, welfare, and supplier profit (the second step).
3. With free entry, more firms enter under wholesale price contracts (the third step, Proposition 1).
4. Each entrant produces less under wholesale prices, but aggregate output is higher when the supplier's bargaining power exceeds β_Q (the fourth step, Proposition 2).
5. Consumer surplus and welfare are higher under wholesale prices at high bargaining power, with cutoffs above β_Q because consumers value variety (the fifth step, Proposition 5).

6. The supplier can gain by giving up bargaining power, and it can prefer wholesale price contracts (the sixth step, Propositions 3 and 4).

The toy model drops several features. It fixes γ, so it cannot show that the cutoffs β_Q, β_{cs}, and β_{tw} rise with product substitutability, or that the supplier prefers wholesale prices only when products are sufficiently differentiated. It fixes the entry cost, so the comparative statics in k of Propositions 3 to 5 are not shown. It treats the number of entrants as continuous, while Section VI of the paper shows similar rankings with an integer number of firms. Finally, it keeps the paper's main assumptions of secret contracts, quantity competition, and demand with love for variety. The paper reports that the main results also hold, under stricter conditions, when entrants observe all contracts before choosing quantities, when they compete in prices, and with a linear demand in which more varieties do not expand the market.

8.10 Connections to Other Chapters

- Chapter 3 also compares linear prices with two-part tariffs, there as prices that competing firms charge consumers rather than as contracts between a supplier and its buyers.
- In Chapter 6, the profit that an incumbent shares with retailers also decides which firms operate downstream, there to keep an entrant out.
- Chapter 7 studies another owner of an input that downstream firms need, whose terms of supply decide whether a downstream rival can cover its entry cost.

References

Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee (2019). "'Nash-in-Nash' Bargaining: A Microfoundation for Applied Work." *Journal of Political Economy*, 127(1), pp. 163–195.

Hart, O., and J. Tirole (1990). "Vertical Integration and Market Foreclosure." *Brookings Papers on Economic Activity: Microeconomics*, 1990, pp. 205–286.

Horn, H., and A. Wolinsky (1988). "Bilateral Monopolies and Incentives for Merger." *RAND Journal of Economics*, 19(3), pp. 408–419.

Milliou, C., and E. Petrakis (2024). "Vertical Contracts and Downstream Entry." *Journal of Industrial Economics*, 72(1), pp. 598–629. https://doi.org/10.1111/joie.12367

O'Brien, D. P., and G. Shaffer (1992). "Vertical Control with Bilateral Contracts." *RAND Journal of Economics*, 23(3), pp. 299–308.

Singh, N., and X. Vives (1984). "Price and Quantity Competition in a Differentiated Duopoly." *RAND Journal of Economics*, 15(4), pp. 546–554.

Chapter 9

Oligopolistic Entry Deterrence: When a High Price Keeps Entrants Out

Companion note to Harrington, J. E., Jr. (1987), "Oligopolistic Entry Deterrence under Incomplete Information," RAND Journal of Economics 18(2): 211–231.

9.1 Introduction

Firms already in a market may use their prices to keep rivals out. Consider two airlines that serve a route and know the cost of operating it, and a third airline that considers entering with the same aircraft. The entrant does not know the cost, but it observes the fare. Can the incumbents use the fare to keep the entrant out, even though they do not collude?

In the single-incumbent model of Milgrom and Roberts (1982), the incumbent charges a low price to signal that its cost is low, so that the entrant expects a tough rival. Here, instead, the incumbents' cost is also the cost that the entrant would face, since it uses the same technology. A high price then signals high costs for everyone, which is bad news for the entrant. Low-cost incumbents can therefore deter entry by producing less and charging the price of high-cost firms. Each incumbent restrains its output because, if it alone expanded, the industry price would fall and the entrant would come in. No incumbent can free ride on its rival's restraint.

Harrington (1987) shows that noncooperative incumbents can deter entry that would occur under complete information, and that they do so with a high price. More incumbents have an ambiguous effect on the probability of entry, while more potential entrants raise the price before entry and lower the probability of entry. Potential competition can thus raise prices before entry occurs.

The toy model has linear demand, two cost levels, and two incumbents. The first step solves the complete-information benchmark, the second step studies how the entrant reads the price, and the third step shows why noncooperative incumbents can deter entry. The fourth step characterizes the pooling equilibrium, in which both cost types charge the same price, and the separating equilibria,

in which the price reveals the cost. The fifth step changes the number of incumbents, the sixth step adds potential entrants, and the seventh step lists what the toy model keeps and drops.

9.2 The Toy Model

Firms and demand. Two incumbents (and, in the fifth step, n incumbents) produce a homogeneous good. In each of two periods, inverse demand is

$$p = 1 - Q, \tag{9.1}$$

where Q is total output. Every active firm has the same constant marginal cost c.

Private information. The common cost is either low, $c_L = 0$, or high, $c_H \in (0, 1)$. The incumbents observe c; the potential entrant does not, and assigns probability ρ to the low cost. An entrant uses the same technology, so its marginal cost is also c.

Timing. In period 1, the incumbents choose outputs simultaneously and noncooperatively. The entrant observes the period-1 price p_1, or equivalently total output $Q_1 = 1 - p_1$, but not individual outputs or costs, and decides whether to enter at a sunk cost $K > 0$. In period 2, the cost becomes common knowledge among the active firms, which play the Cournot equilibrium. There is no discounting.

Cournot benchmark. With k active firms and cost c, each firm produces and earns

$$q = \frac{1 - c}{k + 1}, \qquad p = \frac{1 + kc}{k + 1}, \qquad \pi_k(c) = \frac{(1 - c)^2}{(k + 1)^2}. \tag{9.2}$$

With two incumbents, an entrant who joins them alone earns

$$V(c) = \frac{(1 - c)^2}{16} - K \tag{9.3}$$

in period 2, net of the entry cost. Throughout, $(1 - c_H)^2 < 16K < 1$, so that entry is profitable if and only if the cost is low. The baseline values are $c_H = \frac{2}{5}$, $K = \frac{13}{400}$, and $\rho = \frac{1}{10}$, which give

$$V_L \equiv V(0) = \frac{3}{100}, \qquad V_H \equiv V(c_H) = -\frac{1}{100}.$$

9.3 First Step: The Complete-Information Benchmark

If the entrant observed c, period-1 output would not affect its decision, so each incumbent would produce its Cournot output in period 1: $q_L = \frac{1}{3}$ with low cost and $q_H = \frac{1 - c_H}{3}$ with high cost. Total outputs are $Q_L = \frac{2}{3}$ and $Q_H = \frac{2(1 - c_H)}{3}$, and the prices are

$$p_L = \frac{1}{3}, \qquad p_H = \frac{1 + 2c_H}{3}. \tag{9.4}$$

Entry occurs if and only if $V(c) > 0$, that is, if and only if $c < c^0 \equiv 1 - 4\sqrt{K}$; the assumption $(1 - c_H)^2 < 16K < 1$ places c^0 between 0 and c_H.

Result 1 (complete information, inequality (1) of the paper). Under complete information:

- the entrant enters if and only if the cost is low, so the probability of entry is ρ;
- the incumbents charge $p_L = \frac{1}{3}$ with low cost and $p_H = \frac{1+2c_H}{3}$ with high cost; at the baseline, $p_H = \frac{3}{5}$ and the probability of entry is $\frac{1}{10}$.

Intuitively, a high common cost hurts the entrant in two ways: it makes the incumbents weaker rivals, which helps the entrant, but it also raises the entrant's own cost. In the Cournot profit $\frac{(1-c)^2}{16}$, the second effect dominates, so entry is profitable only when costs are low.

Under incomplete information, if the incumbents kept producing their Cournot outputs, the price would reveal the cost and the outcome of Result 1 would follow. The question is whether low-cost incumbents gain by producing less and charging the high-cost price.

9.4 Second Step: How the Entrant Reads the Price

Beliefs. Let $\mu(Q_1)$ denote the entrant's posterior probability of the low cost after observing total output Q_1. Its expected profit from entering is

$$E(Q_1) = \mu(Q_1)\, V_L + \big(1 - \mu(Q_1)\big) V_H, \tag{9.5}$$

which increases in μ. If both cost types produce the same total output, the entrant learns nothing, $\mu = \rho$, and stays out if and only if $\rho V_L + (1 - \rho)V_H \le 0$, that is, if and only if

$$\rho \;\le\; \bar{\rho} \equiv \frac{-V_H}{V_L - V_H} = \frac{16K - (1 - c_H)^2}{1 - (1 - c_H)^2}. \tag{9.6}$$

At the baseline, $\bar{\rho} = \frac{1}{4}$ and the prior expected profit is $\frac{1}{10} \cdot \frac{3}{100} - \frac{9}{10} \cdot \frac{1}{100} = -\frac{3}{500} < 0$.

Monotone beliefs. Following the paper (Condition 1), beliefs must make the expected entry profit weakly increasing in Q_1 at every output, on and off the equilibrium path: a higher price never makes the entrant more optimistic about entry. In the toy model, this means that $\mu(Q_1)$ is weakly increasing.

Result 2 (the entrant's rule, Condition 1 of the paper). With monotone beliefs:

- the entrant enters if and only if $Q_1 > \hat{Q}$ for some limit output $\hat{Q}$, that is, if and only if the price falls below the limit price $\hat{p} = 1 - \hat{Q}$;
- an output produced by both cost types deters entry if and only if $\rho \le \bar{\rho}$, with $\bar{\rho} = \frac{1}{4}$ at the baseline.

Figure 9.1 depicts the expected entry profit (9.5) as a function of the price p_1 under the beliefs used in the fourth step: $\mu = \rho$ at $p_1 = \frac{3}{5}$, the high-cost price, $\mu = 1$ at any lower price, and $\mu = 0$ at any higher price.

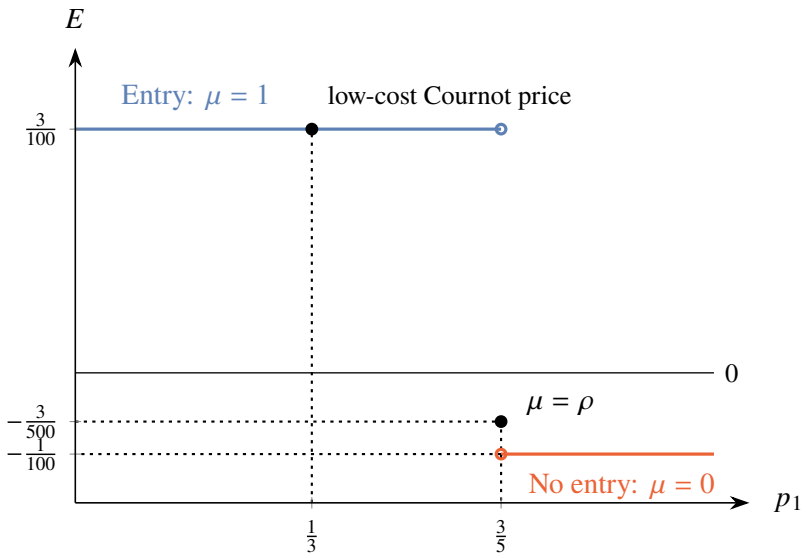


Figure 9.1: Entrant's expected profit from entering as a function of the period-1 price, with monotone beliefs and the high-cost price $\frac{3}{5}$ chosen by both cost types (baseline).

Intuitively, the entrant cannot see the incumbents' cost, but it can see their price, and a high price is what high-cost firms charge. In Milgrom and Roberts (1982), the entrant's cost does not depend on the incumbent's, so a high incumbent cost only means a weaker rival, and the entrant reads a low price as bad news. Here a high price is bad news for the entrant, because the cost it signals is also the entrant's cost. Harrington (1986) obtains the same signal with a single incumbent. The incumbents therefore deter entry by keeping the price high, not low.

9.5 Third Step: Deterring Entry without Collusion

Consider that the entrant enters if and only if $Q_1 > Q_H$, the high-cost Cournot output. With high cost, each incumbent produces q_H and entry is deterred at no cost. With low cost, deterring entry requires total output Q_H instead of Q_L, and each incumbent's choice depends on what its rival does.

If the rival restricts its output. Let the rival produce q_H. If firm i also produces q_H, the price is p_H, it earns $(1-2q_H)q_H$, and it shares the period-2 market with one rival, earning $\frac{1}{9}$. If instead it plays its best response, $\frac{1-q_H}{2}$, it earns $\left(\frac{1-q_H}{2}\right)^2$, the entrant comes in, and it earns $\frac{1}{16}$ in period 2. Since $1-3q_H = c_H$, the period-1 profit given up by restricting is

$$\left(\frac{1-q_H}{2}\right)^2 - (1-2q_H)q_H = \frac{(1-3q_H)^2}{4} = \frac{c_H^2}{4},$$

so the gain from restricting is

$$\phi_1(c_H) = \left(\frac{1}{9} - \frac{1}{16}\right) - \frac{c_H^2}{4} = \frac{7}{144} - \frac{c_H^2}{4}. \tag{9.7}$$

If the rival produces its Cournot output. Let the rival produce $q_L = \frac{1}{3}$. To keep total output at Q_H, firm i must produce $y = Q_H - \frac{1}{3} = \frac{1-2c_H}{3}$ alone. Its period-1 profit falls from $\frac{1}{9}$ to $(1 - Q_H)y$, a loss of $\left(\frac{1}{3} - y\right)^2 = \frac{4c_H^2}{9}$, so the gain from restricting is

$$\phi_2(c_H) = \frac{7}{144} - \frac{4c_H^2}{9}. \tag{9.8}$$

(For $c_H > \frac{1}{2}$, the required output y would be negative, and firm i cannot deter entry on its own.) The gains (9.7) and (9.8) are zero at

$$\bar{c}_1 = \frac{\sqrt{7}}{6} \approx 0.441 \qquad \text{and} \qquad \bar{c}_2 = \frac{\sqrt{7}}{8} \approx 0.331. \tag{9.9}$$

Result 3 (deterrence without collusion, Section 3 of the paper). When the entrant enters if and only if $Q_1 > Q_H$, the low-cost incumbents play a game in which:

- both producing q_H, which deters entry, is a Nash equilibrium if and only if $c_H \le \bar{c}_1$;
- both producing q_L, which invites entry, is a Nash equilibrium if and only if $c_H \ge \bar{c}_2$;
- for $c_H \in [\bar{c}_2, \bar{c}_1]$ both are Nash equilibria; at the baseline $c_H = \frac{2}{5}$, each firm earns $\frac{3}{25} + \frac{1}{9} = \frac{52}{225} \approx 0.231$ over the two periods when both restrict, and $\frac{1}{9} + \frac{1}{16} = \frac{25}{144} \approx 0.174$ when both produce q_L.

Figure 9.2 depicts the two gains as functions of the high cost c_H. At the baseline, $\phi_1 = \frac{31}{3600} > 0$ and $\phi_2 = -\frac{81}{3600} < 0$.

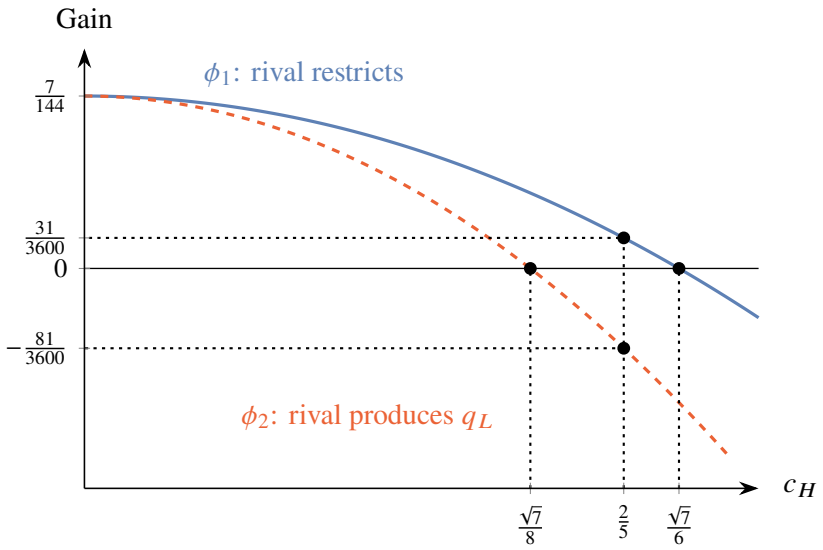


Figure 9.2: Gain to a low-cost incumbent from restricting its output to deter entry, when its rival restricts (ϕ_1, solid) and when its rival produces its Cournot output (ϕ_2, dashed).

Intuitively, the entrant observes only the industry price, so deterrence is all or nothing: if a single incumbent expands, the price falls below the limit price and the entrant comes in. Each incumbent

can therefore trigger entry on its own, and it bears the whole loss from entry, $\frac{1}{9} - \frac{1}{16}$, when it expands; there is no room to free ride on the rival's restraint. This is why noncooperative firms can deter entry. Restraint is easier when the rival also restrains itself, because a firm then gives up only the gap between its best response and q_H; when the rival produces its Cournot output, the firm must absorb the whole output cut alone. For intermediate costs, deterring entry and allowing it are both self-enforcing, and the incumbents face a coordination problem: both earn more when they deter entry, but each restricts only if it expects the other to restrict too.

9.6 Fourth Step: Pooling and Separating Equilibria

Equilibrium selection. The paper imposes three requirements. First, beliefs are monotone, as in the second step. Second, given the entrant's rule, the incumbents' strategies are Pareto optimal among the profiles in which every incumbent best responds for every cost (Definition 2 of the paper); in the third step, this selects deterrence whenever both outcomes are Nash equilibria. Third, the equilibrium must be robust: the entrant's rule must be optimal under every monotone belief, which in the paper places the limit output where the entrant is indifferent about entering (Definition 1 and the discussion that follows it). With two costs, unless the expected entry profit at the limit output is zero, some monotone beliefs make entry unprofitable at outputs slightly above the limit output, so the entrant's rule is not optimal under every monotone belief. At an output chosen by both types this profit equals its prior value, so the requirement rules out every equilibrium of the toy model except in the knife-edge case $\rho = \bar{\rho}$. The toy model keeps the first two requirements and, among the equilibria that satisfy them, selects the one with the largest limit output $\hat{Q}$, that is, the smallest output restriction.

A separating restriction. When pooling cannot deter entry, high-cost firms must choose an output that low-cost firms would not copy even if both of them did. By (9.7), a low-cost firm whose rival produces x gains $\frac{7}{144} - \frac{(1-3x)^2}{4}$ from matching it, which is zero at

$$x_s = \frac{1-\bar{c}_1}{3} = \frac{6-\sqrt{7}}{18} \approx 0.186, \qquad p_s = 1 - 2x_s = \frac{1+2\bar{c}_1}{3} = \frac{3+\sqrt{7}}{9} \approx 0.627. \tag{9.10}$$

That is, high-cost firms charge the Cournot price of two firms whose cost is $\bar{c}_1$.

Result 4 (equilibrium, Propositions 1 and 2 of the paper). In the selected equilibrium:

- if $c_H \leq \bar{c}_1$ and $\rho \leq \bar{\rho}$ (pooling), each incumbent produces q_H for both costs, the price is $p_H = \frac{1+2c_H}{3}$, the limit output is $\hat{Q} = Q_H$, and entry never occurs;
- if $c_H > \bar{c}_1$ (separation at Cournot outputs), each incumbent produces its Cournot output, and entry occurs if and only if the cost is low, with probability ρ;
- if $c_H \leq \bar{c}_1$ and $\rho > \bar{\rho}$ (separation with a high price), low-cost firms produce q_L and entry follows; high-cost firms produce $x_s < q_H$, charge $p_s = \frac{3+\sqrt{7}}{9} > p_H$, and entry does not occur; the limit output is $\hat{Q} = 2x_s$.

At the baseline ($c_H = \frac{2}{5}$, $\rho = \frac{1}{10}$), the equilibrium is pooling: the low-cost price rises from $\frac{1}{3}$ to $\frac{3}{5}$, and the probability of entry falls from $\frac{1}{10}$ to 0.

Figure 9.3 depicts the three regions for $K = \frac{13}{400}$. The curve is $\bar{\rho}$ in (9.6), which rises with c_H because a higher c_H makes entry less attractive in the high-cost state; to the left of $c^0 = 1 - 4\sqrt{K} \approx 0.279$, entry is profitable for both costs and no price can deter it. The dot marks the baseline.

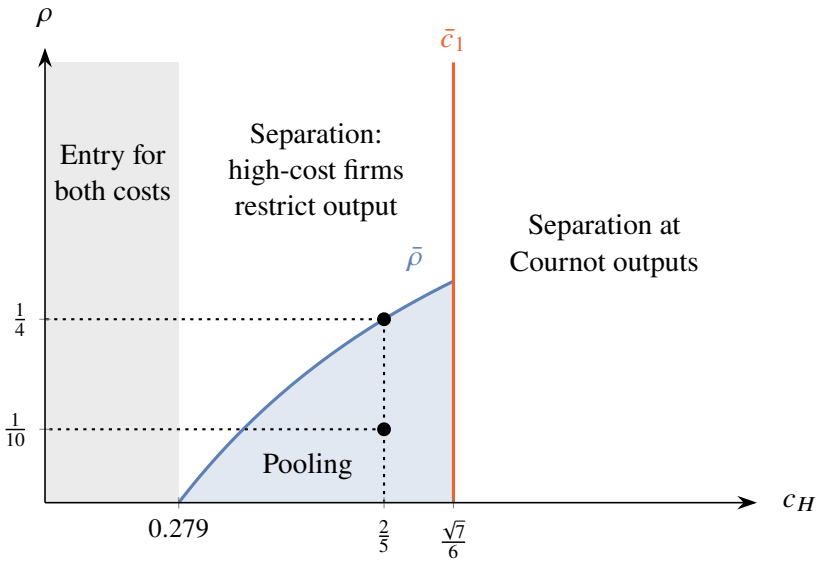


Figure 9.3: Equilibrium regions for each high cost c_H and prior probability ρ of the low cost ($K = \frac{13}{400}$, two incumbents, one potential entrant).

Intuitively, when the gap between the two costs is small, mimicking the high-cost price costs low-cost firms little, and they pool. If the entrant is pessimistic enough, pooling keeps it out, and entry that would occur under complete information is deterred: the low-cost incumbents trade lower profit today for no entry tomorrow. If the entrant is optimistic, pooling cannot keep it out, and the low-cost firms would mimic any price they found worthwhile; to be recognized, high-cost firms raise their price above their Cournot level until mimicking is no longer worthwhile for low-cost firms. In both cases, the signal that keeps the entrant out is a high price. When the cost gap is large, mimicking costs low-cost firms too much, and the outcome is the one under complete information.

9.7 Fifth Step: More Incumbents

With n incumbents, each produces $q_H = \frac{1-c_H}{n+1}$ in the high-cost Cournot equilibrium, and a low-cost firm that matches its rivals' restriction still gives up $\frac{(1-(n+1)q_H)^2}{4} = \frac{c_H^2}{4}$ in period 1. What changes is the loss from entry, $\frac{1}{(n+1)^2} - \frac{1}{(n+2)^2}$, and the entrant's profit, $\frac{(1-c)^2}{(n+2)^2} - K$. Hence

$$\bar{c}_1(n) = 2\sqrt{\frac{1}{(n+1)^2} - \frac{1}{(n+2)^2}} = \frac{2\sqrt{2n+3}}{(n+1)(n+2)}, \qquad \bar{\rho}(n) = \frac{(n+2)^2K - (1-c_H)^2}{1-(1-c_H)^2}, \quad (9.11)$$

where $\bar{\rho}(n)$ is defined when $(1-c_H)^2 < (n+2)^2K < 1$.

Result 5 (number of incumbents, Proposition 3 of the paper). As n rises:

- $\bar{c}_1(n)$ decreases, so pooling requires a smaller cost gap;
- $\bar{\rho}(n)$ increases, and the entrant's profit falls in both states;
- hence the probability of entry can rise or fall: in the example of Table 9.1, it falls from 1 to 0 from $n = 1$ to $n = 2$, while the low-cost price rises from $\frac{1}{2}$ to $\frac{3}{5}$; it rises from 0 to $\frac{1}{10}$ from $n = 2$ to $n = 3$, while both prices fall.

Table 9.1 reports the equilibrium for $n = 1, \ldots, 4$ at $c_H = \frac{2}{5}$, $K = \frac{13}{400}$, and $\rho = \frac{1}{10}$.

Table 9.1: Equilibrium with n incumbents ($c_H = \frac{2}{5}$, $K = \frac{13}{400}$, $\rho = \frac{1}{10}$, one potential entrant). n.a.: both costs ($n = 1$) or neither cost ($n = 4$) attract entry under complete information.

Incumbents	$\bar{c}_1(n)$	$\bar{\rho}(n)$	Entry probability, complete information	Entry probability, equilibrium	Low-cost price	High-cost price
$n = 1$	0.745	n.a.	1	1	$\frac{1}{2}$	$\frac{7}{10}$
$n = 2$	0.441	$\frac{1}{4}$	$\frac{1}{10}$	0	$\frac{3}{5}$	$\frac{3}{5}$
$n = 3$	$\frac{3}{10}$	0.707	$\frac{1}{10}$	$\frac{1}{10}$	$\frac{1}{4}$	$\frac{11}{20}$
$n = 4$	0.221	n.a.	0	0	$\frac{1}{5}$	$\frac{13}{25}$

Intuitively, more incumbents have two opposing effects. Entry becomes less attractive, because the entrant would join a more competitive market; this lowers the probability of entry. But each incumbent also loses less when entry occurs, so it is less willing to restrict output to prevent it; this raises the probability of entry. With a monopoly ($n = 1$), the entrant comes in regardless of its beliefs. With two incumbents, the low-cost state is the only one that attracts entry, and the incumbents hide it by pooling, at a higher price than a monopolist with low cost would charge. With three incumbents, the gain from deterrence no longer covers the cost of mimicking ($\bar{c}_1(3) = \frac{3}{10} < \frac{2}{5}$), so low-cost firms reveal themselves and entry returns.

9.8 Sixth Step: More Potential Entrants

Return to two incumbents and let m potential entrants observe p_1 and decide simultaneously, each with the same rule: enter if and only if $Q_1 > \hat{Q}$. When the low cost is revealed, all m enter, which requires $\frac{1}{(m+3)^2} \geq K$; at $K = \frac{13}{400}$, this holds for $m \leq 2$. At an output that deters entry, each entrant takes the others' staying out as given, so its entry profit is still $V(c)$ in (9.3), and $\bar{\rho}$ does not depend on m. What changes is the incumbents' loss from entry, $\frac{1}{9} - \frac{1}{(m+3)^2}$, which raises the cutoff and lowers the separating output in (9.10):

$$\bar{c}_1(m) = 2\sqrt{\frac{1}{9} - \frac{1}{(m+3)^2}}, \qquad x_s(m) = \frac{1 - \bar{c}_1(m)}{3}, \qquad p_s(m) = \frac{1 + 2\bar{c}_1(m)}{3}. \tag{9.12}$$

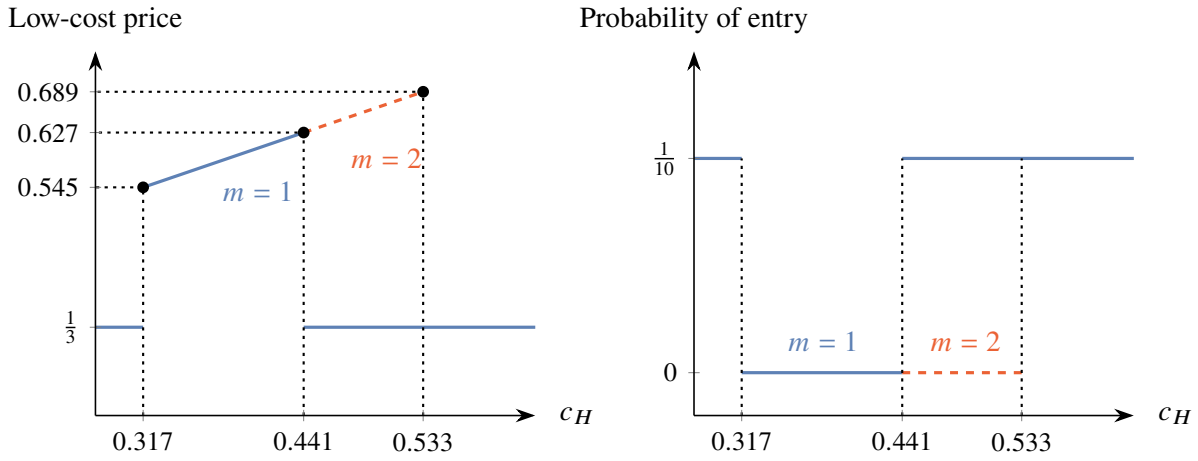
With $m = 1$, $\bar{c}_1 = \frac{\sqrt{7}}{6} \approx 0.441$ and $p_s \approx 0.627$; with $m = 2$, $\bar{c}_1 = \frac{8}{15} \approx 0.533$ and $p_s = \frac{31}{45} \approx 0.689$.

Result 6 (potential competition, Proposition 6 and Corollaries 1 and 2 of the paper). When the number of potential entrants rises from m to $m+1$ (with all of them entering after a low-cost output):

- $\bar{c}_1(m)$ increases and $\bar{\rho}$ is unchanged, so the pooling region expands, and $p_s(m)$ increases;
- for every cost, the period-1 price weakly rises and the probability of entry weakly falls;
- the price rises strictly for low cost if $\bar{c}_1(m) < c_H \le \bar{c}_1(m+1)$ and $\rho \le \bar{\rho}$, where the probability of entry falls from ρ to 0; it rises strictly for high cost if $c_H < \bar{c}_1(m+1)$ and $\rho > \bar{\rho}$.

For example, with $c_H = \frac{1}{2}$ and $\rho = \frac{1}{10}$, the gain (9.7) rises from $\frac{7}{144} - \frac{1}{16} = -\frac{1}{72}$ to $\frac{16}{225} - \frac{1}{16} = \frac{31}{3600}$; the low-cost price rises from $\frac{1}{3}$ to $\frac{2}{3}$, and the probability of entry falls from $\frac{1}{10}$ to 0.

Figure 9.4 depicts the low-cost price and the probability of entry as functions of c_H, for $\rho = \frac{1}{10}$, with one and two potential entrants. For c_H between 0.279 and 0.317, the prior exceeds $\bar{\rho}$, so pooling cannot deter entry; between 0.317 and $\bar{c}_1(m)$, the incumbents pool at the high-cost price.



(a) Period-1 price, low cost

(b) Probability of entry

Figure 9.4: Period-1 price with low cost and probability of entry, with one (solid) and two (dashed) potential entrants ($\rho = \frac{1}{10}$, $K = \frac{13}{400}$, two incumbents). Where the curves coincide, only the solid one is drawn.

Intuitively, a second potential entrant does not change any entrant's calculation at a deterring price, since each one takes the others' staying out as given. It changes the incumbents' calculation: allowing entry now means sharing the market with two newcomers, so deterrence is worth more, and low-cost incumbents mimic the high-cost price for a wider range of costs. When pooling cannot deter entry, high-cost firms must raise their price further to stay recognizable. Potential competition therefore raises prices before entry. The paper finds the same pattern in its example with a continuum of costs (uniform on $[0,1]$, two incumbents, $K = 0.015$): a second potential entrant raises the limit price from 0.769 to 0.796 and lowers the probability of entry from 0.379 to 0.344 (Table 3 of the paper).

9.9 Seventh Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- A high price deters entry because the incumbents' cost is also the entrant's cost, in contrast with Milgrom and Roberts (1982).
- Noncooperative incumbents deter entry, because each one can trigger entry on its own by expanding; deterring and allowing entry can both be self-enforcing, and the Pareto criterion selects deterrence.
- The deterrence of entry that would occur under complete information, the ambiguous effect of the number of incumbents, and the effect of potential competition on prices and on the probability of entry.

What it drops.

- A continuum of costs. In the paper, the incumbents pool over an interval of costs at the Cournot output of the highest cost in the interval, and the limit price makes the entrant indifferent about entering; firms whose cost lies outside the interval produce their Cournot outputs. With two costs, the expected entry profit at a pooled output cannot be adjusted to zero, so the paper's robustness requirement has no counterpart, and the high-cost restriction x_s of the fourth step has no counterpart in the paper.
- Asymmetric equilibria, in which incumbents share the burden of deterrence unequally or potential entrants follow different rules (Proposition 7 of the paper).
- The effect of the entry cost (Proposition 4 of the paper), discounting, and welfare. Deterrence through commitment to output, as in Gilbert and Vives (1986), works through a large output and a low price instead.

9.10 Connections to Other Chapters

- Chapter 10 studies two incumbents that observe demand but disagree about entry, because one sells a substitute for the entrant's good and the other a complement. In this chapter all incumbents lose from entry, so their interests are aligned.
- Chapter 5 studies another way to deter entry, a rebate conditional on exclusivity, and shows that it cannot keep a more efficient entrant out.
- Chapter 24 studies a firm that signals its quality to consumers with its price. In this chapter, prices signal the industry's cost to a potential entrant.

References

Gilbert, R., and Vives, X. (1986). Entry deterrence and the free rider problem. *Review of Economic Studies*, 53(1), 71–83.

Harrington, J. E., Jr. (1986). Limit pricing when the potential entrant is uncertain of its cost function. *Econometrica*, 54(2), 429–437.

Harrington, J. E., Jr. (1987). Oligopolistic entry deterrence under incomplete information. *RAND Journal of Economics*, 18(2), 211–231.

Milgrom, P., and Roberts, J. (1982). Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica*, 50(2), 443–459.

Chapter 10

Limit Pricing when Incumbents Disagree about Entry

Companion note to Schultz, C. (1999), "Limit Pricing when Incumbents Have Conflicting Interests," International Journal of Industrial Organization 17(6): 801–825.

10.1 Introduction

Incumbents do not always agree about entry. Consider a hardware maker, a software firm, and a new software firm that considers entering. The hardware maker gains from entry, since more software makes its hardware more valuable, while the existing software firm loses. Both incumbents know whether demand is high or low; the entrant does not, but it observes their prices before deciding. What can the entrant learn from these prices?

Since Milgrom and Roberts (1982), limit pricing models have usually featured incumbents who all lose from entry, as in Harrington (1987), so that every incumbent wants the entrant to believe that entry is unattractive. Here, the firm selling a substitute for the entrant's good wants to signal low demand, while the firm selling a complement wants to signal high demand. When the two prices tell different stories, the entrant must decide whom to believe, and any belief rewards one of the two firms.

Schultz (1999) shows that separating equilibria, in which prices reveal demand, involve no price distortion: each incumbent charges its full-information price. With the unprejudiced beliefs of Bagwell and Ramey (1991), under which the entrant ignores a deviation by one firm when the other firm's price reveals demand, these equilibria exist only when entry matters little for at least one incumbent. In growing markets, where the stakes from entry are large, only pooling equilibria, in which prices do not depend on demand, may exist. In some of them, the firm that welcomes entry raises its price while the firm that fears entry lowers it.

The toy model has linear demands, two demand states, and a uniformly distributed entry cost. The first step solves the full-information benchmark and the entry stage, the second step characterizes

separating equilibria and whom the entrant believes, and the third step shows that separation requires entry to matter little for one incumbent. The fourth step finds when pooling equilibria exist, the fifth step applies the Intuitive Criterion to pooling prices, and the sixth step lists what the toy model keeps and drops.

10.2 The Toy Model

Incumbents and demand states. Two incumbents, firms 1 and 2, sell goods 1 and 2 at zero marginal cost over two periods. Demand is high ($S = H$) or low ($S = L$), and the state is the same in both periods. Both incumbents observe the state; the entrant does not, and assigns prior probability $\rho^0 \in (0, 1)$ to the low state. In period 1, the demand for good i in state S is

$$q_i = \alpha_S - p_i, \qquad i = 1, 2, \tag{10.1}$$

with $\alpha_L = 1$ and $\alpha_H = \eta$, where $1 < \eta < 2$. Period-1 prices lie in $[0, 1]$, so demand is positive in both states (the paper also bounds period-1 prices). Goods 1 and 2 are independent, so each incumbent's period-1 profit, $\pi_i(p_i, S) = p_i(\alpha_S - p_i)$, depends only on its own price.

Entry and period 2. After observing (p_1, p_2), a potential entrant, firm 3, decides whether to enter at a sunk cost k that only it knows; the incumbents know that k is uniformly distributed on $[0, \bar{k}]$. In period 2, demand grows by the factor $\gamma > 0$: the market grows if $\gamma > 1$ and shrinks if $\gamma < 1$ (γ also absorbs discounting). Without entry, the demand for good i is $\gamma\alpha_S - p_i$. With entry,

$$q_1 = (1 + \theta_1)\gamma\alpha_S - p_1, \qquad q_2 = (1 - \theta_2)\gamma\alpha_S - p_2, \qquad q_3 = \gamma\alpha_S - p_3, \tag{10.2}$$

where $\theta_1, \theta_2 \in (0, 1)$. Entry raises the demand for good 1 and lowers the demand for good 2, as when good 3 is a complement of good 1 and a substitute for good 2; θ_1 and θ_2 measure how much entry matters for each incumbent. In the baseline, $\theta_1 = \theta_2 = \frac{1}{4}$. In period 2, each firm prices as a monopolist on its own good. The upper bound $\bar{k}$ equals the entrant's period-2 profit in the high state, so the entrant enters for sure when it knows that demand is high.

Timing and beliefs. (1) The incumbents observe the state and simultaneously set p_1 and p_2. (2) The entrant observes both prices, forms the belief $\rho(p_1, p_2)$ that the state is low, and enters if its expected profit exceeds k. (3) Period-2 profits are realized. An equilibrium is *separating* if the price pair differs across states and *pooling* if it does not. Beliefs are *unprejudiced* (Definition 2 of the paper): if one incumbent charges different prices in the two states and the other incumbent deviates to a price it never charges in equilibrium, the entrant ignores the deviation and infers the state from the non-deviating firm's price. If a firm deviates to the price it charges in the other state, the entrant cannot tell which firm deviated, and beliefs are unrestricted.

10.3 First Step: The Full-Information Benchmark and the Entry Stage

Period 1. Maximizing $p_i(\alpha_S - p_i)$ gives the full-information price $p_i^*(S) = \frac{\alpha_S}{2}$ and profit $\frac{\alpha_S^2}{4}$. Since demand is linear with unit slope, a firm that charges p_i instead of its best reply loses

$$\pi_i\left(\tfrac{\alpha_S}{2}, S\right) - \pi_i(p_i, S) = \left(p_i - \frac{\alpha_S}{2}\right)^2. \tag{10.3}$$

Period 2. A monopolist facing $a - p$ earns $\frac{a^2}{4}$. Without entry, each incumbent earns $\frac{\gamma^2\alpha_S^2}{4}$; with entry, firm 1 earns $\frac{(1+\theta_1)^2\gamma^2\alpha_S^2}{4}$, firm 2 earns $\frac{(1-\theta_2)^2\gamma^2\alpha_S^2}{4}$, and the entrant earns $\pi_e^E(S) = \frac{\gamma^2\alpha_S^2}{4}$. Firm 1's gain from entry and firm 2's loss from entry are

$$G_1(S) = \frac{\theta_1(2+\theta_1)}{4}\,\gamma^2\alpha_S^2, \qquad G_2(S) = \frac{\theta_2(2-\theta_2)}{4}\,\gamma^2\alpha_S^2, \tag{10.4}$$

which equal $\frac{9}{64}\gamma^2\alpha_S^2$ and $\frac{7}{64}\gamma^2\alpha_S^2$ in the baseline.

Entry. With belief ρ, the entrant's expected profit is $\rho\frac{\gamma^2}{4} + (1-\rho)\frac{\gamma^2\eta^2}{4}$. Since k is uniform on $[0, \frac{\gamma^2\eta^2}{4}]$, the probability of entry is

$$\xi(\rho) = \frac{\rho\frac{\gamma^2}{4} + (1-\rho)\frac{\gamma^2\eta^2}{4}}{\frac{\gamma^2\eta^2}{4}} = 1 - \lambda\rho, \qquad \text{where} \quad \lambda \equiv 1 - \frac{1}{\eta^2}. \tag{10.5}$$

Result 1 (full-information benchmark). In the toy model:

- each incumbent's full-information price is $p_i^*(S) = \frac{\alpha_S}{2}$, so the price pair is $(\frac{1}{2}, \frac{1}{2})$ in state L and $(\frac{\eta}{2}, \frac{\eta}{2})$ in state H;
- $G_1(S) > 0$ and $G_2(S) > 0$ in (10.4), and both are larger in state H than in state L;
- the probability of entry $\xi(\rho) = 1 - \lambda\rho$ decreases in the belief ρ, from $\xi(0) = 1$ to $\xi(1) = \frac{1}{\eta^2}$.

Intuitively, firm 1 wants the entrant to believe that demand is high, since a lower ρ raises the probability of entry, while firm 2 wants it to believe that demand is low. Both firms hold the same information, but each would like to send the opposite message. Result 1 reproduces assumptions (1), (2), and (16) of the paper: firm 1 gains and firm 2 loses from entry in every state, entry is more attractive in the high state, and the stakes from entry are larger in the high state.

10.4 Second Step: Separating Equilibria and Whom the Entrant Believes

Prices. In a separating equilibrium the entrant learns the state, so entry occurs with probability $\xi(1)$ in state L and $\xi(0)$ in state H. If firm 1 charged the same price in both states, firm 2 would

separate, the entrant would read the state from firm 2's price, and firm 1 could move to its best reply in each state without changing beliefs. If firm 1 separated but did not charge $\frac{1}{2}$ in state L, moving to $\frac{1}{2}$ would raise its period-1 profit and could only improve the entrant's belief, which is already the worst one for firm 1. The same arguments apply to firm 2.

Result 2 (Lemma 1 of the paper). In every separating equilibrium with unprejudiced beliefs:

- both incumbents separate and charge their full-information prices, $p_i(L) = \frac{1}{2}$ and $p_i(H) = \frac{\eta}{2}$;
- the only deviations that can change the entrant's belief are firm 1 charging $\frac{\eta}{2}$ in state L and firm 2 charging $\frac{1}{2}$ in state H; both lead to the price pair $(\frac{\eta}{2}, \frac{1}{2})$.

Figure 10.1 depicts the equilibrium price pairs and the two deviations. Firm 1's deviation moves the pair horizontally from the low-state point, and firm 2's deviation moves it vertically from the high-state point; both end at $(\frac{\eta}{2}, \frac{1}{2})$, where firm 1's price says "high demand" and firm 2's price says "low demand."

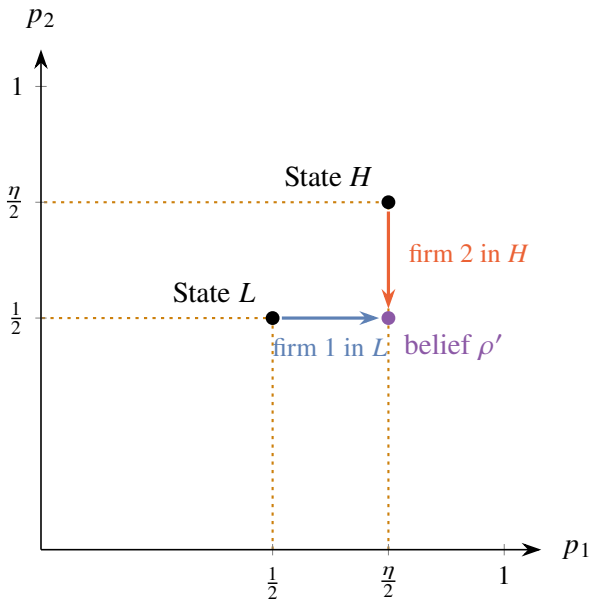


Figure 10.1: Separating price pairs and the two deviations that reach the same off-path pair $(\frac{\eta}{2}, \frac{1}{2})$.

Intuitively, a separating equilibrium involves no price distortion: every price is a best reply, so a deviation can be profitable only if it changes the entrant's belief. Under unprejudiced beliefs, a deviation to any other price is ignored, because the rival's price reveals the state. Only the off-path pair $(\frac{\eta}{2}, \frac{1}{2})$ creates a conflict, since each firm's price contradicts the other's.

Who should the entrant believe? Let ρ' denote the entrant's belief after observing $(\frac{\eta}{2}, \frac{1}{2})$. Mimicking the other state costs each firm, by (10.3),

$$CD_1(L) = \pi_1\left(\tfrac{1}{2}, L\right) - \pi_1\left(\tfrac{\eta}{2}, L\right) = \frac{(\eta-1)^2}{4}, \qquad CD_2(H) = \pi_2\left(\tfrac{\eta}{2}, H\right) - \pi_2\left(\tfrac{1}{2}, H\right) = \frac{(\eta-1)^2}{4}. \tag{10.6}$$

These are the costs of deviation (5) and (6) of the paper.[1] Firm 1 does not deviate in state L, where the entrant otherwise believes $\rho = 1$, if

$$[\xi(\rho') - \xi(1)]\, G_1(L) = \lambda(1-\rho')\, G_1(L) \;\le\; CD_1(L), \tag{10.7}$$

and firm 2 does not deviate in state H, where the entrant otherwise believes $\rho = 0$, if

$$[\xi(0) - \xi(\rho')]\, G_2(H) = \lambda\rho'\, G_2(H) \;\le\; CD_2(H). \tag{10.8}$$

Substituting (10.4), (10.5), and (10.6), conditions (10.7) and (10.8) become $\rho' \ge 1 - \frac{a_1}{\gamma^2}$ and $\rho' \le \frac{a_2}{\gamma^2}$, where

$$a_1 = \frac{\eta^2(\eta-1)}{(\eta+1)\,\theta_1(2+\theta_1)}, \qquad a_2 = \frac{\eta-1}{(\eta+1)\,\theta_2(2-\theta_2)}. \tag{10.9}$$

A belief $\rho' \in [0,1]$ satisfying both exists if and only if $a_1 + a_2 \ge \gamma^2$, that is, if and only if $\gamma \le \gamma_S(\eta)$, where

$$\gamma_S(\eta) = \sqrt{\frac{\eta-1}{\eta+1}\left[\frac{\eta^2}{\theta_1(2+\theta_1)} + \frac{1}{\theta_2(2-\theta_2)}\right]}. \tag{10.10}$$

Dividing by λ, the condition $a_1 + a_2 \ge \gamma^2$ is condition (9) of the paper,

$$\xi(0) - \xi(1) \;\le\; \frac{CD_1(L)}{G_1(L)} + \frac{CD_2(H)}{G_2(H)}.$$

Result 3 (Theorem 1 of the paper). A separating equilibrium with unprejudiced beliefs exists if and only if $\gamma \le \gamma_S(\eta)$. Moreover:

- it can be sustained with the entrant believing firm 1 ($\rho' = 0$) if and only if $\gamma^2 \le a_1$;
- it can be sustained with the entrant believing firm 2 ($\rho' = 1$) if and only if $\gamma^2 \le a_2$;
- if $\max\{a_1, a_2\} < \gamma^2 \le a_1 + a_2$, it requires an intermediate belief $\rho' \in [1 - \frac{a_1}{\gamma^2}, \frac{a_2}{\gamma^2}]$.

In the baseline with $\eta = \frac{3}{2}$, $a_1 = \frac{4}{5}$, $a_2 = \frac{16}{35}$, and $\gamma_S = \sqrt{\frac{44}{35}} \approx 1.121$.

Figure 10.2 depicts the beliefs ρ' that sustain a separating equilibrium in the baseline with $\eta = \frac{3}{2}$. Firm 1 does not mimic the high state above the blue curve, and firm 2 does not mimic the low state below the red curve. For $\gamma \le 0.676$ every belief works; for $0.676 < \gamma \le 0.894$ the entrant can still believe firm 1 entirely; for $0.894 < \gamma \le 1.121$ only intermediate beliefs work; and at $\gamma_S = 1.121$ the only belief left is $\rho' = \frac{4}{11}$.

Intuitively, the entrant must assign some belief to $(\frac{\eta}{2}, \frac{1}{2})$, and that belief cannot be both 0 and 1, so at least one firm can move the belief in its favor by mimicking the other state. A separating equilibrium survives only if mimicking costs more than it brings for that firm. When the stakes are small, the entrant can believe one firm entirely, because that firm would never lie; when the stakes grow, the entrant must hold an intermediate belief so that each firm's deviation moves the belief

[1] In the paper, the second profit in (6) is written as $\pi_1(p_1^*(H), p_2^*(L), H)$; since $CD_2(H)$ is firm 2's cost, it should read $\pi_2(p_1^*(H), p_2^*(L), H)$.

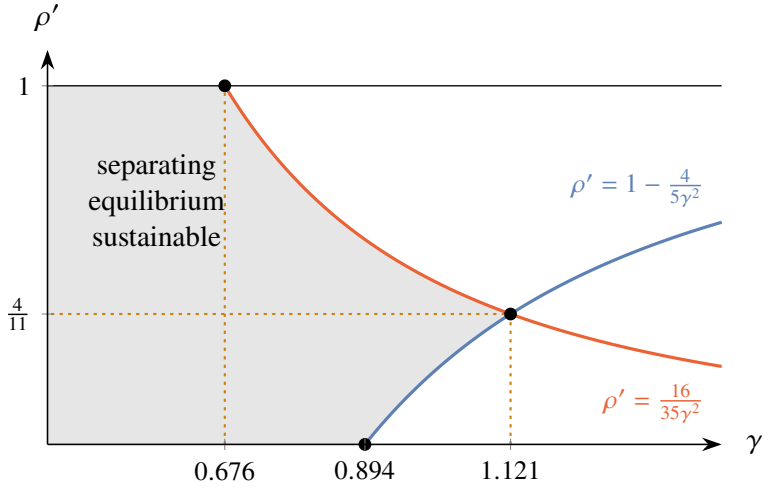


Figure 10.2: Beliefs after the off-path pair $(\frac{\eta}{2}, \frac{1}{2})$ that sustain a separating equilibrium ($\eta = \frac{3}{2}$, $\theta_1 = \theta_2 = \frac{1}{4}$).

by too little to be worthwhile. In Bagwell and Ramey (1991), both incumbents want the entrant to believe that entry is unattractive, so the entrant can treat a mismatched price pair as evidence of the attractive state, no incumbent gains from creating it, and separating equilibria always exist. With conflicting interests, any such belief rewards one of the two firms.

10.5 Third Step: Separation Requires Entry to Matter Little for One Incumbent

The threshold (10.10) adds two terms, one per incumbent: the first falls as entry matters more for firm 1 (θ_1 higher), and the second falls as entry matters more for firm 2 (θ_2 higher). The deviation cost (10.6) grows with the difference between the states.

Result 4 (comparative statics of separation). The threshold $\gamma_S(\eta)$ in (10.10):

- decreases in θ_1 and in θ_2, and increases in η;
- tends to infinity as $\theta_1 \to 0$ or as $\theta_2 \to 0$, so for any η and γ a separating equilibrium exists when θ_1 or θ_2 is small enough;
- tends to zero as $\eta \to 1$, so for any γ no separating equilibrium exists when η is close enough to 1.

Figure 10.3 depicts, for $\eta = \frac{3}{2}$, the pairs (θ_1, θ_2) for which a separating equilibrium exists, in a stationary market ($\gamma = 1$) and in a growing market ($\gamma = \frac{3}{2}$). The baseline $(\frac{1}{4}, \frac{1}{4})$ supports separation when $\gamma = 1$ but not when $\gamma = \frac{3}{2}$. In the growing market, separation holds for every θ_2 when $\theta_1 \leq 0.104$ and for every θ_1 when $\theta_2 \leq 0.049$, but fails at the baseline.

Intuitively, the entrant needs only one incumbent it can rely on. If entry hardly affects firm 1, firm 1 never gives up period-1 profit to manipulate the entrant, so the entrant can read the state from

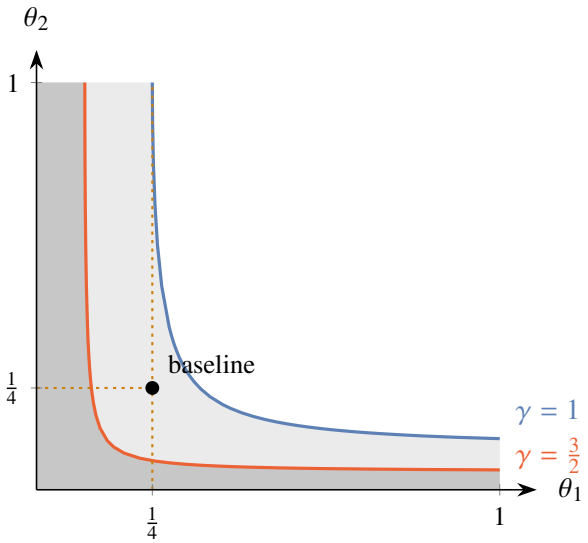


Figure 10.3: Pairs (θ_1, θ_2) for which a separating equilibrium exists when $\eta = \frac{3}{2}$: below and to the left of the blue curve when $\gamma = 1$ (light shading), and of the red curve when $\gamma = \frac{3}{2}$ (dark shading).

firm 1's price and ignore firm 2. The same holds if entry hardly affects firm 2. Separation fails when entry matters for both incumbents, as when the entrant's good is a close substitute for one incumbent's good and a strong complement of the other's, or when the market grows, which raises both stakes at once. A larger difference between the states helps separation, because mimicking the other state then requires a larger price change.

10.6 Fourth Step: Pooling Equilibria

Band conditions. In a pooling equilibrium both incumbents charge the same prices (p_1^0, p_2^0) in both states, and the entrant keeps its prior ρ^0. Consider off-path beliefs that punish the deviator: $\rho = 1$ after a deviation by firm 1 and $\rho = 0$ after a deviation by firm 2. A deviating firm then charges its best reply $\frac{\alpha_S}{2}$, so, by (10.3), firm 1 does not deviate if

$$\left(p_1^0 - \tfrac{1}{2}\right)^2 \le \lambda(1-\rho^0)\, G_1(L) \quad \text{and} \quad \left(p_1^0 - \tfrac{\eta}{2}\right)^2 \le \lambda(1-\rho^0)\, G_1(H), \tag{10.11}$$

since $\xi(\rho^0) - \xi(1) = \lambda(1-\rho^0)$, and firm 2 does not deviate if

$$\left(p_2^0 - \tfrac{\eta}{2}\right)^2 \le \lambda\rho^0\, G_2(H) \quad \text{and} \quad \left(p_2^0 - \tfrac{1}{2}\right)^2 \le \lambda\rho^0\, G_2(L), \tag{10.12}$$

since $\xi(0) - \xi(\rho^0) = \lambda\rho^0$. These are conditions (12) to (15) of the paper. Each inequality defines an interval, a *band*, centered on the full-information price of that state. By (10.4), the radius of each band in state H is η times its radius in state L, so firm 1's bands overlap if and only if $\frac{\eta-1}{2} \le (1+\eta)\sqrt{\lambda(1-\rho^0)G_1(L)}$, that is, $\gamma \ge \gamma_1(\eta)$, and firm 2's bands overlap if and only if

$\gamma \geq \gamma_2(\eta)$, where

$$\gamma_1(\eta) = \frac{\eta - 1}{(\eta+1)\sqrt{\lambda(1-\rho^0)\,\theta_1(2+\theta_1)}}, \qquad \gamma_2(\eta) = \frac{\eta - 1}{(\eta+1)\sqrt{\lambda\rho^0\,\theta_2(2-\theta_2)}}. \tag{10.13}$$

When a firm's two bands overlap, the overlap contains a price between $\frac{1}{2}$ and $\frac{\eta}{2}$, and hence in $[0, 1]$, since the low-state band contains $\frac{1}{2}$ and the high-state band contains $\frac{\eta}{2}$.

Result 5 (Theorem 2 of the paper). A pooling equilibrium exists if and only if $\gamma \geq \gamma_P(\eta)$, where $\gamma_P(\eta) \equiv \max\{\gamma_1(\eta), \gamma_2(\eta)\}$. In the baseline with $\rho^0 = \frac{1}{2}$,

$$\gamma_P(\eta) = \gamma_2(\eta) = \sqrt{\frac{32\,\eta^2(\eta-1)}{7(\eta+1)^3}},$$

and $\gamma_P(\eta) < \gamma_S(\eta)$ for every $\eta \in (1, 2)$. Hence:

- if $\gamma < \gamma_P(\eta)$, only separating equilibria exist;
- if $\gamma_P(\eta) \leq \gamma \leq \gamma_S(\eta)$, both separating and pooling equilibria exist;
- if $\gamma > \gamma_S(\eta)$, only pooling equilibria exist.

At $\eta = \frac{3}{2}$, $\gamma_P = \sqrt{\frac{288}{875}} \approx 0.574$ and $\gamma_S \approx 1.121$.

Figure 10.4 depicts the three regions in the (η, γ) plane, as in Figure 2 of the paper. For a given difference between the states, only separating equilibria exist in slowly growing or shrinking markets, only pooling equilibria exist in fast-growing markets, and both exist in between. The dot marks the example of the fifth step, $(\eta, \gamma) = (\frac{3}{2}, 2)$.

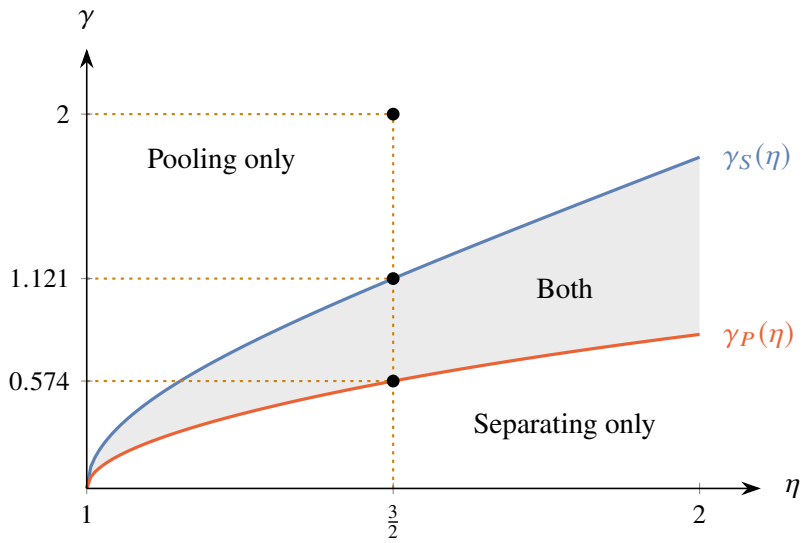


Figure 10.4: Existence of separating and pooling equilibria ($\theta_1 = \theta_2 = \frac{1}{4}$, $\rho^0 = \frac{1}{2}$).

Intuitively, a pooling equilibrium requires each incumbent to give up period-1 profit rather than let the entrant learn the state, and a separating equilibrium requires at least one incumbent to refuse to

do so. Both incentives are driven by the stakes G_1 and G_2, which grow with γ^2: when the future matters little, the incumbents set their full-information prices, and when the future is important, both are willing to manipulate the entrant and neither can be relied upon to reveal the state. Since $\gamma_P < \gamma_S$, a pure-strategy equilibrium exists for every (η, γ) in the baseline of the toy model, as in the paper's example of Section 6.1.

10.7 Fifth Step: Pooling Prices that Survive the Intuitive Criterion

The band conditions use the harshest off-path beliefs. The paper also requires *intuitive beliefs*: if a deviation could be profitable for a firm only in one state, the entrant should believe that state. Consider the example $\eta = \frac{3}{2}$, $\gamma = 2$, $\rho^0 = \frac{1}{2}$, and $\theta_1 = \theta_2 = \frac{1}{4}$, which lies in the pooling-only region of Figure 10.4. Then $\lambda = \frac{5}{9}$, $\lambda\rho^0 = \lambda(1-\rho^0) = \frac{5}{18}$, $G_1(L) = \frac{9}{16}$, $G_1(H) = \frac{81}{64}$, $G_2(L) = \frac{7}{16}$, and $G_2(H) = \frac{63}{64}$. By (10.11) and (10.12), the pooling prices are

$$p_1^0 \in \left[\tfrac{3}{4} - \tfrac{3\sqrt{10}}{16},\ \tfrac{1}{2} + \tfrac{\sqrt{10}}{8}\right] \approx [0.157, 0.895], \qquad p_2^0 \in \left[\tfrac{3}{4} - \tfrac{\sqrt{70}}{16},\ \tfrac{1}{2} + \tfrac{\sqrt{70}}{24}\right] \approx [0.227, 0.849]. \tag{10.14}$$

Firm 1. Firm 1 would like to convince the entrant that demand is high. A deviation that convinces it would raise firm 1's expected period-2 profit by $\lambda\rho^0 G_1(S)$. Condition (19) of the paper shows that no deviation is profitable for firm 1 only in the high state if and only if firm 1's low-state profit at the pooling price is below what it would gain in the low state from that belief:

$$p_1^0\left(1 - p_1^0\right) < \lambda\rho^0\, G_1(L) = \frac{5}{32}, \quad \text{that is,} \quad p_1^0 < \frac{1}{2} - \frac{\sqrt{6}}{8} \approx 0.194 \ \text{ or } \ p_1^0 > \frac{1}{2} + \frac{\sqrt{6}}{8} \approx 0.806. \tag{10.15}$$

Firm 2. Firm 2 would like to convince the entrant that demand is low, which raises its expected period-2 profit by $c_S \equiv \lambda(1-\rho^0)G_2(S)$, with $c_L = \frac{35}{288}$ and $c_H = \eta^2 c_L = \frac{35}{128}$. Its candidate deviations are price cuts, since the paper shows that a price increase cannot be profitable for firm 2 only in state L. The deepest cut it would accept in state L is to the price $\tilde{p} > 0$ that solves $(p_2^0 - \tilde{p})(1 - p_2^0 - \tilde{p}) = c_L$. In state H, the same cut costs $(p_2^0 - \tilde{p})(\eta - p_2^0 - \tilde{p})$, which exceeds the low-state cost by $(\eta - 1)(p_2^0 - \tilde{p})$, while the gain c_H exceeds c_L by $(\eta^2 - 1)c_L$. The deepest cut is the hardest one to accept in state H, so every cut that firm 2 would accept in state L is also acceptable in state H, and no deviation is profitable only in state L, if and only if $(\eta - 1)(p_2^0 - \tilde{p}) < (\eta^2 - 1)c_L$, that is, $p_2^0 - \tilde{p} < (\eta + 1)c_L$, which gives

$$p_2^0 < \frac{1}{2} - \frac{1 - (\eta+1)^2 c_L}{2(\eta+1)} = \frac{2603}{5760} \approx 0.452. \tag{10.16}$$

This is condition (22) of the paper. Its simpler sufficient condition (23) holds only for $p_2^0 < 0.212$, which lies outside firm 2's band in (10.14).

Result 6 (Theorem 3 of the paper). With $\eta = \frac{3}{2}$, $\gamma = 2$, $\rho^0 = \frac{1}{2}$, and $\theta_1 = \theta_2 = \frac{1}{4}$:

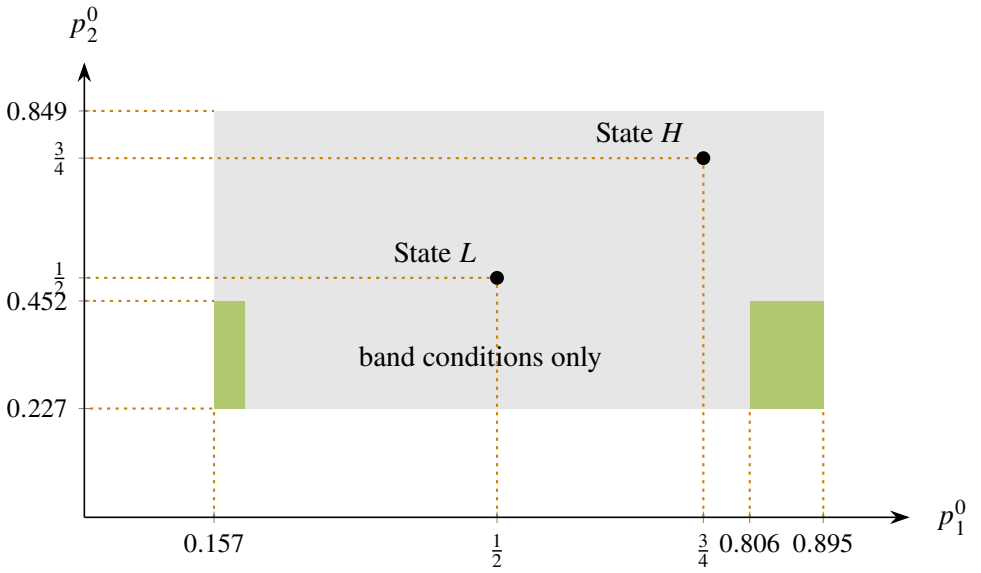


Figure 10.5: Pooling prices for $\eta = \frac{3}{2}$, $\gamma = 2$, $\rho^0 = \frac{1}{2}$: the light rectangle satisfies the band conditions, and the green blocks also satisfy the intuitive criterion.

- a price pair (p_1^0, p_2^0) is a pooling equilibrium with intuitive beliefs if and only if $p_1^0 \in [0.157, 0.194) \cup (0.806, 0.895]$ and $p_2^0 \in [0.227, 0.452)$;
- firm 2's pooling price lies below both of its full-information prices, $\frac{1}{2}$ and $\frac{3}{4}$, and firm 1's pooling prices in the upper interval lie above both;
- with $\eta = \frac{3}{2}$ and $\rho^0 = \frac{1}{2}$, firm 1's upper interval is nonempty if and only if $\gamma > \frac{4}{\sqrt{5}} \approx 1.789$; at $\gamma = \frac{3}{2}$, no pooling equilibrium satisfies the intuitive criterion.

Figure 10.5 depicts the pooling prices of (10.14) (light rectangle) and the pairs that also satisfy (10.15) and (10.16) (green blocks), together with the full-information price pairs.

Intuitively, firm 1 can keep a pooling price only if no price exists that firm 1 would choose in the high state but not in the low state to signal high demand. Condition (10.15) rules out such a price when firm 1 earns little in the low state at the pooling price, so that it would give up all of its low-state profit to be believed; this happens near the low-state choke price, or at very low prices. Firm 2 prices low to look like the low state, as in standard limit pricing, while firm 1 in the upper interval prices high to look like the high state and attract the entrant. The two price distortions go in opposite directions, so the threat of entry need not benefit consumers. In addition, the entry decision is distorted: entry occurs with probability $\xi(\frac{1}{2}) = \frac{13}{18}$ in both states, instead of $\frac{4}{9}$ in the low state and 1 in the high state under full information.

10.8 Sixth Step: What the Toy Model Keeps and Drops

Table 10.1 reports which equilibria exist in the baseline with $\eta = \frac{3}{2}$ and $\rho^0 = \frac{1}{2}$ for three growth rates.

Table 10.1: Equilibria of the toy model ($\eta = \frac{3}{2}$, $\theta_1 = \theta_2 = \frac{1}{4}$, $\rho^0 = \frac{1}{2}$).

Market growth	Separating	Pooling	Pooling with intuitive beliefs
Shrinking, $\gamma = \frac{1}{2}$	yes	no	no
Stationary, $\gamma = 1$	yes	yes	no
Growing, $\gamma = 2$	no	yes	yes

What the toy model keeps.

- Two informed incumbents with opposite interests about entry, and an entrant that must decide whom to believe when their prices disagree.
- Separating equilibria at full-information prices, which exist only if mimicking the other state is costly, relative to the stakes from entry, for at least one incumbent.
- Pooling as the only outcome in growing markets or when the states are close, with pooling prices that can move in opposite directions for the two incumbents. As the paper notes, unlike in Bagwell and Ramey (1991), the incumbents would not coordinate a joint deviation even if they could, since they disagree about entry.

What it drops.

- Strategic interaction between the incumbents in period 1: here their goods are independent, while in the paper they are complements and best replies slope.
- The period-2 price game: entry here shifts the incumbents' demands instead of adding a third firm to a differentiated Bertrand market (Section 6.1 of the paper) or a Cournot market (Section 6.2); the vertical structure of Section 6.3 is also left out.
- More than one incumbent of each type, coordinated deviations, and mixed strategies, which the paper discusses but does not solve.

10.9 Connections to Other Chapters

- In Chapter 9, several incumbents share private information about a common cost and all lose from entry, so a high price that signals high costs can keep the entrant out even though they do not collude.
- The last section of Chapter 30 studies several senders with opposed interests who compete to persuade a receiver, the information-design counterpart of the two incumbents in this chapter.
- Chapter 24 studies a single firm that signals its quality with its price, where separating requires a price distortion.

References

Bagwell, K., and Ramey, G. (1991). Oligopoly limit pricing. *RAND Journal of Economics*, 22(2), 155–172.

Harrington, J. E. (1987). Oligopolistic entry deterrence under incomplete information. *RAND Journal of Economics*, 18(2), 211–231.

Milgrom, P., and Roberts, J. (1982). Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica*, 50(2), 443–459.

Schultz, C. (1999). Limit pricing when incumbents have conflicting interests. *International Journal of Industrial Organization*, 17(6), 801–825.

Part III

Mergers

The eight chapters in this part study how a competition authority should review a horizontal merger, and each takes up a difficulty that a simple test of the merger's effect on consumers leaves out. Chapter 11 asks how the authority should treat cost savings that only the merging firms know and must prove at a cost: the best rule approves every merger, blocks every merger, or approves only the mergers whose efficiencies are proven. Chapter 12 measures the price effect that such a rule weighs against the efficiencies, the upward pricing pressure of a merger, which pass-through turns into a predicted price change. Chapters 13 and 14 study what the authority does not see when it reviews a single merger. In Chapter 13, mergers arrive in sequence, yet an authority that approves each merger if and only if it does not lower consumer surplus today does as well as one that knows every future merger. In Chapter 14, a target chooses among several partners, and larger mergers should face a tougher standard.

Chapters 15 to 17 study how a merger changes the incentive to innovate. In Chapter 15, firms improve the quality of their products. Price coordination raises the reward from innovating, but the merged firm also counts the sales that an improved product takes from its other product, and this second force tends to dominate. In Chapter 16, firms invest to lower their costs, so the return to investment grows with output; a merger lowers output, so the merging firms invest less, and efficiency gains large enough to restore investment can still leave consumers worse off. In Chapter 17, labs race to discover a new product. A merged firm that keeps both labs equally busy searches less, but it may instead concentrate its effort in one lab, avoid duplicated discoveries, and raise the probability that the product is found.

Chapter 18 turns to markets in which switching costs lock consumers in, such as banking or insurance. The average markup remains proportional to the Herfindahl index, with a factor that depends on the switching cost and on how patient firms and consumers are. A merger of two low-cost firms raises prices the most, while a merger of two high-cost firms can raise welfare without any cost savings.

Chapter 11

Should Merger Control Allow an Efficiency Defense?

Companion note to Lagerlöf, J. N. M., and P. Heidhues (2005), "On the Desirability of an Efficiency Defense in Merger Control," International Journal of Industrial Organization 23(9–10): 803–827.

11.1 Introduction

Merging firms often claim that the merger will lower their costs. Williamson (1968) showed that such cost savings can outweigh the loss from having one competitor less, and the merger guidelines in the US and the EU let merging firms make this argument, known as an efficiency defense. Fisher (1987) warned, however, that the defense invites rent seeking: firms spend resources to persuade the authority. Should merger control allow an efficiency defense?

Lagerlöf and Heidhues (2005) study a model in which the merging firms know whether their efficiencies are high or low, but must spend resources to produce hard, verifiable evidence of them. Society commits in advance to a rule, which the paper calls an institution, that approves the merger depending on the evidence it receives. Evidence helps society separate good mergers from bad ones, but it is costly, and the firms produce it to obtain approval, not to inform society. In particular, firms with low efficiencies never reveal them.

The best institution is always one of three: approve every merger (laissez-faire), block every merger (the strict regime), or approve the merger if and only if the firms prove high efficiencies (the hard evidence regime). The last one, the only one with an efficiency defense, is best when the interests of society and the firms are aligned enough, for instance when high efficiencies lower the price, and when society is pessimistic about the merger. When society is optimistic, an efficiency defense can be harmful even when firms spend little on evidence, since it blocks good mergers whenever no evidence is found.

The toy model uses a quadratic cost of evidence. The first step shows where the gains from a merger

come from in a Cournot market, the second step solves for the firms' evidence production, the third step shows that the optimal institution is one of these three, the fourth step ranks them, the fifth step allows for a general cost of evidence, and the sixth step lists what the toy model keeps and drops.

11.2 The Toy Model

Players and information. Two firms have proposed to merge and act as one unit, the insiders. The post-merger efficiency is either high, H, or low, L. The insiders know it, while society assigns probability $p \in (0,1)$ to H. The insiders' knowledge is soft: they cannot credibly communicate it unless they find hard evidence.

Payoffs. Relative to blocking the merger, approving it changes the insiders' profits by $\Delta\pi_i$ and total surplus by ΔW_i when efficiency is $i \in \{L, H\}$. We normalize the insiders' gain from a high-efficiency merger to one and consider

$$\Delta\pi_H = 1, \qquad \Delta\pi_L > 0, \qquad \Delta W_H = w > 0, \qquad \Delta W_L = -\frac{1}{2}, \tag{11.1}$$

so the insiders always want to merge, while society wants only high-efficiency mergers. Parameter w measures how much society values a high-efficiency merger relative to the insiders.

Evidence. Insiders who know that efficiency is i choose the probability $\tau_i \in [0,1]$ of finding hard evidence of it, at cost $C(\tau_i) = k\tau_i^2$, where $k > \frac{1}{2}$. Evidence found can be disclosed or hidden. The cost of evidence counts fully in welfare, which is total surplus minus this cost.

Institutions and timing. (1) Society commits to a merger control institution $z = (z_L, z_H, z_N)$: the probabilities of approving the merger after evidence of low efficiency, after evidence of high efficiency, and after no evidence. (2) The insiders learn their type and choose τ_i. (3) Insiders who found evidence decide whether to disclose it. (4) Society applies z. Three institutions play a key role: laissez-faire, LF, with $z = (1,1,1)$; the strict regime, SR, with $z = (0,0,0)$; and the hard evidence regime, HE, with $z = (0,1,0)$. An institution allows for an efficiency defense if evidence changes the approval probability, $z_L \neq z_N$ or $z_H \neq z_N$; of the three, only HE does.

11.3 First Step: Who Gains from a Merger

The paper does not model the market, but it illustrates the payoffs with a Cournot market (its Figure 1). Consider three symmetric firms with inverse demand $P = a - Q$ and marginal cost c, with $a - c = 1$. Before the merger, each firm produces $\frac{1}{4}$ and earns $\frac{1}{16}$, and consumer surplus is $\frac{9}{32}$. Two firms merge, and the merged firm's marginal cost falls to $c - e$, with $e \in [0,1)$. In the resulting Cournot duopoly, the merged firm produces $\frac{1+2e}{3}$, the outsider produces $\frac{1-e}{3}$, and the price is $c + \frac{1-e}{3}$. Therefore, the merger changes the insiders' profits, the outsider's profit, consumer

surplus, and total surplus by

$$\begin{aligned} \Delta\pi_I(e) &= \frac{(1+2e)^2}{9} - \frac{1}{8}, & \Delta\pi_O(e) &= \frac{(1-e)^2}{9} - \frac{1}{16}, \\ \Delta CS(e) &= \frac{(2+e)^2}{18} - \frac{9}{32}, & \Delta W(e) &= \frac{176e^2 + 128e - 7}{288}. \end{aligned} \tag{11.2}$$

Result 1 (gains from a merger, Figure 1 of the paper). In the Cournot market:

- the insiders gain from merging if and only if

$$e > e^\circ = \frac{3\sqrt{2} - 4}{8} \approx 0.03;$$

- total surplus rises if and only if

$$e > e' = \frac{\sqrt{333} - 16}{44} \approx 0.05;$$

- the price falls if and only if $e > e'' = \frac{1}{4}$; at this threshold, consumer surplus and the outsider's profit change sign in opposite directions, and total surplus rises by more than the insiders' profits if and only if $e > \frac{1}{4}$.

Figure 11.1 depicts the four gains in (11.2). A low efficiency $e_L \in (e^\circ, e')$ and a high efficiency $e_H > e'$ produce the signs in (11.1). Following Farrell and Shapiro (1990), the merger has negative externalities if $\Delta W_H < \Delta\pi_H$, that is, $w < 1$, and positive externalities if $w > 1$. In the Cournot market, a high-efficiency merger has positive externalities if and only if it lowers the price, $e_H > \frac{1}{4}$.

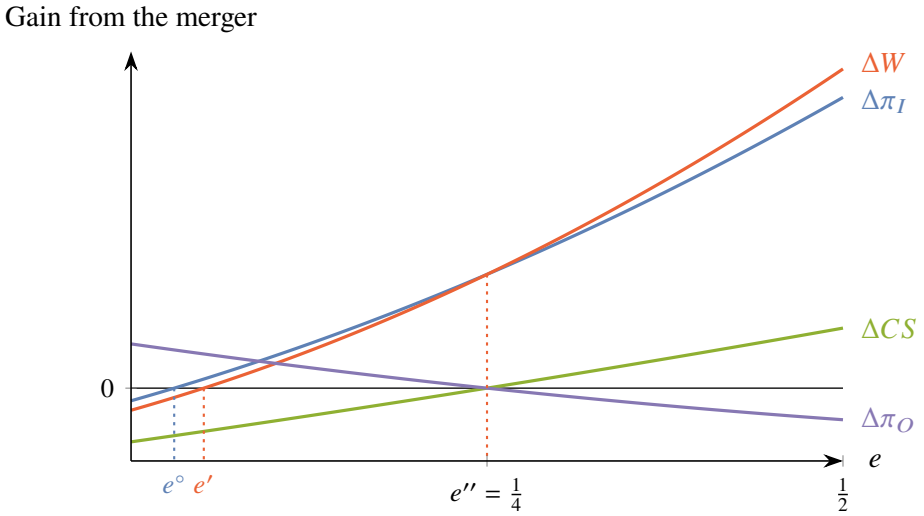


Figure 11.1: Gains from the merger for the insiders ($\Delta\pi_I$), society (ΔW), consumers (ΔCS), and the outsider ($\Delta\pi_O$) in a three-firm Cournot market.

Intuitively, the merger removes a competitor, which raises the price unless the cost savings are large. The insiders gain from both the cost savings and the higher price, so they want to merge even

when total surplus falls, $e^\circ < e < e'$. The difference between society's and the insiders' gains, $\Delta CS + \Delta\pi_O$, is the effect on everybody else; it is positive only when the price falls, since a lower price helps consumers by more than it hurts the outsider.

11.4 Second Step: Evidence Production

Insiders who found evidence of efficiency i disclose it if $z_i > z_N$ and hide it if $z_i < z_N$. Society never sets $z_L > z_N$, since that would approve more low-efficiency mergers and pay for evidence it does not want, so low-efficiency insiders find nothing worth disclosing. Insiders who know that efficiency is high, facing $z_H \geq z_N$, choose τ to maximize their expected gain

$$[\tau z_H + (1-\tau) z_N]\, \Delta\pi_H - k\tau^2,$$

whose first-order condition $z_H - z_N = 2k\tau$ yields

$$\tau_H^* = \frac{z_H - z_N}{2k}. \tag{11.3}$$

Result 2 (evidence production, equation (3) of the paper). Given an institution z with $z_L \leq z_N \leq z_H$:

- low-efficiency insiders do not invest in evidence, $\tau_L^* = 0$;
- high-efficiency insiders invest $\tau_H^* = \frac{z_H - z_N}{2k}$, which increases in the gap $z_H - z_N$;
- no evidence is produced under LF or SR, and the most is produced under HE, $\tau^{HE} = \frac{1}{2k}$, at cost $C\left(\tau^{HE}\right) = \frac{1}{4k}$.

Intuitively, evidence is worth producing only when it changes the approval probability. Its marginal benefit is the gap $z_H - z_N$ times the insiders' gain from the merger, so an institution that ignores evidence, as LF and SR do, induces none.

11.5 Third Step: The Optimal Institution Is a Corner

Normalize total surplus without the merger to zero, so expected welfare under SR is $EW^{SR} = 0$, and under LF it is $EW^{LF} = pw - \frac{1-p}{2}$. Under an institution z, the merger is approved with probability

$$\tau_H^* z_H + \left(1 - \tau_H^*\right) z_N$$

when efficiency is high and with probability z_N when it is low. Substituting (11.3), expected welfare becomes

$$EW(z) = (1 - z_N)\, EW^{SR} + z_N\, EW^{LF} + p\, \frac{(z_H - z_N)^2 (2w - 1)}{4k}. \tag{11.4}$$

The last term is the net value of evidence. By (11.3), the insiders' benefit from evidence, $\tau_H^*(z_H - z_N)$, equals twice its cost; society values the same approvals at w per unit, so evidence is worth $2wC$ to society and costs C.

Result 3 (corner solution, Proposition 1 of the paper). The optimal institution is SR, LF, or HE:

- if $w < \frac{1}{2}$, the last term in (11.4) is negative unless $z_H = z_N$, so society induces no evidence, and welfare is linear in z_N, which yields SR or LF;
- if $w > \frac{1}{2}$, welfare is convex in (z_H, z_N), so its maximum over $0 \le z_N \le z_H \le 1$ is at a vertex: SR, LF, or HE;
- if $w \neq \frac{1}{2}$, an institution with $0 < z_H - z_N < 1$ is never optimal.

Figure 11.2 depicts the welfare gain over SR of the institutions $(0, z_H, 0)$, with $p = \frac{1}{4}$ and $k = 1$. When $w = \frac{3}{4}$, the gain $\frac{z_H^2}{32}$ rises in z_H, so society chooses $z_H = 1$, which is HE; when $w = \frac{1}{3}$, the gain $-\frac{z_H^2}{48}$ falls, so society chooses $z_H = 0$, which is SR.

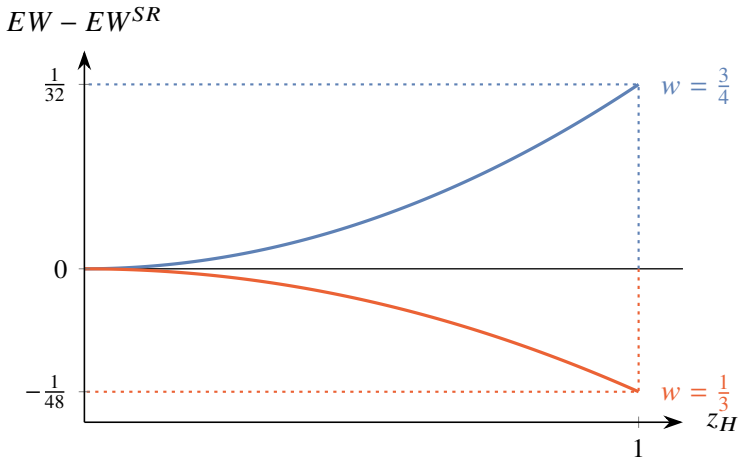


Figure 11.2: Welfare gain over SR of the institution $(0, z_H, 0)$, with $p = \frac{1}{4}$ and $k = 1$.

Intuitively, both the insiders' cost of evidence and society's benefit from it scale with the same gap $z_H - z_N$, so the sign of the net value of evidence does not depend on how much evidence society induces: society either encourages evidence as much as it can or not at all. Moreover, an institution with $z_H < 1$ throws away costly information, since it sometimes learns that the merger is efficient and still blocks it.

11.6 Fourth Step: Ranking the Three Institutions

By (11.4), the gains of LF and HE over SR are

$$EW^{LF} - EW^{SR} = pw - \frac{1-p}{2} \qquad \text{and} \qquad EW^{HE} - EW^{SR} = \frac{p(2w-1)}{4k}. \tag{11.5}$$

LF dominates SR if and only if $p > \bar{p}$, where

$$\bar{p} \equiv \frac{1}{2w+1}.$$

As in the paper, Regime I, $p < \bar{p}$, is the case where society is pessimistic about the merger, and Regime II, $p > \bar{p}$, the case where it is optimistic.

Result 4 (optimal institution, Propositions 2 and 3 of the paper). With a quadratic cost of evidence:

- in Regime I, SR is optimal if $w < \frac{1}{2}$ and HE if $w > \frac{1}{2}$;
- in Regime II, LF is optimal if $w < \frac{1}{2}$; if $w > \frac{1}{2}$, HE is optimal for $k < \hat{k}$ and LF for $k > \hat{k}$, where
$$\hat{k} \equiv \frac{p(2w-1)}{4\left[pw - \frac{1-p}{2}\right]};$$
- in Regime II, HE can be optimal only if $p < \frac{1}{2}$, since otherwise $\hat{k} \leq \frac{1}{2}$;
- positive externalities, $w > 1$, are sufficient for HE to dominate SR.

Figure 11.3 depicts the optimal institution with $p = \frac{1}{4}$, where $\bar{p} = p$ at $w = \frac{3}{2}$ and
$$\hat{k} = \frac{2w-1}{2(2w-3)}.$$
For instance, with $w = 2$ and $k = 1$, society is optimistic, but $k < \hat{k} = \frac{3}{2}$: HE yields $\frac{3}{16}$ over SR and LF yields $\frac{1}{8}$, so HE is optimal.

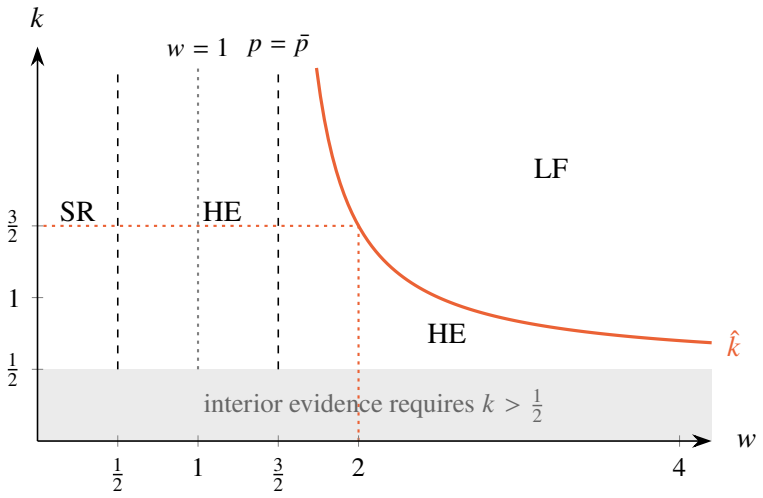


Figure 11.3: Optimal institution with $p = \frac{1}{4}$. To the left of $w = \frac{3}{2}$, society is pessimistic (Regime I); to the right, optimistic (Regime II).

Intuitively, with a quadratic cost the insiders spend half of their benefit from evidence, so evidence is worth its cost to society if and only if society values a high-efficiency merger at more than half as much as the insiders do, $w > \frac{1}{2}$. If a high-efficiency merger lowers the price, $w > 1$, society values it more than the insiders, who then produce too little evidence, and society should encourage evidence as much as it can. The two regimes differ because of the insiders' vested interest. When society is pessimistic, it blocks the merger without evidence and needs to learn about the high state, which the insiders are glad to prove. When society is optimistic, it would rather learn about the low state, which the insiders never reveal; HE then blocks a good merger whenever no evidence is found, which is more likely when k is large.

The paper's Example. With $w = 1$ and $p = \frac{1}{2}$ (the paper's numbers, $\Delta W_H = \Delta\pi_H = 200$ and $\Delta W_L = -100$, divided by 200), we obtain $\bar{p} = \frac{1}{3}$ and $\hat{k} = \frac{1}{2}$, so LF $\succ$ HE $\succ$ SR for every $k > \frac{1}{2}$. The expected cost of evidence under HE, $\frac{1}{8k}$, vanishes as k grows, yet LF's advantage, $\frac{1}{4} - \frac{1}{8k}$, grows. An efficiency defense can thus be harmful even when firms spend almost nothing on it: it blocks too many good mergers.

11.7 Fifth Step: Extension to a General Cost Elasticity

Consider now $C(\tau) = k\tau^{\eta}$, whose elasticity $\eta > 1$ is constant, with $k > k^{\circ} \equiv \frac{1}{\eta}$, so that evidence production under HE is interior. The baseline is $\eta = 2$. The first-order condition under HE gives

$$\tau^{HE} = \left(\frac{1}{k\eta}\right)^{\frac{1}{\eta-1}}$$

and

$$\eta = \frac{\tau^{HE}\Delta\pi_H}{C\left(\tau^{HE}\right)}, \tag{11.6}$$

so η is the ratio of the insiders' benefit from evidence to its cost (equation (9) of the paper). Therefore,

$$EW^{HE} - EW^{SR} = p\,C\left(\tau^{HE}\right)(\eta w - 1).$$

Result 5 (Propositions 2 and 3 of the paper). With $C(\tau) = k\tau^{\eta}$:

- HE dominates SR if and only if $\eta > \frac{1}{w} = \frac{\Delta\pi_H}{\Delta W_H}$;
- in Regime I, SR is optimal if $\eta < \frac{1}{w}$ and HE if $\eta > \frac{1}{w}$;[1]
- in Regime II, LF is optimal if $\eta < \frac{1}{w}$, or if $\eta > \frac{1}{w}$ and $k > \hat{k}$; HE is optimal if $\eta > \frac{1}{w}$ and $k < \hat{k}$, where

$$\hat{k} \equiv \left(\frac{1}{\eta}\right)^{\eta}\left[\frac{p(\eta w - 1)}{pw - \frac{1-p}{2}}\right]^{\eta-1},$$

  which reduces to the cutoff in Result 4 when $\eta = 2$.

Figure 11.4 depicts the optimal ranking in the (η, k) plane, as in Figure 2 of the paper, with negative externalities, $w = \frac{1}{3}$, so HE dominates SR if and only if $\eta > 3$. Then $\bar{p} = \frac{3}{5}$; Panel A sets $p = \frac{1}{2}$ (Regime I) and Panel B sets $p = \frac{2}{3}$ (Regime II). In both panels, $\hat{k}$ crosses k° at

$$\eta = \frac{p\,\Delta\pi_H}{(1-p)(-\Delta W_L)},$$

which is 2 in Panel A and 4 in Panel B.

[1] The published statement of Proposition 3 reads "HE if $\eta < \Delta\pi_H/\Delta W_H$" in this case; the inequality should be reversed, as Proposition 2 and the discussion before Proposition 3 show.

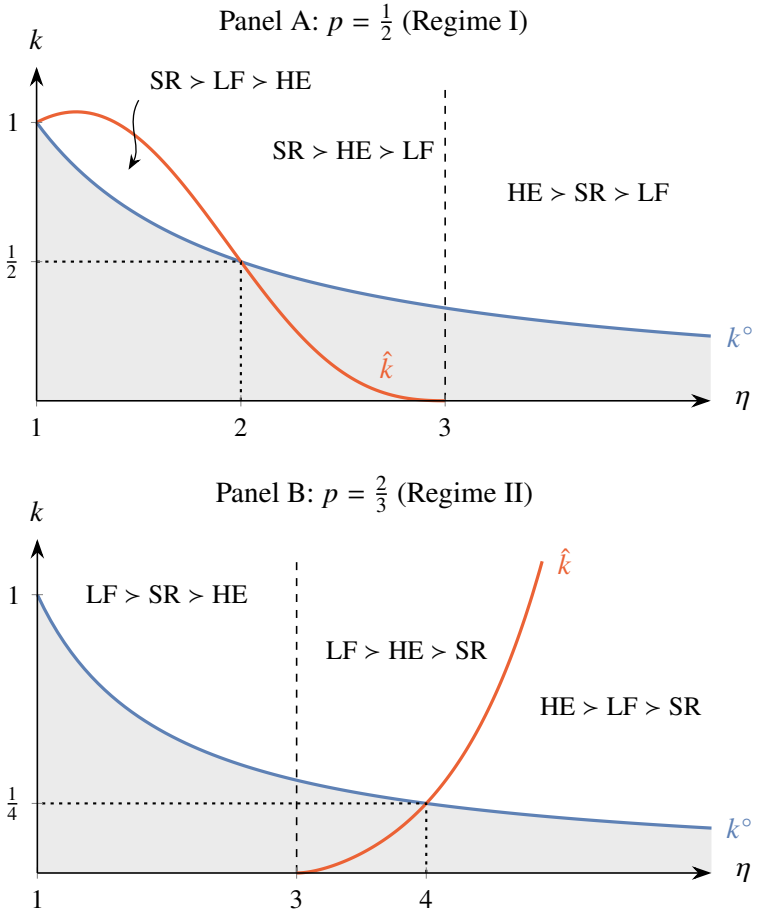


Figure 11.4: Ranking of the three institutions with $w = \frac{1}{3}$ and $\Delta W_L = -\frac{1}{2}$, for $k > k^\circ$; the shaded area, $k < k^\circ$, is excluded. The dashed line is $\eta = \frac{1}{w} = 3$.

Intuitively, by (11.6), a larger η means that the insiders' benefit from evidence exceeds its cost by more, so evidence costs little relative to the information it delivers, and HE can dominate SR even with negative externalities. The asymmetry between the panels is the one in the fourth step: when society is optimistic, HE also needs a small k, so that the insiders find evidence with a high enough probability and good mergers are unlikely to be blocked. As p approaches one, LF is always optimal, while as p approaches zero, HE is optimal if and only if $\eta > \frac{1}{w}$.

Section 5 of the paper turns these results into a three-step procedure for an antitrust authority, assuming that the region of Panel B where HE is optimal does not arise in practice. First, if it expects the merger to raise total surplus, it approves it. Second, if a surplus-increasing merger would also lower the price, it asks for evidence and approves the merger if and only if the evidence shows high efficiencies. Third, it asks for evidence only if the cost elasticity is large relative to the merger's price increase, and otherwise blocks the merger.

11.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Williamson's trade-off between market power and efficiencies, with the paper's Cournot illustration.
- The insiders' private information about efficiencies and their costly evidence production.
- The corner result: society either induces no evidence or as much as it can.
- The roles of the alignment between society and insiders, of the prior, and of the cost elasticity, including the asymmetry between pessimistic and optimistic regimes and the paper's Example.

What it drops.

- General convex costs with a weakly increasing elasticity, under which the corner result (Proposition 1) still holds.
- Low-efficiency insiders who fabricate evidence (Section 6.4 and Appendix A of the paper), which makes the conditions for HE more stringent.
- A consumer welfare standard, other parties who produce evidence, and costly verification of reports (Sections 6.2, 6.3, and 6.5 of the paper).

11.9 Connections to Other Chapters

- In Chapter 18, a merger of two high-cost firms raises welfare without any cost savings, because it shifts sales toward the low-cost rivals.
- The fifth step of Chapter 12 adds cost efficiencies to the pricing-pressure analysis of a merger.
- Chapter 13 studies an authority that reviews mergers one at a time under a consumer surplus standard, and shows that a myopic rule is optimal.

- Chapter 14 studies an authority that sees only the merger the firms propose, and shows that larger mergers should face a tougher standard.
- In Chapter 31, applicants to an agency also spend resources on the evidence they submit, but to make bad projects hard to evaluate.

References

Farrell, J., and C. Shapiro (1990). Horizontal mergers: An equilibrium analysis. *American Economic Review*, 80(1), 107–126.

Fisher, F. M. (1987). Horizontal mergers: Triage and treatment. *Journal of Economic Perspectives*, 1(2), 23–40.

Lagerlöf, J. N. M., and P. Heidhues (2005). On the desirability of an efficiency defense in merger control. *International Journal of Industrial Organization*, 23(9–10), 803–827.

Williamson, O. E. (1968). Economies as an antitrust defense: The welfare tradeoffs. *American Economic Review*, 58(1), 18–36.

Chapter 12

Pricing Pressure and Pass-Through: A First-Order Approach to Mergers

Companion note to Jaffe, S., and E. G. Weyl (2013), "The First-Order Approach to Merger Analysis," American Economic Journal: Microeconomics 5(4): 188–218.

12.1 Introduction

Consider two coffee shops on the same street that merge. Before the merger, when one shop raised its price, some customers walked across the street, and the shop lost those sales. After the merger, those customers still buy from the merged firm. Raising either price is therefore less costly than before, and prices rise. By how much do they rise, and how much do consumers lose?

Farrell and Shapiro (2010) measured this incentive with the upward pricing pressure, UPP: the diversion ratio, that is, the share of the sales lost by one product that the partner's product captures, times the partner's markup. UPP works like a new cost: every sale of one product now gives up the markup the firm would earn if that customer bought the partner's product instead. The 2010 US Horizontal Merger Guidelines incorporated this measure, but UPP only tells us whether prices rise, not by how much.

Jaffe and Weyl (2013) turn UPP into a quantitative prediction. Since the pricing pressure acts like a cost increase, the price change is approximately the pricing pressure times a pass-through rate, the share of a cost increase that the firm passes on to its price. Pass-through depends on the curvature of demand, which data from before the merger do not reveal, so two markets with the same UPP can see very different price changes. The paper also extends UPP to conduct other than Bertrand competition and converts price changes into changes in consumer and total surplus.

The toy model is a symmetric duopoly in which the curvature of demand can vary while prices, quantities, markups, and diversion ratios before the merger stay fixed. The first step describes the premerger market, the second step computes the pricing pressure, the third step turns it into price changes through pass-through, the fourth and fifth steps allow for Cournot conduct and cost

efficiencies, the sixth step computes welfare changes, and the seventh step lists what the toy model keeps and drops.

12.2 The Toy Model

Firms and demand. Two single-product firms, $i, j \in \{1, 2\}$, sell differentiated substitutes at constant marginal cost $c = \frac{1}{2}$. Demand for product i is

$$Q_i(P_i, P_j) = \left[1 + 2\theta\left((1 - P_i) - \frac{1}{5}(1 - P_j)\right)\right]^{\frac{1}{\theta}}, \tag{12.1}$$

where $\theta > 0$ indexes the curvature of demand. With $\theta = 1$, demand is linear, and as θ approaches zero, it becomes exponential:

$$Q_i = \frac{13}{5} - 2P_i + \frac{2}{5}P_j \qquad \text{and} \qquad Q_i = e^{2\left[(1-P_i)-\frac{1}{5}(1-P_j)\right]},$$

respectively; larger values of θ make it more concave.

Merger. The two firms merge and the merged firm sets both prices to maximize joint profits. The merger leaves marginal costs unchanged, except in the fifth step.

12.3 First Step: The Premerger Market

Under Bertrand competition, firm i maximizes $(P_i - c)Q_i$. As in the paper, write its first-order condition so that it is linear in marginal cost,

$$f_i(P) \equiv -\left(\frac{\partial Q_i}{\partial P_i}\right)^{-1} Q_i - (P_i - c) = \frac{1}{2} + \theta\left[(1 - P_i) - \frac{1}{5}(1 - P_j)\right] - \left(P_i - \frac{1}{2}\right) = 0. \tag{12.2}$$

Result 1 (premerger equilibrium). For every $\theta > 0$:

- both firms charge $P^0 = 1$ and sell $Q^0 = 1$, with markup $m = P^0 - c = \frac{1}{2}$;
- the demand slopes at these prices are

$$\frac{\partial Q_i}{\partial P_i} = -2 \qquad \text{and} \qquad \frac{\partial Q_i}{\partial P_j} = \frac{2}{5},$$

  so the diversion ratio, the share of the sales lost by product i that product j captures, is

$$D = -\frac{\partial Q_j/\partial P_i}{\partial Q_i/\partial P_i} = \frac{1}{5}.$$

Intuitively, θ changes only the curvature of demand away from the premerger prices. An analyst who observes prices, quantities, markups, and diversion ratios before the merger cannot tell the members of this family apart.

12.4 Second Step: Pricing Pressure

After the merger, raising P_1 also raises the profits on product 2. Dividing the merged firm's first-order condition for P_1 by $-\frac{\partial Q_1}{\partial P_1}$ gives $h_1(P) = f_1(P) + g_1(P)$, where the generalized pricing pressure (GePP) is the change in the first-order condition,

$$g_1(P) \equiv h_1(P) - f_1(P) = -\frac{\partial Q_2/\partial P_1}{\partial Q_1/\partial P_1}\,(P_2 - c). \tag{12.3}$$

Result 2 (upward pricing pressure, Proposition 1 of the paper). Under Bertrand conduct:

- the pricing pressure on each product equals the UPP, the diversion ratio times the partner's markup,
$$g = D\,(P^0 - c) = \frac{1}{5}\cdot\frac{1}{2} = \frac{1}{10};$$
- the pricing pressure is the same for every θ.

Intuitively, after the merger each unit that product 1 loses to product 2 is no longer a loss for the firm: it earns product 2's markup. A sale of product 1 thus carries an opportunity cost of $\frac{1}{10}$, which acts like a tax of $\frac{1}{10}$ on each product. Since UPP depends only on premerger observables, it cannot depend on θ.

12.5 Third Step: From Pricing Pressure to Price Changes

Theorem 1 of the paper approximates the price change by the pricing pressure times the merger pass-through matrix,

$$\rho_M = -\left(\frac{\partial h}{\partial P}\right)^{-1},$$

that is,

$$\Delta P \approx -\left(\frac{\partial h}{\partial P}(P^0)\right)^{-1} g(P^0). \tag{12.4}$$

In our family, the first-order condition h is linear along symmetric prices,

$$h(P) = \frac{9}{10} - \frac{4}{5}P + \frac{4\theta}{5}(1 - P),$$

so (12.4) is exact. A symmetric pressure is passed through at rate $\frac{5}{4(1+\theta)}$, and the postmerger price is $P^M = 1 + \frac{1}{8(1+\theta)}$.

Result 3 (price changes and pass-through, Theorem 1 and Section IV of the paper). With the same UPP of $\frac{1}{10}$:

- the merger raises both prices by

$$\Delta P = \frac{5}{4(1+\theta)} \cdot \frac{1}{10} = \frac{1}{8(1+\theta)},$$

  that is, by $\frac{1}{8}$ with exponential demand, $\frac{1}{12}$ with $\theta = \frac{1}{2}$, $\frac{1}{16}$ with linear demand, and $\frac{1}{24}$ with $\theta = 2$;
- the premerger pass-through rate of a tax on both products is $\frac{5}{5+4\theta}$, the merger pass-through rate is $\frac{5}{4(1+\theta)}$, and the postmerger pass-through rate is $\frac{1}{1+\theta}$;
- the postmerger rate is below the premerger rate, which is below the merger rate, so using either observable rate understates the price change; with linear demand, they predict $\frac{1}{18}$ and $\frac{1}{20}$ instead of $\frac{1}{16}$.

Figure 12.1 depicts the merged firm's first-order condition along symmetric prices for $\theta = \frac{1}{2}$ and $\theta = 2$. Both lines start from the same pricing pressure, $\frac{1}{10}$, at $P^0 = 1$, but the flatter one, with higher pass-through, reaches zero at $\frac{13}{12}$, while the steeper one reaches zero at $\frac{25}{24}$.

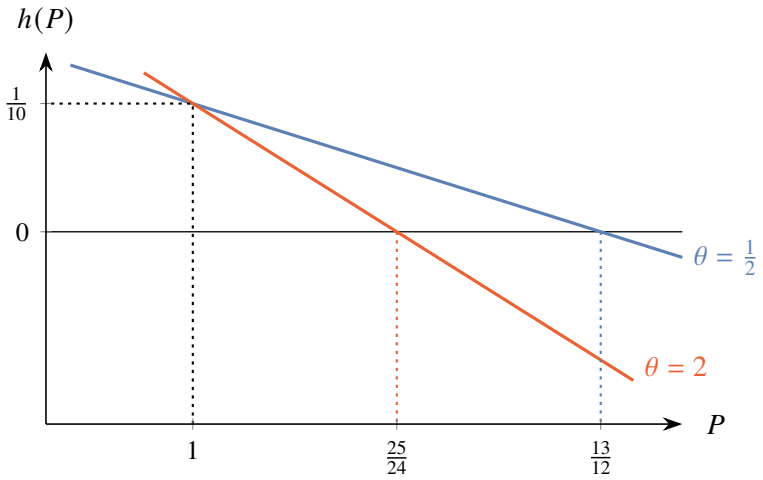


Figure 12.1: The merged firm's first-order condition along symmetric prices. Its value at $P^0 = 1$ is the pricing pressure; it crosses zero at the postmerger price.

Figure 12.2 depicts the price change and its two approximations based on observable pass-through rates as functions of θ. All three fall as demand becomes more concave.

Intuitively, the pricing pressure is a new opportunity cost, so the price responds to it as it would to a cost increase: by the pass-through rate. Pass-through depends on the curvature of demand, which premerger observables do not reveal, so two markets with the same UPP can experience very different price changes. The merger pass-through rate differs from both observable rates because the opportunity cost enters the merged firm's first-order conditions as a direct shift, whereas a real cost shock would also lower the partner's markup. When diversion is small, the three rates are close, and the paper argues that the premerger rate is then a reasonable approximation.

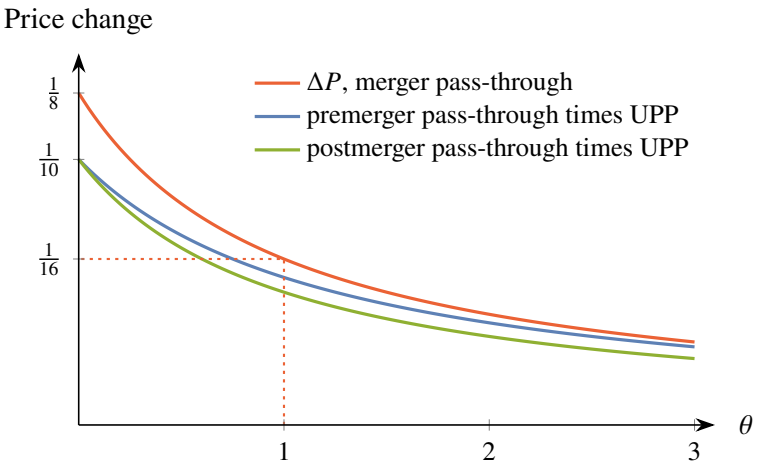


Figure 12.2: Price change caused by the merger (red) and its approximations using the premerger (blue) and postmerger (green) pass-through rates, with the same UPP of $\frac{1}{10}$.

12.6 Fourth Step: Extension to Cournot Conduct

Consider now that, before the merger, the firms compete in quantities, with linear demand, $\theta = 1$. In the paper's framework, firm i then conjectures that its rival raises P_j to hold Q_j fixed, by $\frac{dP_j}{dP_i} = \frac{1}{5}$, so its residual demand is steeper,

$$\frac{dQ_i}{dP_i} = -2 + \frac{2}{5} \cdot \frac{1}{5} = -\frac{48}{25}.$$

The Cournot equilibrium has $P^0 = \frac{89}{88}$ and $Q^0 = \frac{54}{55}$. After the merger to monopoly, the partner no longer accommodates, so the merged firm faces the slope -2. By Proposition 1 of the paper, the pricing pressure is

$$g = D\,(P^0 - c) - \left[\left(\frac{1}{-2}\right) - \left(\frac{1}{-\frac{48}{25}}\right)\right] Q^0 = \frac{9}{88} - \frac{9}{440} = \frac{9}{110}. \tag{12.5}$$

Result 4 (Cournot conduct, Example 2 of the paper). With premerger Cournot conduct and linear demand:

- the pricing pressure is the UPP evaluated at the Cournot markup, $\frac{9}{88}$, minus an end-of-accommodation term, $\frac{9}{440}$, so $g = \frac{9}{110}$;
- the merger raises prices from $\frac{89}{88}$ to $\frac{17}{16}$, by $\frac{9}{176}$, which equals the merger pass-through rate $\frac{5}{8}$ times $\frac{9}{110}$;
- ignoring the end-of-accommodation term would overstate the price change, predicting $\frac{45}{704}$.

Intuitively, a Cournot competitor expects its rival to raise its price when it does, which makes its own demand less elastic; after the merger this accommodation is gone, and the merged firm's demand for each product becomes more elastic, which partly offsets the pricing pressure from

diversion. Since Cournot prices are already higher, the merger to monopoly raises them by less, $\frac{9}{176}$ rather than $\frac{1}{16} = \frac{11}{176}$ under Bertrand.

12.7 Fifth Step: Extension to Cost Efficiencies

Return to Bertrand conduct with linear demand, and let the merger lower the marginal cost of both products by e. A lower cost of product 2 raises its markup, and with it the pricing pressure on product 1, while a lower cost of product 1 lowers the pressure one for one. Hence the pricing pressure becomes $\frac{1}{10} - (1 - D)\, e$, and

$$\Delta P = \frac{5}{8}\left[\frac{1}{10} - \frac{4}{5}\, e\right] = \frac{1}{16} - \frac{e}{2}. \tag{12.6}$$

Result 5 (efficiencies, Section II.C of the paper). With cost reduction e:

- prices do not rise if and only if $e \geq e^\star$, where

  $$e^\star = \frac{g}{1 - D} = \frac{1}{8}$$

  is Werden's (1996) compensating marginal cost reduction;
- comparing efficiencies with the UPP itself, as Farrell and Shapiro (2010) propose, is more permissive: a cost reduction $e = \frac{1}{10}$ passes this test, yet prices rise by $\frac{1}{80}$.

Figure 12.3 depicts the price change as a function of the cost reduction. It falls at rate $\frac{1}{2}$, the merger pass-through rate $\frac{5}{8}$ times $1 - D = \frac{4}{5}$, and reaches zero at $e^\star = \frac{1}{8}$.

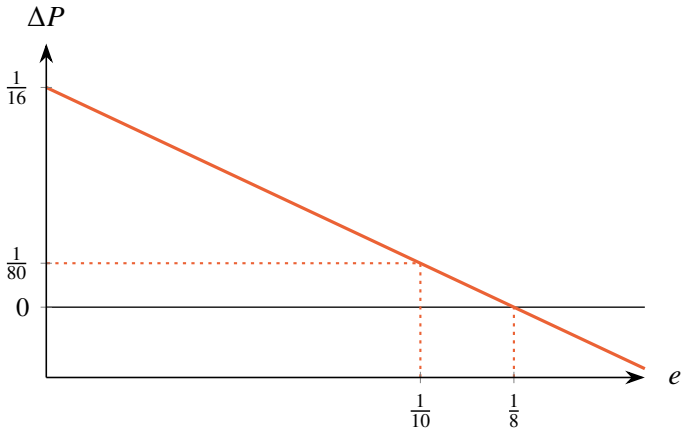


Figure 12.3: Price change caused by the merger as a function of the cost reduction e, with linear demand and Bertrand conduct.

Intuitively, the merged firm passes its cost savings through to prices, but each saving on product 2 also raises the value of the sales that product 1 diverts to it. The efficiency needed to offset the pricing pressure is therefore larger than the pressure itself, by the factor $\frac{1}{1-D} = \frac{5}{4}$.

12.8 Sixth Step: Welfare Changes

To a first order, consumers lose the price change times the quantities they buy, and total surplus falls by the lost quantity times the markups:

$$\Delta CS \approx -\Delta P^T Q, \qquad \Delta SS \approx \left(\frac{\partial Q}{\partial P}\Delta P\right)^T (P - c). \tag{12.7}$$

Result 6 (welfare, Section V of the paper). With linear demand and Bertrand conduct:

- each quantity falls by $\frac{1}{10}$, from 1 to $\frac{9}{10}$;
- consumer surplus falls by approximately

  $$2 \cdot \frac{1}{16} \cdot 1 = \frac{1}{8};$$

  the exact loss is $\frac{19}{160}$, since quantities fall as prices rise;
- total surplus falls by approximately

  $$2 \cdot \frac{1}{10} \cdot \frac{1}{2} = \frac{1}{10}$$

  with premerger markups; the exact loss, $\frac{17}{160}$, uses the average of the premerger and postmerger markups, $\frac{17}{32}$;
- the merged firm's profits rise by $\frac{1}{80}$.

Intuitively, quantities are the natural weights for aggregating price changes: a price increase on a product that sells more hurts consumers more. The approximation errors here, $\frac{1}{160}$ in each case, come from the second-order terms that (12.7) drops; they are small relative to the losses because the price change is small relative to the markups. When every product faces the same pricing pressure, as in this symmetric merger, the consumer surplus formula reduces to

$$\rho \sum_i Q_i \, UPP_i = \frac{5}{8} \cdot 2 \cdot \frac{1}{10} = \frac{1}{8},$$

where ρ is the merger pass-through rate of a common shock, here $\frac{5}{8}$.

12.9 Seventh Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- UPP as the opportunity cost of diverted sales and its generalization to non-Bertrand conduct through the end-of-accommodation term.
- The conversion of pricing pressure into price changes through merger pass-through, and the ranking of premerger, merger, and postmerger pass-through rates.
- Compensating marginal cost reductions and the welfare formulas.

What it drops.

- Multiproduct firms and asymmetric mergers, where the first-order approximation is generally not exact unless demand is linear.
- General conjectural variations and the result that accommodation by rivals can raise or lower the pricing pressure (Example 3 and Appendix C of the paper).
- The bound on the approximation error, and nonconstant marginal costs.

12.10 Connections to Other Chapters

- Chapter 18 links the average markup to the Herfindahl index in a market where consumers face switching costs, and uses this link to predict the price effects of mergers.
- Chapter 1 shows that the curvature of demand also determines the welfare effects of price discrimination.
- Chapter 11 studies how merger control should treat efficiency claims that the merging firms must prove at a cost.
- Chapter 15 adds an innovation stage to a merger between sellers of substitutes. Price coordination, the force that the pricing pressure of this chapter measures, raises the reward from innovating, but the merged firm also takes into account the sales that an improved product takes from its other product.

References

Farrell, J., and C. Shapiro (2010). Antitrust evaluation of horizontal mergers: An economic alternative to market definition. *The B.E. Journal of Theoretical Economics*, 10(1), Article 9.

Jaffe, S., and E. G. Weyl (2013). The first-order approach to merger analysis. *American Economic Journal: Microeconomics*, 5(4), 188–218.

Werden, G. J. (1996). A robust test for consumer welfare enhancing mergers among sellers of differentiated products. *Journal of Industrial Economics*, 44(4), 409–413.

Chapter 13

Reviewing Mergers One at a Time: When a Myopic Policy Is Optimal

Companion note to Nocke, V., and M. D. Whinston (2010), "Dynamic Merger Review," Journal of Political Economy 118(6): 1200–1251.

13.1 Introduction

Antitrust authorities review mergers one at a time. Mergers, however, come in sequences: approving one merger today changes the market in which the next one will be evaluated, and may change which mergers firms propose later. Should an authority that cares about consumers look ahead when it reviews a merger?

Start with a single merger. Merging partners stop competing with each other, which pushes them to cut output, while the merger's efficiency gain, a lower marginal cost, pushes them to produce more. In a Cournot market with linear demand, the merger leaves consumers as well off as before if and only if the efficiency gain equals the partners' margin before the merger. Now add a second merger. A merger that helps consumers raises total output and lowers the price, which lowers the margins of the other firms, so their merger needs a smaller efficiency gain to help consumers. In this sense, mergers that help consumers are complementary: each one makes the others more helpful.

Nocke and Whinston (2010) show that, when each firm has at most one possible merger partner, the authority does not need to look ahead. A myopic policy, which approves a merger if and only if it does not lower consumer surplus in the current market, maximizes discounted consumer surplus, even compared with an authority that knows every future merger. Complementarity is one reason; the other is that a rejected merger can be proposed again and approved later. The result supports the common practice of judging each merger in the current market, but it fails if the authority maximizes aggregate surplus instead.

The toy model has four firms and two possible mergers that share no firm. The first step studies a single merger, the second step shows how two mergers interact, the third step solves a two-period

review problem, the fourth step shows that the argument fails under an aggregate surplus standard, and the fifth step lists what the toy model keeps and drops.

13.2 The Toy Model

Market. Four firms produce a homogeneous good and compete in quantities. Inverse demand is $P = 2 - Q$, every firm has marginal cost 1, and consumer surplus is $CS = \frac{Q^2}{2}$, which rises with total output Q.

Possible mergers. Firms 1 and 2 can merge (merger M_1), and firms 3 and 4 can merge (merger M_2). The mergers are disjoint: no firm can join both. Merger M_k lowers the merged firm's marginal cost from 1 to $1 - \Delta_k$, where $\Delta_k \geq 0$ is its efficiency gain.

Merger review. Firms propose mergers, and an antitrust authority approves or rejects each proposal to maximize discounted consumer surplus. A rejected merger can be proposed again later. A merger is *CS-increasing*, *CS-neutral*, or *CS-decreasing* if, given the current market, it raises, leaves unchanged, or lowers consumer surplus, and *weakly CS-increasing* if it does not lower it.

13.3 First Step: A Single Merger

Before any merger, each firm produces $\frac{1}{5}$, so $Q = \frac{4}{5}$, $P = \frac{6}{5}$, each firm earns $\frac{1}{25}$, and $CS = \frac{8}{25}$. After merger M_1 alone, three firms compete; the merged firm produces $\frac{1+3\Delta_1}{4}$, each of firms 3 and 4 produces $\frac{1-\Delta_1}{4}$, and

$$Q = \frac{3 + \Delta_1}{4}, \qquad P = \frac{5 - \Delta_1}{4}. \tag{13.1}$$

Result 1 (one merger, Lemma 1 and Corollary 1 of the paper). A single merger:

- is CS-neutral if and only if its efficiency gain equals the premerger margin of each merging firm, $\Delta_1 = P - 1 = \frac{1}{5}$; it is CS-increasing for larger gains and CS-decreasing for smaller ones;
- at $\Delta_1 = \frac{1}{5}$, leaves every output unchanged except that the merged firm produces $\frac{2}{5}$, the sum of the partners' outputs, at a margin of $\frac{2}{5}$, the sum of their margins;
- is strictly profitable whenever it is weakly CS-increasing (at $\Delta_1 = \frac{1}{5}$, joint profit rises from $\frac{2}{25}$ to $\frac{4}{25}$), while without an efficiency gain it lowers the partners' joint profit, the merger paradox of Salant et al. (1983); it is profitable if and only if

$$\Delta_1 \geq \frac{4\sqrt{2} - 5}{15} \approx 0.044.$$

Intuitively, a merger lets the partners internalize the effect of their output on each other's price, which makes them cut output; with linear demand this market-power effect equals the partners' margin. Only an efficiency gain of the same size keeps the merged firm producing as much as

the partners did. The required gain falls as the premerger output Q rises, because the margin $P - 1 = 1 - Q$ falls. This is the key property for the rest of the analysis: a merger that raises total output elsewhere makes every other merger easier to approve.

13.4 Second Step: Two Mergers

If both mergers occur, two firms with costs $1 - \Delta_1$ and $1 - \Delta_2$ compete and total output is

$$Q = \frac{2 + \Delta_1 + \Delta_2}{3}.$$

Comparing with (13.1), once M_1 has taken place, merger M_2 is weakly CS-increasing if and only if

$$\Delta_2 \;\geq\; \frac{1 - \Delta_1}{4}, \tag{13.2}$$

the margin of firms 3 and 4 in the market after M_1. This threshold equals $\frac{1}{5}$ when $\Delta_1 = \frac{1}{5}$ and falls below $\frac{1}{5}$ when M_1 is CS-increasing.

Result 2 (complementarity, Propositions 1 and 2 of the paper). With two disjoint mergers:

- a merger that is weakly CS-increasing in isolation remains weakly CS-increasing (and profitable) after another merger that is weakly CS-increasing in isolation;
- a merger that is CS-decreasing in isolation remains CS-decreasing after another merger that is weakly CS-decreasing in isolation;
- a CS-increasing merger can turn a merger that is CS-decreasing in isolation into a weakly CS-increasing one: with $\Delta_2 = \frac{3}{5}$, the threshold for M_1 falls from $\frac{1}{5}$ to $\frac{1}{10}$.

Figure 13.1 depicts the set of mergers that maximizes consumer surplus for each pair of efficiency gains (Δ_1, Δ_2), as in Figure 5 of the paper. Both mergers should occur in the shaded region, which contains the square where both are weakly CS-increasing in isolation ($\Delta_1, \Delta_2 \geq \frac{1}{5}$) and also two slivers where one of them is CS-decreasing in isolation. The dot marks the gains $(\frac{1}{10}, \frac{3}{5})$ used in the third step.

Intuitively, a CS-increasing merger raises total output and lowers the price, so the other merging pair produces less and has a lower margin; its merger then creates less market power and needs a smaller efficiency gain. Mergers that help consumers therefore make other such mergers more helpful, never less. Conversely, a merger that hurts consumers lowers output and raises the margins of others, so it makes other harmful mergers more harmful.

13.5 Third Step: Dynamic Merger Review

Consider two periods, as in Example 1 of the paper. In period 1 only merger M_1 is feasible, with $\Delta_1 = \frac{1}{10}$, so it is CS-decreasing in isolation. In period 2 merger M_2 becomes feasible, with $\Delta_2 = \frac{3}{5}$ or

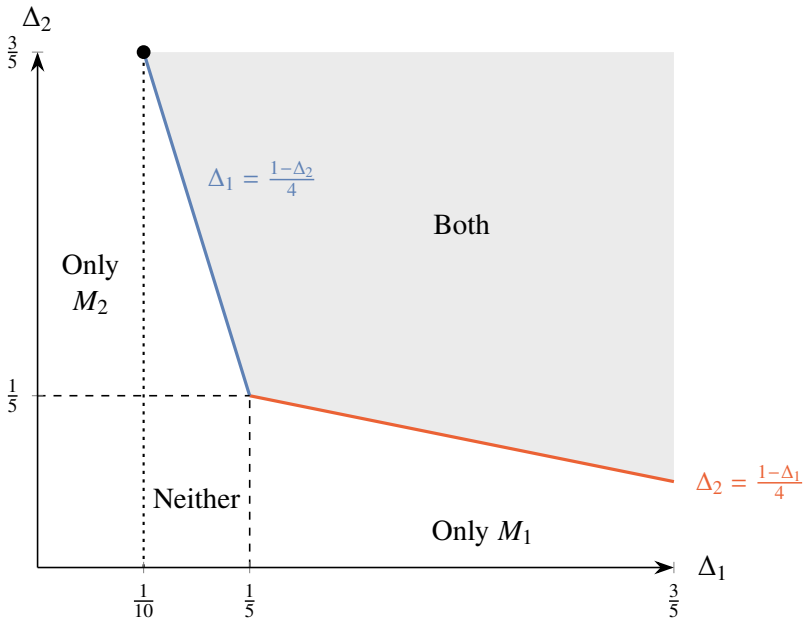


Figure 13.1: Consumer-surplus-maximizing set of mergers for each pair of efficiency gains.

$\Delta_2 = 0$ with equal probabilities. Consumer surplus counts in both periods without discounting. A *myopic* policy approves, in each period, the largest set of proposed mergers that maximizes current consumer surplus.

Under the myopic policy, the authority rejects M_1 in period 1 (consumer surplus would fall from $\frac{8}{25}$ to $\frac{961}{3200}$). In period 2, if $\Delta_2 = \frac{3}{5}$, it approves M_2, which raises output to $\frac{9}{10}$; by (13.2), M_1 is now CS-neutral, so it approves it as well, and both mergers occur. If $\Delta_2 = 0$, it approves neither. Table 13.1 reports consumer surplus under this policy and under early approval of M_1.

	$\Delta_2 = \frac{3}{5}$		$\Delta_2 = 0$	
Policy	Period 1	Period 2	Period 1	Period 2
Myopic: reject M_1 in period 1	$\frac{8}{25}$	$\frac{81}{200}$ (both)	$\frac{8}{25}$	$\frac{8}{25}$ (neither)
Approve M_1 in period 1	$\frac{961}{3200}$	$\frac{81}{200}$ (both)	$\frac{961}{3200}$	$\frac{961}{3200}$ (only M_1)

Table 13.1: Consumer surplus in each period under two policies, for each realization of Δ_2.

Result 3 (optimality of myopic review, Proposition 3 of the paper). In the example:

- the myopic policy yields the highest consumer surplus in every period for each realization of Δ_2, so even an authority that knew Δ_2 in advance could not do better;
- approving M_1 early lowers expected two-period consumer surplus by

$$\frac{3}{2} \cdot \frac{63}{3200} = \frac{189}{6400};$$

- firms propose every feasible merger: in period 1 merger M_1 would raise the partners' joint

profit from $\frac{2}{25}$ to $\frac{169}{1600}$, and in period 2, after M_2, it raises it from $\frac{1}{50}$ to $\frac{1}{25}$; merger M_2 raises its partners' joint profit from $\frac{2}{25}$ to $\frac{49}{100}$.

Intuitively, two features make looking ahead unnecessary. First, by Result 2, approving a weakly CS-increasing merger today is never regretted: later weakly CS-increasing mergers only make it more helpful. Second, rejecting a CS-decreasing merger costs nothing, because the merger remains feasible and can be approved later if other mergers make it helpful, as M_1 does once M_2 occurs. Firms cooperate with this policy because every merger it approves is profitable (Result 1). The example also shows how merger waves arise: one efficient merger can make several others acceptable at the same time.

13.6 Fourth Step: Extension to an Aggregate Surplus Standard

Consider now an authority that maximizes aggregate surplus, consumer surplus plus profits. Let $\Delta_1 = \frac{1}{20}$: in isolation, M_1 lowers consumer surplus but raises aggregate surplus by $\frac{79}{12800}$, because it moves output toward the low-cost merged firm. Once M_2 has taken place with gain Δ_2, the effect of M_1 on aggregate surplus becomes

$$g(\Delta_2) = -\frac{31}{288}\Delta_2^2 + \frac{67}{720}\Delta_2 - \frac{1}{1800}. \tag{13.3}$$

Result 4 (aggregate surplus, Section IV.K of the paper). With $\Delta_1 = \frac{1}{20}$:

- merger M_2 raises aggregate surplus in isolation for every $\Delta_2 \in [\frac{1}{5}, 1]$;
- after M_2, merger M_1 still raises aggregate surplus if

$$\Delta_2 < \frac{67 + 3\sqrt{485}}{155} \approx 0.859,$$

  but lowers it for larger Δ_2; with $\Delta_2 = \frac{9}{10}$ it lowers aggregate surplus by $\frac{23}{5760}$;
- hence mergers that raise aggregate surplus are not complementary, and an authority that approved M_1 in period 1 may regret it once a very efficient M_2 arrives.

Figure 13.2 depicts the effect (13.3) of M_1 on aggregate surplus once M_2 has taken place, for Δ_2 between $\frac{1}{5}$ and 1. It is positive for moderate gains of M_2 and turns negative once Δ_2 exceeds 0.859.

Intuitively, aggregate surplus rewards shifting output to low-cost firms. A small cost cut by M_1 does that when M_1's partners are as efficient as their rivals, but once M_2 has made firms 3 and 4 far more efficient, the merged firm M_1 becomes the high-cost producer, and the output it gains from merging comes partly at the expense of the more efficient rival. The complementarity that makes myopic review optimal is specific to the consumer surplus standard, which tracks only total output.

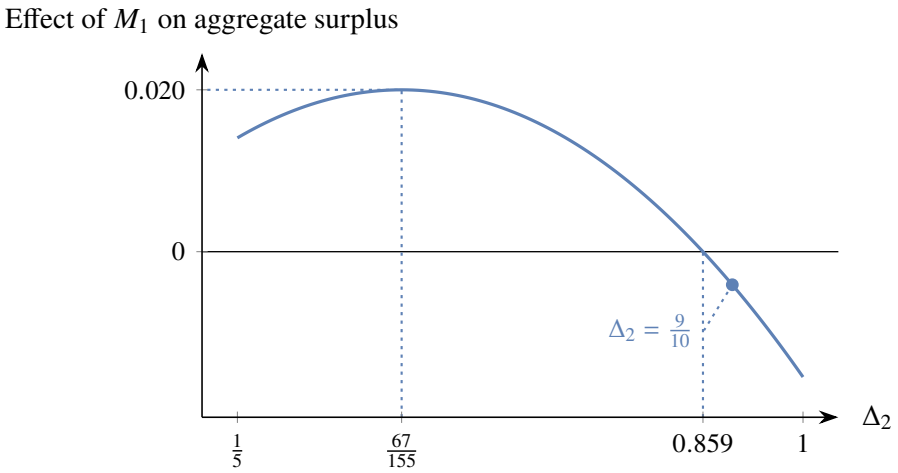


Figure 13.2: Effect of merger M_1 (with $\Delta_1 = \frac{1}{20}$) on aggregate surplus after merger M_2, as a function of M_2's efficiency gain.

13.7 Fifth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The Farrell and Shapiro (1990) test for a single merger: the efficiency gain must offset the partners' margin.
- Complementarity of weakly CS-increasing mergers and the domino effect behind merger waves.
- The optimality of myopic review, firms' incentives to propose every feasible merger, and the failure of complementarity under an aggregate surplus standard.

What it drops.

- General demand, asymmetric firms, and many possible mergers arriving at random dates (Section III of the paper).
- Firms that learn about rivals' mergers only when they are proposed, price competition, differentiated products, demand shifts, entry, and merger costs (Section IV of the paper).
- Mergers that are not disjoint, where a firm can choose among partners; Nocke and Whinston (2013) study that problem.

13.8 Connections to Other Chapters

- Chapter 14 studies the same authority when a target can choose among several partners, and shows that larger mergers must then raise consumer surplus by a strictly positive amount.
- Chapter 11 studies how merger control should treat efficiency claims that the merging firms must prove at a cost.

- Chapter 12 predicts the price effect of a single merger from premerger data, the evaluation that the authority performs at each date in this chapter.

References

Farrell, J., and C. Shapiro (1990). Horizontal mergers: An equilibrium analysis. *American Economic Review*, 80(1), 107–126.

Nocke, V., and M. D. Whinston (2010). Dynamic merger review. *Journal of Political Economy*, 118(6), 1200–1251.

Nocke, V., and M. D. Whinston (2013). Merger policy with merger choice. *American Economic Review*, 103(2), 1006–1033.

Salant, S. W., S. Switzer, and R. J. Reynolds (1983). Losses from horizontal merger: The effects of an exogenous change in industry structure on Cournot-Nash equilibrium. *Quarterly Journal of Economics*, 98(2), 185–199.

Chapter 14

Choosing a Partner: Why Larger Mergers Deserve a Tougher Standard

Companion note to Nocke, V., and M. D. Whinston (2013), "Merger Policy with Merger Choice," American Economic Review 103(2): 1006–1033.

14.1 Introduction

A firm that wants to sell itself can often choose among several buyers. Consider a small firm, the target, that can merge either with a small rival or with a large one. The firms propose only one of these mergers, and the antitrust authority sees the proposed merger, not the alternative that the firms gave up. How strict should the authority be with the merger it sees?

The answer depends on which merger the firms prefer. Firms propose the merger that raises their joint profit the most. Part of a merger's profit gain comes from the partners' combined margin on the output they already sell, and a large acquirer sells more and earns a higher margin. Therefore, for the same effect on consumers, a merger with the large rival is more profitable than one with the small rival. Firms favor the larger merger too much from the consumers' point of view: the target can choose it even when the smaller merger would help consumers more.

Nocke and Whinston (2013) show how the authority should respond. An authority that maximizes expected consumer surplus should approve the smallest merger whenever it does not hurt consumers, which is the plain consumer surplus standard. Larger mergers, however, should raise consumer surplus by a strictly positive amount, and more so the larger the merger. Rejecting a large merger that helps consumers only slightly steers the target toward the smaller merger, which helps them more. This tougher standard is valuable only when an alternative merger exists.

The toy model is a three-firm Cournot market in which a target can merge with a small or a large rival. The first step computes the effect of each merger on consumers and on the merging firms, the second step shows why firms favor the larger merger, the third step derives the optimal approval policy, the fourth step shows how it depends on the alternatives available, and the fifth step lists

what the toy model keeps and drops.

14.2 The Toy Model

Market. Three firms produce a homogeneous good and compete in quantities, with inverse demand $P = 2 - Q$. Firm 0 (the target) and firm 1 (the small acquirer) have marginal cost 1; firm 2 (the large acquirer) has marginal cost $\frac{3}{4}$. Consumer surplus is $\frac{Q^2}{2}$.

Possible mergers. The target can merge with firm 1 (merger M_1) or with firm 2 (merger M_2), but not with both. Each merger is feasible with probability $\frac{1}{2}$, independently. A feasible merger M_k comes with a post-merger marginal cost x, drawn from a uniform distribution on $\left[\frac{1}{2}, 1\right]$ for M_1 and on $\left[\frac{1}{2}, \frac{3}{4}\right]$ for M_2.

Proposal and review. The antitrust authority commits to cost thresholds a_1 and a_2: it approves merger M_k if and only if $x \leq a_k$. Among the feasible mergers that would be approved, the target proposes the one that raises the joint profit of the two merging firms the most, provided this gain is positive (the paper derives this rule from a game in which the target makes take-it-or-leave-it offers). The authority sees only the proposed merger and chooses a_1 and a_2 to maximize expected consumer surplus.

14.3 First Step: Each Merger's Effect on Consumers and on the Merging Firms

Before any merger, firms 0 and 1 produce $\frac{3}{16}$ each and firm 2 produces $\frac{7}{16}$, so $Q = \frac{13}{16}$, $P = \frac{19}{16}$, and consumer surplus is $\frac{169}{512}$; firm 2 has $\frac{7}{13}$ of the market and firm 1 has $\frac{3}{13}$. After merger M_k with post-merger cost x, two firms remain, and total output is

$$Q_1(x) = \frac{1}{3}\left(\frac{13}{4} - x\right), \qquad Q_2(x) = \frac{1}{3}(3 - x). \tag{14.1}$$

Result 1 (consumer surplus and profit effects). For each merger:

- consumer surplus does not change if and only if the merged firm's cost equals $\hat{c}_1 = \frac{13}{16}$ for M_1 and $\hat{c}_2 = \frac{9}{16}$ for M_2, that is, the acquirer's cost minus the target's margin $\frac{3}{16}$; lower costs raise consumer surplus;
- at these CS-neutral costs, the merger raises the merging firms' joint profit by $\frac{9}{128}$ for M_1 and by $\frac{21}{128}$ for M_2;
- the joint profit gain falls with the post-merger cost x for both mergers.

Intuitively, as in Farrell and Shapiro (1990), a merger leaves consumers unaffected only if it lowers cost enough to offset the lost competition; here the merged firm must produce what the two partners

produced together. A merger that consumers are indifferent to is still profitable, because the partners stop competing with each other.

14.4 Second Step: Firms Favor the Larger Merger

Two mergers with the same effect on consumers lead to the same total output Q. Substituting the costs that deliver a given Q into the profit gains, the larger merger's gain exceeds the smaller one's by

$$\frac{32Q - 23}{32} \geq \frac{3}{32} \qquad \text{for every } Q \geq \frac{13}{16}. \tag{14.2}$$

Result 2 (proposal bias). Among mergers that do not hurt consumers, a merger with the large acquirer raises the merging firms' joint profit by more than a merger with the small acquirer that has the same effect on consumer surplus. Hence, when both mergers are approvable, the target can prefer the large merger even when the small one would help consumers more.

Figure 14.1 depicts each merger's effect on the merging firms' joint profit (horizontal axis) and on consumer surplus (vertical axis) as its post-merger cost ranges over its support, as in Figure 1 of the paper. Each curve slopes upward, since a lower cost raises both, and above the horizontal axis the curve for M_2 lies to the right of the curve for M_1.

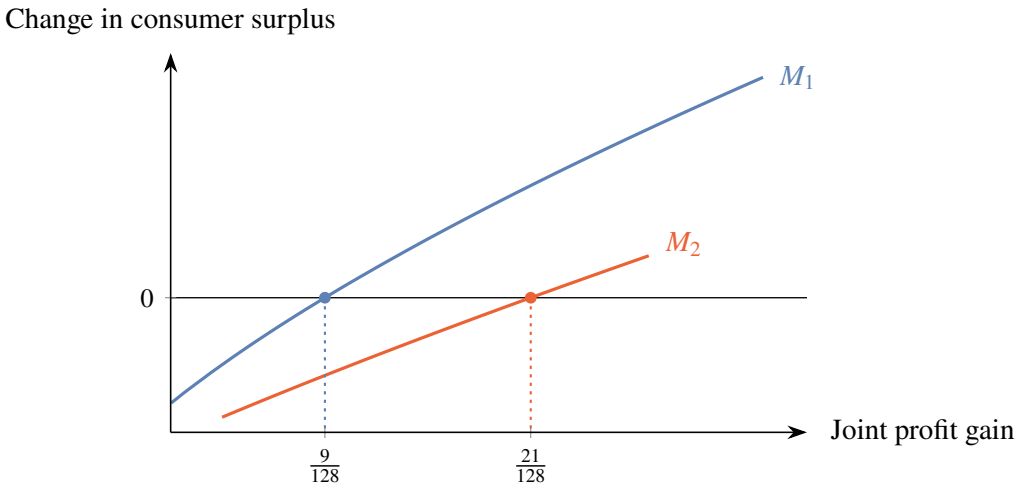


Figure 14.1: Joint profit gain and consumer surplus change of each merger as its post-merger cost varies; the dots mark the CS-neutral costs.

Intuitively, a merger's profit gain comes partly from the partners' combined margin on the output they already sell. The large acquirer sells more and earns a higher margin, so for the same effect on market output it captures a larger profit gain. The firms' ranking of mergers therefore differs systematically from the consumers' ranking, in favor of the larger merger.

14.5 Third Step: The Optimal Approval Policy

A threshold a_k above $\hat{c}_k$ would approve some mergers that hurt consumers, so the authority never sets one. The question is whether it should approve every merger that does not hurt consumers. Lowering a_2 slightly below $\hat{c}_2$ has two effects. When M_2 is the only feasible merger, the authority loses a merger that helps consumers slightly. When M_1 is also feasible and approvable, rejecting the marginal M_2 leads the target to propose M_1 instead, which helps consumers more. Near $\hat{c}_2$ the first loss is close to zero, so the second effect dominates.

Result 3 (tougher standard for the larger merger, the paper's main result). With each merger feasible with probability $\frac{1}{2}$:

- the optimal threshold for the small merger is its CS-neutral cost, $a_1^* = \frac{13}{16}$: the authority approves M_1 if and only if it does not lower consumer surplus;
- the optimal threshold for the large merger is $a_2^* \approx 0.541$, below $\hat{c}_2 = \frac{9}{16} = 0.5625$: the authority approves M_2 only if it raises consumer surplus by at least 0.0058, about 1.7 percent of premerger consumer surplus;
- relative to approving every merger that does not lower consumer surplus, the optimal policy raises the expected gain in consumer surplus from 0.0140 to 0.0141.

These values are computed numerically, by integrating over the cost draws of both mergers. The gain in the last bullet is small in this example because the large merger is distorted only when both mergers are feasible and approvable; the direction of the policy, not its size, is the point.

14.6 Fourth Step: How the Standard Depends on the Alternatives

Let the small merger be feasible with probability r_1, keeping the large merger's probability at $\frac{1}{2}$.

Result 4 (role of the alternative). The optimal threshold for the large merger:

- equals $\hat{c}_2 = \frac{9}{16}$ when $r_1 = 0$, so with no alternative the authority applies the plain consumer surplus standard;
- falls as r_1 rises, to 0.541 at $r_1 = \frac{1}{2}$;
- reaches the bottom of the cost support, $\frac{1}{2}$, when r_1 is close to one, so that the large merger is essentially never approved, as the paper notes can happen for the largest mergers.

Figure 14.2 depicts the optimal threshold a_2^* as a function of r_1.

Intuitively, the tougher standard is a way to steer the target toward the better merger. It is valuable only when an alternative merger exists, and the more likely that alternative, the more the authority gains from rejecting a large merger that helps consumers only slightly.

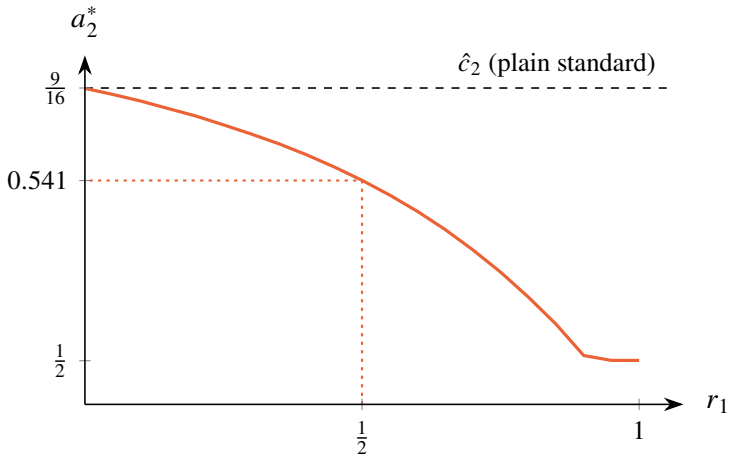


Figure 14.2: Optimal cost threshold for the large merger as a function of the probability that the small merger is feasible.

14.7 Fifth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The Farrell and Shapiro (1990) cost threshold for a merger to leave consumers unaffected.
- Firms' bias toward larger mergers when they choose among mutually exclusive mergers.
- The optimal policy: the plain consumer surplus standard for the smallest merger and a tougher standard for larger mergers, which depends on the available alternatives.

What it drops.

- General demand, more than two acquirers, and the full proof that the minimum acceptable consumer surplus gain increases with the acquirer's size.
- Other bargaining protocols among the firms, including efficient bargaining, and an aggregate surplus standard.
- Mergers that arrive over time, studied in Nocke and Whinston (2010), where disjoint mergers make a myopic policy optimal.

14.8 Connections to Other Chapters

- Chapter 13 studies a sequence of possible mergers in which each firm has at most one partner, and shows that a myopic consumer surplus rule is optimal.
- Chapter 11 studies how merger control should treat efficiency claims that the merging firms must prove at a cost.
- Chapter 15 studies another way in which a merger can harm consumers: through the merging firms' incentive to innovate.

References

Farrell, J., and C. Shapiro (1990). Horizontal mergers: An equilibrium analysis. *American Economic Review*, 80(1), 107–126.

Nocke, V., and M. D. Whinston (2010). Dynamic merger review. *Journal of Political Economy*, 118(6), 1200–1251.

Nocke, V., and M. D. Whinston (2013). Merger policy with merger choice. *American Economic Review*, 103(2), 1006–1033.

Chapter 15

Do Mergers Slow Down Innovation? Price Coordination and the Innovation Externality

Companion note to Federico, G., G. Langus, and T. Valletti (2018), "Horizontal Mergers and Product Innovation," International Journal of Industrial Organization 59: 1–23.

15.1 Introduction

Merger review usually focuses on prices. Competition authorities, however, also worry that a merger between two rival innovators slows down the introduction of new products. Consider two firms that each work on a better version of their product. Will they work less on it once they merge?

Federico et al. (2018) model firms that first invest to improve their products and then compete in prices. A merger changes the merging firms' incentive to innovate through two forces. The first is price coordination: the merged firm no longer competes on price between its two products, which softens competition and tends to raise the reward from improving either of them. The second is the innovation externality: an improved product takes sales away from rival products, and after the merger one of those rivals is the merged firm's own other product. Before the merger, each firm ignored this loss, since it fell on a competitor; after the merger, the firm counts it as a cost of innovating.

In the paper's simulations, the innovation externality dominates. The merging firms innovate less, the rival that stays out of the merger innovates somewhat more, and both total innovation and consumer surplus fall. Consumers therefore lose twice, from higher prices and from fewer new products, and in the toy model only a large cut in research costs makes up for both losses.

The toy model has three products and linear demand, so that prices have closed forms after every outcome of the research stage. The first step solves price competition, the second step solves the innovation stage before the merger, the third step splits the merger's effect on innovation incentives

into the two forces, the fourth step computes the merger's equilibrium effects, the fifth step adds research efficiencies, and the sixth step lists what the toy model keeps and drops.

15.2 The Toy Model

Demand. Three firms each sell one differentiated product at zero marginal cost. Inverse demands are

$$p_i = A_i - q_i - \gamma \sum_{j \neq i} q_j, \tag{15.1}$$

where $\gamma = \frac{1}{2}$ measures how close substitutes the products are, and the quality index A_i equals 1 for the current product and $\frac{3}{2}$ for an improved product. Consumer surplus is

$$CS = \frac{1}{2}\left(\sum_i q_i^2 + \gamma \sum_{i \neq j} q_i q_j\right).$$

Innovation and timing. In stage 1, each firm chooses the probability x_i that it improves its product, at cost $\frac{\kappa}{2}x_i^2$ with $\kappa = \frac{1}{2}$; outcomes are independent across products, and an improved product replaces the old one. In stage 2, firms observe all qualities and set prices. A merger between firms 1 and 2 puts both products, and both innovation decisions, under one owner; firm 3 is the outsider.

15.3 First Step: Price Competition after Innovation

In each of the $2^3 = 8$ quality profiles, prices solve the firms' first-order conditions, a linear system. With no innovation, each firm charges $\frac{1}{4}$, sells $\frac{3}{8}$, and earns $\frac{3}{32}$, and consumer surplus is $\frac{27}{64}$. After the merger, the merged firm charges $\frac{7}{22}$ for each product and the outsider charges $\frac{3}{11}$; consumer surplus falls to $\frac{177}{484}$.

Result 1 (product market, Section 2.1 of the paper). In every quality profile:

- the merger raises the prices of the merging products and lowers consumer surplus;
- an improvement of product 2 lowers the profit earned on product 1, both before and after the merger.

Intuitively, the merged firm stops competing for the customers who switch between products 1 and 2, so it raises both prices. An improved product attracts customers from all rivals, so before the merger it hurts a competitor, and after the merger it also hurts the merged firm's own other product.

15.4 Second Step: Innovation before the Merger

A firm's marginal benefit from raising x_i is its expected profit gain from an improvement, averaged over the rivals' innovation outcomes. Before the merger, firm i chooses x_i to solve

$$E_{k_{-i}}[\pi_i(k_{-i}, 1) - \pi_i(k_{-i}, 0)] = \kappa\, x_i, \tag{15.2}$$

where k_{-i} lists which rival products improved, as in equation (5) of the paper.

Result 2 (pre-merger innovation). The symmetric equilibrium before the merger is unique, with innovation probability $x = \frac{1599}{3604} \approx 0.444$ for each product; the expected profit gain from an improvement is 0.222.

15.5 Third Step: Price Coordination and the Innovation Externality

After the merger, the merged firm's first-order condition for product 1 counts the change in the profit of both of its products. Evaluating the difference between the post- and pre-merger marginal benefits at the pre-merger probabilities gives the paper's downward innovation pressure (DIP), its equation (7):

$$DIP = \underbrace{E[\Delta\tilde{\pi}_1 - \Delta\pi_1]}_{\text{price coordination}} + \underbrace{E[\Delta\tilde{\pi}_2]}_{\text{innovation externality}}, \tag{15.3}$$

where $\Delta\pi_1$ and $\Delta\tilde{\pi}_1$ are the gains in product 1's profit from improving it before and after the merger, and $\Delta\tilde{\pi}_2$ is the change in product 2's profit after the merger when product 1 improves.

Result 3 (downward innovation pressure, Observation 1 of the paper). At the pre-merger probabilities:

- price coordination raises the expected gain from improving product 1 by 0.0221;
- the innovation externality lowers it by 0.0573;
- the DIP is -0.0352, about one sixth of the pre-merger gain 0.222, and the externality outweighs price coordination in each of the four profiles of the other products.

Figure 15.1 depicts the two channels and their sum as functions of the closeness of the products, γ, between 0.05 and 0.6 (all outputs remain positive for $\gamma < 0.617$). Both channels grow with γ, but the externality grows faster, so the downward pressure is stronger for closer substitutes.

Intuitively, price coordination raises the reward from an improvement because the merged firm can steer customers toward the improved product by pricing the other one higher. The innovation externality is a new cost of innovating: part of the improved product's sales comes from the merged firm's own other product, so the merged firm gains less from each improvement than an independent firm would. As in the upward pricing pressure analysis of Farrell and Shapiro (2010), these forces are larger when the merging products are closer substitutes.

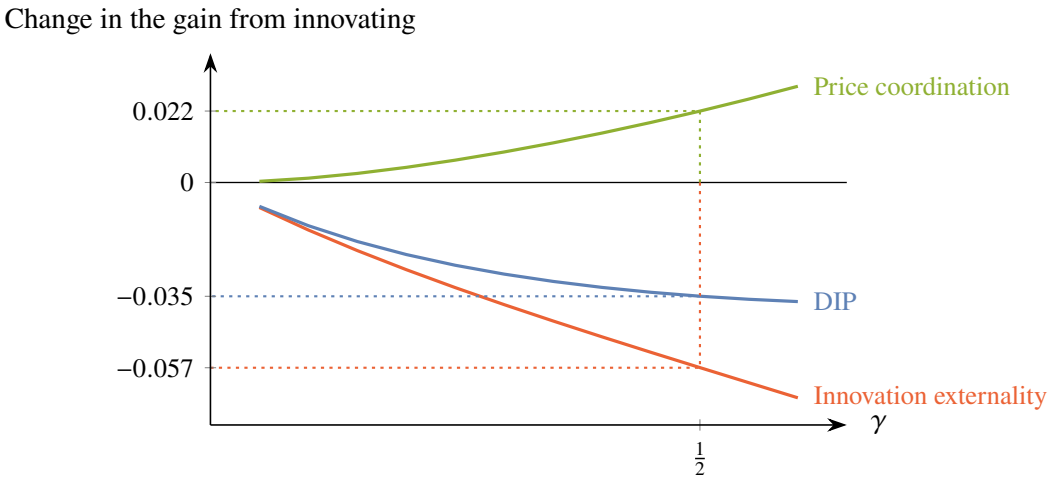


Figure 15.1: The two channels of the merger's effect on the incentive to improve product 1, and their sum, at pre-merger innovation probabilities.

15.6 Fourth Step: Equilibrium Effects of the Merger

After the merger, the merged firm chooses a common probability for its two products and the outsider chooses its own, each solving its first-order condition.

Result 4 (equilibrium, Observations 2 to 5 of the paper). The merger:

- lowers the innovation probability of each merging product from 0.444 to 0.380;
- raises the outsider's innovation probability from 0.444 to 0.480;
- lowers the expected number of improved products from 1.331 to 1.240; the probabilities of at least one, two, and three improvements fall from 0.828, 0.416, and 0.087 to 0.800, 0.371, and 0.069, so the pre-merger distribution first-order stochastically dominates;
- lowers expected consumer surplus from 0.661 to 0.569: higher prices alone, holding innovation at its pre-merger level, would lower it to 0.579, and lower innovation removes another 0.010;
- is profitable: the merging firms' expected profit net of research costs rises from 0.226 to 0.248.

Figure 15.2 depicts the change in the innovation probabilities of the merging products and of the outsider as functions of γ, as in Figure 1 of the paper. The merging products' drop deepens with γ and always exceeds the outsider's rise.

Intuitively, the merged firm innovates less because of the innovation externality. The outsider innovates more because the merging products are now priced higher, which makes an improvement more rewarding for a rival, but this response does not make up for the merging firms' cut. Consumers lose twice: from higher prices for given products and from fewer improved products.

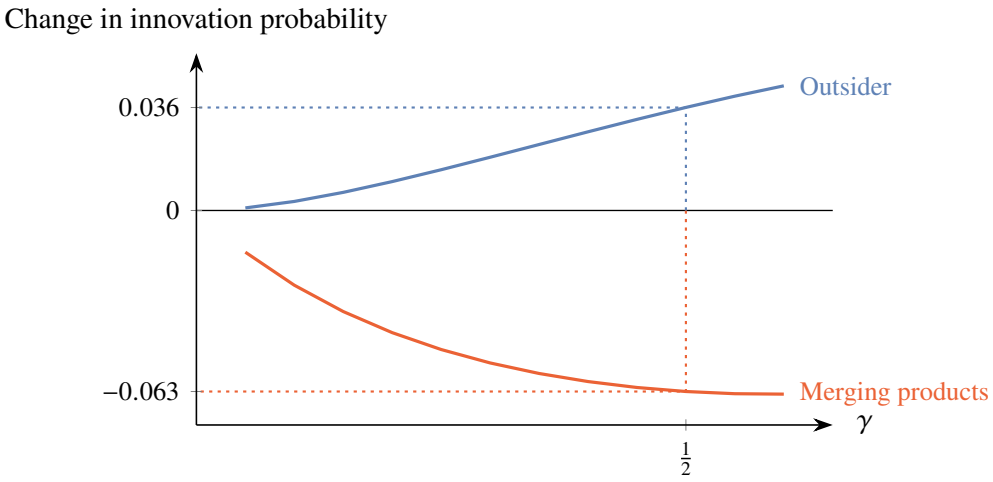


Figure 15.2: Change in equilibrium innovation probabilities caused by the merger.

15.7 Fifth Step: Extension to Research Efficiencies

Consider now that the merger lowers the merged firm's research cost parameter from $\kappa = \frac{1}{2}$ to κ_m, for instance by combining research teams, while the outsider's cost is unchanged.

Result 5 (research efficiencies, Section 4.3 of the paper). With research efficiencies:

- a cost cut of about 16 percent, $\kappa_m \approx 0.419$, restores the merging products' innovation probability to its pre-merger level 0.444;
- restoring expected consumer surplus requires a much larger cut, about 58 percent, $\kappa_m \approx 0.210$, at which the merging products improve with probability 0.776 and the outsider with probability 0.431.

Intuitively, efficiencies in research can offset the innovation externality, but consumers also bear the higher prices that follow the merger in every quality profile. Only a merger that makes improved products much more likely compensates consumers for the loss of price competition.

15.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Uncertain product innovation followed by price competition, with an improved product that replaces the old one.
- The decomposition of the merger's effect on innovation incentives into price coordination and the innovation externality, and the dominance of the latter.
- The equilibrium responses of insiders and outsiders, the fall in total innovation and consumer surplus, and the role of research efficiencies.

What it drops.

- The paper's demand system, in which quality also changes how strongly consumers respond to prices, its four-firm baseline, and its effort technology $x(w) = 1 - e^{-w}$ with linear cost.
- Nested logit and CES demand, other parameter values, and knowledge spillovers (Sections 4.1 to 4.3 of the paper).
- Cost-reducing (process) innovation, which the paper does not study either.

15.9 Connections to Other Chapters

- Chapter 16 studies cost-reducing investment, which affects the merging partner only through prices, so the innovation externality of this chapter disappears and a merger lowers investment through lower output.
- Chapter 17 studies research labs that race to discover a new product, where a merger can also avoid duplicated discoveries.
- Chapter 12 measures the price coordination of this chapter with the upward pricing pressure, and turns it into a predicted price change through pass-through.
- Chapter 39 studies Arrow's replacement effect: a firm that already earns a profit on a product gains less from replacing it. The innovation externality of this chapter is a similar force inside the merged firm.
- In Chapter 44, firms in a research joint venture coordinate their abatement R&D and invest more. In this chapter, the merging firms coordinate their prices and innovate less.
- In Chapter 40, an incumbent buys an innovator and shuts down its project, because the new product would take sales from the incumbent's existing product.

References

Farrell, J., and C. Shapiro (2010). Antitrust evaluation of horizontal mergers: An economic alternative to market definition. *The B.E. Journal of Theoretical Economics*, 10(1), Article 9.

Federico, G., G. Langus, and T. Valletti (2018). Horizontal mergers and product innovation. *International Journal of Industrial Organization*, 59, 1–23.

Chapter 16

Do Mergers Raise Investment? Prices, Cost-Reducing Investment, and Efficiencies

Companion note to Motta, M., and E. Tarantino (2021), "The Effect of Horizontal Mergers, when Firms Compete in Prices and Investments," International Journal of Industrial Organization 78: 102774.

16.1 Introduction

Firms compete in prices, but they also invest. A mobile network operator, for example, builds capacity that lowers its cost of serving each customer. In several mobile telephony mergers in Europe, the industry argued that consolidation would strengthen firms' incentive to invest, as Motta and Tarantino (2021) discuss. Is that right?

The value of a lower marginal cost is proportional to output: a firm that sells more units saves more from each cent of cost reduction. A merger ends the price competition between the two merging firms, which we call the insiders, so the merged firm raises their prices and sells fewer units of each product. With lower output, investment is worth less, and the insiders invest less. The firms that stay out of the merger, the outsiders, gain customers and invest more. Efficiency gains, that is, savings in investment costs from combining the insiders' investment programs, push the other way.

Motta and Tarantino (2021) show that, absent efficiency gains, a merger to monopoly raises prices and lowers investment. With more firms, the insiders' cut in investment outweighs the outsiders' increase in standard demand specifications, so total investment and consumer surplus fall. Only large enough efficiency gains make the merger good for consumers, and the gains that restore investment fall short of those that restore consumer surplus. For merger policy, a promise of more investment is not enough: it must rest on merger-specific efficiencies.

The toy model uses linear demand, a linear marginal cost, and a quadratic investment cost, so every equilibrium has a closed form. The first step shows that investment equals output. The second step studies a merger to monopoly, and the third step adds efficiency gains. The fourth step turns

to three firms, the fifth step adds efficiency gains and a network-sharing agreement, the sixth step studies quality-enhancing investment and connects with Federico et al. (2018), and the seventh step lists what the toy model keeps and drops.

16.2 The Toy Model

Demand. There are $n \in \{2, 3\}$ firms, each selling one differentiated product. Inverse demands are

$$p_i = 3 - q_i - \gamma \sum_{j \neq i} q_j, \qquad \gamma = \frac{1}{4}, \tag{16.1}$$

where γ measures how close substitutes the products are. These demands come from a representative consumer with quasi-linear utility, so consumer surplus is

$$CS = \frac{1}{2}\Big(\sum_i q_i^2 + \gamma \sum_{i \neq j} q_i q_j \Big). \tag{16.2}$$

This is the linear demand system of the paper's parametric examples. Inverting (16.1), the direct demands are

$$q_i = \frac{16}{15}(3 - p_i) - \frac{4}{15}(3 - p_j) \quad (n = 2), \qquad q_i = \frac{10}{9}(3 - p_i) - \frac{2}{9}\sum_{j \neq i}(3 - p_j) \quad (n = 3). \tag{16.3}$$

Costs. A firm that invests $x_i \geq 0$ has marginal cost $c(x_i) = 2 - x_i$ and bears the investment cost $F(x_i) = \frac{x_i^2}{2}$, so its profit is

$$\pi_i = \big(p_i - 2 + x_i\big)\, q_i - \frac{x_i^2}{2}. \tag{16.4}$$

In every equilibrium below, investment is below 2, so marginal cost stays positive.

Timing and the merger. Each firm chooses its price and its investment at the same time, so rivals do not observe investments when they set prices. In a merger, firms 1 and 2 (the insiders) choose both prices and both investments to maximize their joint profit, and firm 3, if present, is the outsider. The merger may bring efficiency gains: the insiders' investment cost becomes

$$(1 - \lambda)\, \frac{x_1^2 + x_2^2}{2}, \qquad 0 \leq \lambda < \frac{1}{3}, \tag{16.5}$$

where λ measures the efficiency gains; $\lambda = 0$ is a merger without them. The upper bound $\frac{1}{3}$ keeps the merged firm's profit strictly concave in its prices and investments. The paper writes the merged firm's investment cost as $F(x_1) + F(x_2) - \lambda G(x_1, x_2)$, and (16.5) is the case $G = \frac{x_1^2 + x_2^2}{2}$.

16.3 First Step: Investment Follows Output

Given its price and its rivals' prices, a firm chooses x_i to maximize $x_i q_i - \frac{x_i^2}{2}$, the cost saving on its q_i units minus the investment cost. The first-order condition is $q_i - x_i = 0$, and the maximized value is $\frac{q_i^2}{2}$.

Result 1 (investment equals output). For an independent firm:

- investment equals output, $x_i = q_i$, so marginal cost is $c_i = 2 - q_i$;
- profit, after the optimal investment, is

$$\Pi_i = (p_i - 2)\, q_i + \frac{q_i^2}{2};$$

- the price satisfies

$$q_i + (p_i - c_i)\frac{\partial q_i}{\partial p_i} = 0, \qquad \text{that is,} \qquad p_i - c_i = \frac{q_i}{\left|\partial q_i / \partial p_i\right|}. \tag{16.6}$$

For the merged firm with efficiency gains, the same steps give $x_i = \frac{q_i}{1-\lambda}$.

Intuitively, a cost cut of one unit saves the firm one unit on each unit it sells, so the marginal return to investment is output. Condition (16.6) is the standard pricing rule with the marginal cost evaluated at the chosen investment: by the envelope theorem, the effect of the price on the optimal investment does not change profit. The paper derives this reduction in general, with any $c(x)$ and $F(x)$ satisfying its assumptions: investment is an increasing function of output, and the game becomes one in prices only. Since the marginal return to investment is output, anything that lowers a firm's output also lowers its incentive to invest.

16.4 Second Step: A Merger to Monopoly

Benchmark. With two independent firms, (16.3) gives $\left|\partial q_i / \partial p_i\right| = \frac{16}{15}$, and at a symmetric price $q = \frac{4}{5}(3 - p)$. Using $c = 2 - q$ in (16.6),

$$p^b = \frac{37}{19}, \qquad x^b = q^b = \frac{16}{19}, \qquad c^b = \frac{22}{19}, \qquad p^b - c^b = \frac{15}{19}. \tag{16.7}$$

Each firm earns $\frac{15}{19} \cdot \frac{16}{19} - \frac{1}{2}\left(\frac{16}{19}\right)^2 = \frac{112}{361}$, and consumer surplus is $CS = \frac{5}{4}q^2 = \frac{320}{361}$. The price is below 2, the cost without investment, because investment lowers marginal cost to $\frac{22}{19}$.

Merger. The merged firm's first-order condition for p_1 adds the effect of p_1 on the profit from product 2,

$$q_1 + (p_1 - c_1)\frac{\partial q_1}{\partial p_1} + (p_2 - c_2)\frac{\partial q_2}{\partial p_1} = 0. \tag{16.8}$$

At the benchmark, the first two terms sum to zero and the last term equals $\frac{15}{19} \cdot \frac{4}{15} = \frac{4}{19} > 0$, so the merged firm wants higher prices. At a symmetric price, (16.8) becomes $q - \left(\frac{16}{15} - \frac{4}{15}\right)(p - c) = 0$, that is, $p - c = \frac{5}{4}\, q$, with $c = 2 - q$ and $q = \frac{4}{5}(3 - p)$.

Result 2 (merger to monopoly). Without efficiency gains, the merger:

- raises the price from $\frac{37}{19} \approx 1.947$ to $p^m = \frac{13}{6} \approx 2.167$;
- lowers investment and output from $\frac{16}{19} \approx 0.842$ to $x^m = q^m = \frac{2}{3}$, so marginal cost rises from $\frac{22}{19}$ to $\frac{4}{3}$;
- lowers consumer surplus from $\frac{320}{361} \approx 0.886$ to $\frac{5}{9} \approx 0.556$;
- raises the profit per product from $\frac{112}{361} \approx 0.310$ to $\frac{5}{6} \cdot \frac{2}{3} - \frac{1}{2}\left(\frac{2}{3}\right)^2 = \frac{1}{3}$.

Intuitively, the merged firm internalizes that a lower price for one product takes customers from the other, so it raises both prices. Higher prices mean lower output, and by Result 1 lower output means lower investment. Lower investment raises marginal cost, which pushes the price up further, so the two effects reinforce each other. The paper proves this result for any demand and cost functions that satisfy its assumptions.

16.5 Third Step: Efficiency Gains in a Merger to Monopoly

With efficiency gains, the merged firm invests $x = \frac{q}{1-\lambda}$, so its marginal cost is $c = 2 - \frac{q}{1-\lambda}$, while its pricing condition is still $p - c = \frac{5}{4}q$. Solving with $q = \frac{4}{5}(3 - p)$,

$$p^m(\lambda) = 2 + \frac{1 - 5\lambda}{2(3 - 5\lambda)}, \qquad x^m(\lambda) = \frac{2}{3 - 5\lambda}, \qquad q^m(\lambda) = \frac{2(1 - \lambda)}{3 - 5\lambda}. \tag{16.9}$$

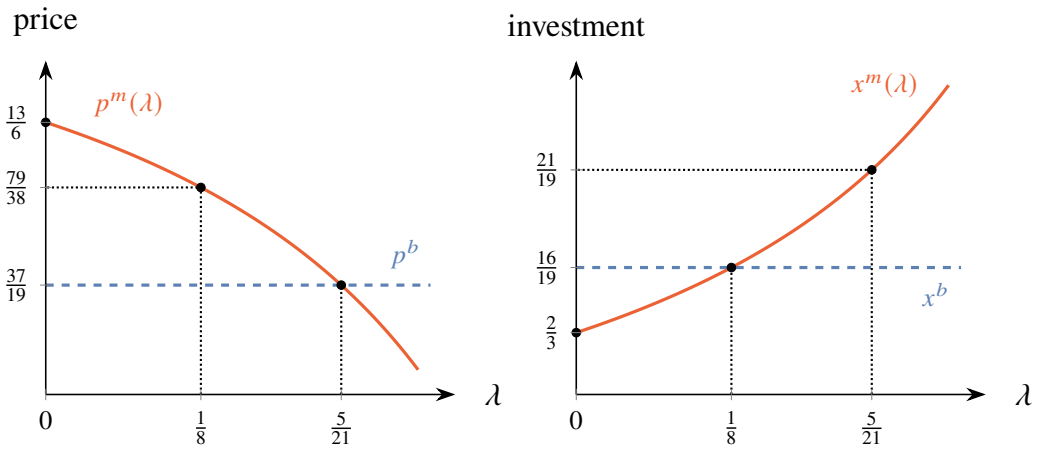
The price is decreasing in λ, with slope $-\frac{5}{(3-5\lambda)^2}$, and investment is increasing, with slope $\frac{10}{(3-5\lambda)^2}$. Since $CS = \frac{5}{4}q^2$ at symmetric prices, consumer surplus moves inversely with the price. Define λ_b as the efficiency gain at which investment returns to its benchmark level, and $\widetilde{\lambda}$ as the gain at which the price returns to its benchmark level:

$$x^m(\lambda_b) = \frac{16}{19} \quad \text{gives} \quad \lambda_b = \frac{1}{8}, \qquad\qquad p^m(\widetilde{\lambda}) = \frac{37}{19} \quad \text{gives} \quad \widetilde{\lambda} = \frac{5}{21}.$$

Result 3 (how large must efficiency gains be). In a merger to monopoly:

- if $\lambda < \frac{1}{8}$, the merger lowers investment and raises the price;
- if $\frac{1}{8} \le \lambda < \frac{5}{21}$, the merger weakly raises investment but still raises the price and lowers consumer surplus; at $\lambda = \frac{1}{8}$ the price is $\frac{79}{38} > \frac{37}{19}$ and output is $\frac{14}{19} < \frac{16}{19}$;
- if $\lambda \ge \frac{5}{21}$, the merger weakly lowers the price and weakly raises consumer surplus; at $\lambda = \frac{5}{21}$ output is back at $\frac{16}{19}$ and investment is $\frac{21}{19}$.

Figure 16.1 depicts the merged firm's price and investment as functions of λ, against their benchmark levels. In panel (a) the price curve crosses the benchmark price only at $\widetilde{\lambda} = \frac{5}{21}$, while in panel (b) the investment curve crosses the benchmark investment earlier, at $\lambda_b = \frac{1}{8}$.



(a) Price (b) Investment

Figure 16.1: Merger to monopoly: the merged firm's price and investment (solid) against the efficiency gain λ, and the benchmark levels with two independent firms (dashed).

Intuitively, a lower investment cost makes the merged firm invest more for each unit it sells, which lowers its marginal cost and its price. Restoring investment, however, does not undo the price effect of the merger. At λ_b the merged firm invests as much as each independent firm did, so its marginal cost equals the benchmark marginal cost, $\frac{22}{19}$, but it still applies the higher markup $p - c = \frac{5}{4}q$ instead of $\frac{15}{16}q$. To bring the price down to $\frac{37}{19}$, investment must exceed its benchmark level by enough to offset the higher markup. Therefore, evidence that a merger will raise investment does not show that it will benefit consumers: for λ between $\frac{1}{8}$ and $\frac{5}{21}$, investment rises and consumers lose.

16.6 Fourth Step: Three Firms, Insiders and Outsiders

With three firms, the outsider reacts to the merger. Its pricing and investment conditions do not change, so it is affected only through the insiders' choices.

Benchmark. By (16.3), $\left|\partial q_i/\partial p_i\right| = \frac{10}{9}$, and at a symmetric price $q = \frac{2}{3}(3-p)$. Condition (16.6) with $c = 2 - q$ gives $p^b = \frac{27}{14}$ and $x^b = q^b = \frac{5}{7}$.

Merger. Firms 1 and 2 merge without efficiency gains. The insiders solve (16.8) for each product, with $x_i = q_i$, and the outsider solves (16.6). With $x_i = q_i$, the pricing conditions form a linear system whose solution is the same for the two insiders. Writing I for an insider and O for the outsider,

$$p_I = \frac{369}{178}, \qquad p_O = \frac{171}{89}, \qquad x_I = q_I = \frac{52}{89}, \qquad x_O = q_O = \frac{70}{89}. \tag{16.10}$$

Table 16.1 reports the marginal costs, margins, and profits, together with total investment $X = \sum_i x_i$ and consumer surplus (16.2), which with an insider pair and an outsider reads $CS = \frac{1}{2}\left[\frac{5}{2}q_I^2 + q_O^2 + q_I q_O\right]$.

Variable	Benchmark, each firm	Merger, each insider	Merger, outsider
Investment $x = q$	$\frac{5}{7}$	$\frac{52}{89}$	$\frac{70}{89}$
Marginal cost $c = 2 - x$	$2 - \frac{5}{7} = \frac{9}{7}$	$2 - \frac{52}{89} = \frac{126}{89}$	$2 - \frac{70}{89} = \frac{108}{89}$
Price p	$\frac{27}{14}$	$\frac{369}{178}$	$\frac{171}{89}$
Margin $p - c$	$\frac{27}{14} - \frac{9}{7} = \frac{9}{14}$	$\frac{369}{178} - \frac{126}{89} = \frac{117}{178}$	$\frac{171}{89} - \frac{108}{89} = \frac{63}{89}$
Profit $(p - c)\, q - \frac{x^2}{2}$	$\frac{9}{14} \cdot \frac{5}{7} - \frac{25}{98} = \frac{10}{49}$	$\frac{117}{178} \cdot \frac{52}{89} - \frac{1352}{7921} = \frac{1690}{7921}$	$\frac{63}{89} \cdot \frac{70}{89} - \frac{2450}{7921} = \frac{1960}{7921}$
Total investment X	$3 \cdot \frac{5}{7} = \frac{15}{7}$	$2 \cdot \frac{52}{89} + \frac{70}{89} = \frac{174}{89}$	
Consumer surplus	$\frac{9}{4}\left(\frac{5}{7}\right)^2 = \frac{225}{196}$	$\frac{1}{2}\left[\frac{5}{2} \cdot \frac{2704}{7921} + \frac{4900}{7921} + \frac{3640}{7921}\right] = \frac{7650}{7921}$	

Table 16.1: Three firms: the benchmark and the merger of firms 1 and 2 without efficiency gains, with $\gamma = \frac{1}{4}$, $c(x) = 2 - x$, and $F(x) = \frac{x^2}{2}$.

Result 4 (merger with an outsider). Without efficiency gains, the merger of firms 1 and 2:

- raises the insiders' price from $\frac{27}{14} \approx 1.929$ to $\frac{369}{178} \approx 2.073$ and lowers their investment from $\frac{5}{7} \approx 0.714$ to $\frac{52}{89} \approx 0.584$;
- raises the outsider's investment to $\frac{70}{89} \approx 0.787$ and lowers its price to $\frac{171}{89} \approx 1.921$;
- lowers total investment from $\frac{15}{7}$ to $\frac{174}{89}$, a fall of $\frac{117}{623} \approx 0.188$, and consumer surplus from $\frac{225}{196} \approx 1.148$ to $\frac{7650}{7921} \approx 0.966$;
- raises the profits of the insiders, from $\frac{20}{49} \approx 0.408$ to $\frac{3380}{7921} \approx 0.427$ in total, and of the outsider, from $\frac{10}{49}$ to $\frac{1960}{7921} \approx 0.247$.

The aggregate. The paper proves the general version of this result with the tools of aggregative games (see Anderson et al., 2020). Define firm i's action as $a_i = 3 - p_i$, so that a higher action is a lower price, and the aggregate as $A = \sum_i a_i$. By (16.3), $q_i = \frac{10}{9}a_i - \frac{2}{9}(A - a_i)$, so each firm's profit depends only on its own action and on A, and total output is $Q = \frac{2}{3}A$. A firm's inclusive best reply $\widetilde{r}(A)$ is the action that is a best reply to the others' actions when the aggregate, including its own action, equals A. In the toy model,

$$\widetilde{r}(A) = \frac{45 - A}{39}, \qquad \widetilde{r}^{\,m}(A) = \frac{A + 36}{42}, \tag{16.11}$$

for an independent firm and for each insider, respectively. An equilibrium aggregate is a fixed point: $3\,\widetilde{r}(A) = A$ before the merger and $2\,\widetilde{r}^{\,m}(A) + \widetilde{r}(A) = A$ after it, which gives $A^b = \frac{45}{14}$ and $A^m = \frac{261}{89}$.

Figure 16.2 depicts the two fixed points. The merger lowers the insiders' inclusive best reply, $\widetilde{r}^{\,m}(A) < \widetilde{r}(A)$ for every $A < 6$, so the sum of the inclusive best replies shifts down and crosses the

45-degree line at a lower aggregate, $A^m = \frac{261}{89} \approx 2.933 < A^b = \frac{45}{14} \approx 3.214$. Since total output and total investment both equal $\frac{2}{3}A$, they fall by the same amount.

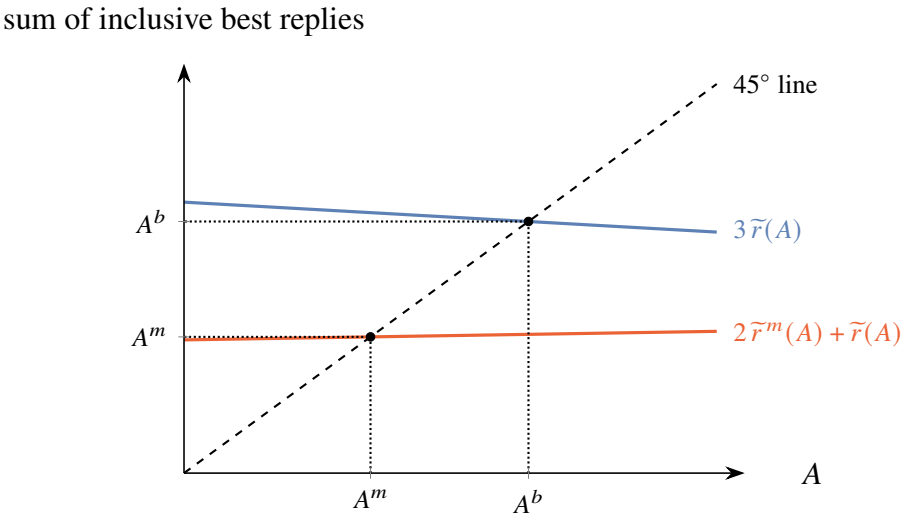


Figure 16.2: The equilibrium aggregate $A = \sum_i (3 - p_i)$ before the merger (blue) and after it (red), as fixed points of the sum of inclusive best replies, with three firms and $\gamma = \frac{1}{4}$.

Intuitively, the insiders raise their prices, so they sell less and invest less. The outsider sees its rivals' prices rise and gains customers, so by Result 1 it invests more. Its own price responds to two forces: the insiders' higher prices pull it up, as in a standard merger, while its lower marginal cost pulls it down. In the toy model the second force wins. The outsider's response, however, is smaller than the insiders' cut, so the aggregate, total output, and total investment fall. With linear demand, consumer surplus does not depend on the aggregate alone, so the paper shows by computation that it falls; for demands with the independence of irrelevant alternatives, such as logit and CES, a lower aggregate is enough.

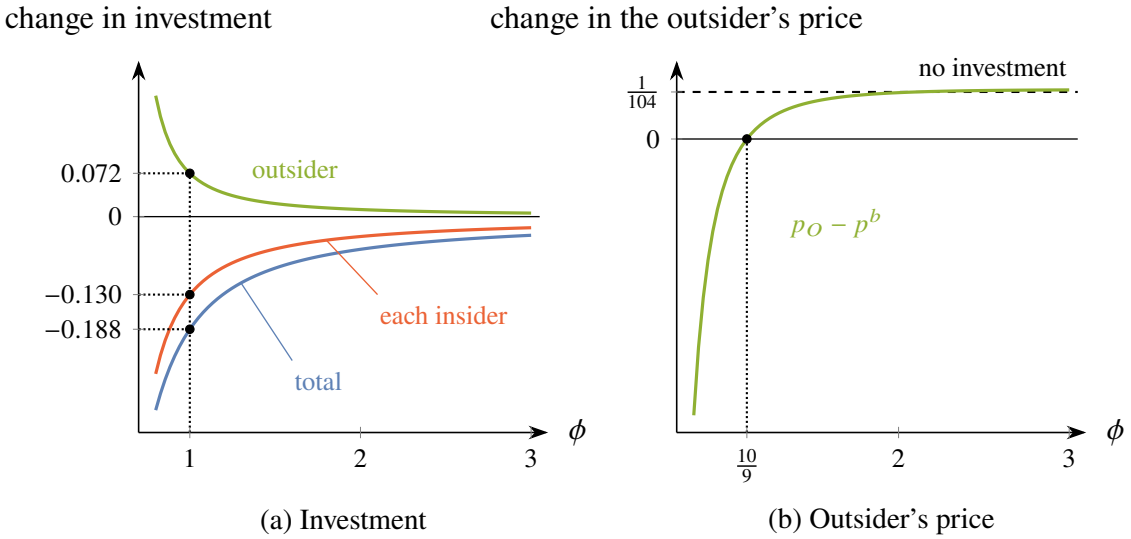
The outsider's price. The direction of the outsider's price depends on how costly investment is. Consider that the investment cost is $F(x) = \phi \frac{x^2}{2}$ with $\phi > 0$, so that $x_i = \frac{q_i}{\phi}$; the base model has $\phi = 1$. Solving the benchmark and the merger again, the change in the outsider's price is

$$p_O - p^b = \frac{9\,\phi^2\,(9\phi - 10)}{2\,(12\phi - 5)\,(351\phi^2 - 342\phi + 80)}. \tag{16.12}$$

Result 5 (the outsider's price). For every $\phi \in [0.8, 3]$, the merger lowers the insiders' investment, raises the outsider's investment, lowers total investment and consumer surplus, and raises the insiders' profits. The outsider's price:

- falls if $\phi < \frac{10}{9}$, which includes the base model, where it falls by $\frac{27}{14} - \frac{171}{89} = \frac{9}{1246}$;
- rises if $\phi > \frac{10}{9}$, and the rise approaches $\frac{1}{104}$ as ϕ grows, the rise in a standard merger without investment, in which the insiders' price rises by $\frac{5}{104}$.

Figure 16.3 depicts the changes in investment, in panel (a), and in the outsider's price, in panel (b), as functions of ϕ. The marked values in panel (a) are the changes in the base model, $\phi = 1$, implied by Table 16.1: $\frac{52}{89} - \frac{5}{7} = -\frac{81}{623}$ per insider, $\frac{70}{89} - \frac{5}{7} = \frac{45}{623}$ for the outsider, and $-\frac{117}{623}$ in total.



(a) Investment (b) Outsider's price

Figure 16.3: Effect of the merger on investment and on the outsider's price against the investment cost parameter ϕ, with three firms, $\gamma = \frac{1}{4}$, and no efficiency gains. The dashed line in panel (b) is the change in the outsider's price when firms cannot invest.

Intuitively, when investment costs little, a given gain in output brings a large investment response and a large cut in marginal cost, which outweighs the upward pull of the insiders' prices. When investment is expensive, the outsider's investment barely responds, and the toy model behaves like a standard merger with price competition, in which all prices rise (Deneckere and Davidson, 1985). For every ϕ in the figure, however, the insiders' cut in investment exceeds the outsider's increase, so total investment falls. The paper finds the same ambiguity across demand systems: in its examples, the outsider's price falls in some specifications, such as linear demand, and rises in others, such as logit.

16.7 Fifth Step: Efficiencies and Network Sharing with Three Firms

The merger with efficiency gains. With the insiders' cost (16.5), each insider invests $x_I = \frac{q_I}{1-\lambda}$. Solving the three pricing conditions again, total investment and consumer surplus are

$$X^m(\lambda) = \frac{6\,(29 - 25\lambda)}{89 - 161\lambda}, \qquad CS^m(\lambda) = \frac{10\,(1853\lambda^2 - 2298\lambda + 765)}{(89 - 161\lambda)^2}. \tag{16.13}$$

A network-sharing agreement. As an alternative to the merger, firms 1 and 2 can sign a network-sharing agreement (NSA), as mobile operators do, or form a research joint venture: they choose

their investments jointly and obtain the same cost savings (16.5), but each sets its own price to maximize its own profit. Each member's price then satisfies (16.6), as in the benchmark, while its investment is $x_i = \frac{q_i}{1-\lambda}$, as in the merger. Solving,

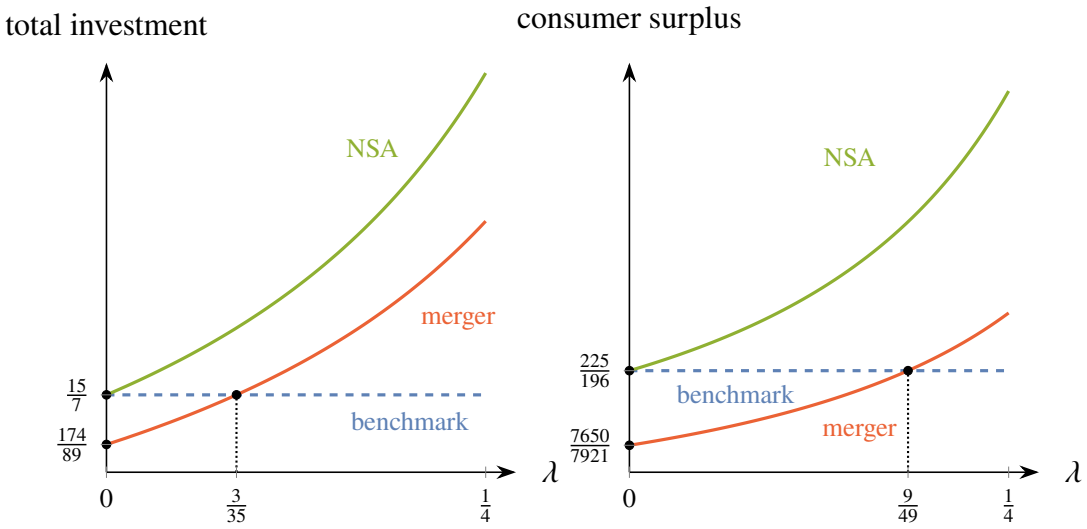
$$X^N(\lambda) = \frac{15\,(13 - 11\lambda)}{91 - 181\lambda}, \qquad CS^N(\lambda) = \frac{25\,(3881\lambda^2 - 4602\lambda + 1521)}{4\,(91 - 181\lambda)^2}. \tag{16.14}$$

At $\lambda = 0$ the NSA coincides with the benchmark, $X^N(0) = \frac{15}{7}$ and $CS^N(0) = \frac{225}{196}$.

Result 6 (efficiency gains and the NSA). With three firms:

- the merger restores total investment at $\lambda = \frac{3}{35}$ and consumer surplus only at $\lambda = \frac{9}{49}$, where both prices return to $\frac{27}{14}$, each insider sells its benchmark output $\frac{5}{7}$, and it invests $\frac{5}{7}$ divided by $1 - \frac{9}{49}$, that is, $\frac{7}{8}$; for $\frac{3}{35} < \lambda < \frac{9}{49}$, the merger raises total investment and lowers consumer surplus;
- for every $\lambda \in (0, \frac{1}{4}]$, the NSA yields higher total investment and higher consumer surplus than both the merger and the benchmark.

Figure 16.4 depicts total investment, in panel (a), and consumer surplus, in panel (b), under the benchmark, the merger, and the NSA, as functions of λ. The merger curves start below the benchmark, at $\frac{174}{89}$ and $\frac{7650}{7921}$, and cross it at $\lambda = \frac{3}{35} \approx 0.086$ and $\lambda = \frac{9}{49} \approx 0.184$, respectively, while the NSA curves start at the benchmark and stay above both.



(a) Total investment (b) Consumer surplus

Figure 16.4: Total investment and consumer surplus against the efficiency gain λ under the benchmark (dashed), the merger, and the network-sharing agreement, with three firms and $\gamma = \frac{1}{4}$.

Intuitively, the merger combines two changes: joint investment, which brings the cost savings, and joint pricing, which raises prices. Only the first is good for consumers. The NSA keeps the first and drops the second, so for every efficiency gain in Figure 16.4 it does better than the merger, and with positive gains it does better than the status quo. The paper shows that the NSA weakly raises the

aggregate relative to the benchmark for any demand in its class, and that it gives consumers more than the merger whenever the merger's efficiency gains are too small to protect them. Therefore, efficiencies that an NSA can also deliver are not specific to the merger and should not count in its favor.

16.8 Sixth Step: Quality-Enhancing Investment

Consider now that investment raises the quality of a product instead of lowering its cost. Marginal cost is constant at 2, the investment cost is still $\frac{x_i^2}{2}$, and a product of quality x_i has inverse demand

$$p_i = 3 + x_i - q_i - \gamma \sum_{j \neq i} q_j, \tag{16.15}$$

so that quality shifts the consumer's willingness to pay, as in Häckner (2000). Define the hedonic price $h_i = p_i - x_i$, the price net of quality. Demand (16.15) in terms of hedonic prices is the toy-model demand (16.1), and profit becomes

$$(p_i - 2)\, q_i - \frac{x_i^2}{2} = \big(h_i - (2 - x_i)\big)\, q_i - \frac{x_i^2}{2},$$

which is profit (16.4) with the hedonic price in place of the price.

Result 7 (quality-enhancing investment). With demand (16.15):

- every equilibrium has the same investments, quantities, profits, and consumer surplus as the cost-reducing toy model, and prices equal the cost-model prices plus investment;
- with three firms, the merger lowers the insiders' quality from $\frac{5}{7}$ to $\frac{52}{89}$, raises the outsider's quality to $\frac{70}{89}$, and raises prices from $\frac{37}{14}$ to $\frac{473}{178}$ for the insiders and to $\frac{241}{89}$ for the outsider;
- the insider's first-order condition for quality at the merger equilibrium reads

$$\underbrace{\frac{10}{9}\Big(\frac{473}{178} - 2\Big)}_{=\frac{65}{89}} \underbrace{- \frac{2}{9}\Big(\frac{473}{178} - 2\Big)}_{=-\frac{13}{89}} = \frac{52}{89} = x_I, \tag{16.16}$$

  while before the merger it reads $\frac{10}{9}\big(\frac{37}{14} - 2\big) = \frac{5}{7} = x^b$.

Intuitively, raising quality by one unit and the price by one unit leaves demand unchanged and raises the margin by one unit, which is what a one-unit cost cut does at an unchanged price. A quality investment therefore works like a cost reduction, and the conclusions of the second to fifth steps carry over. Condition (16.16) shows the two forces at work. The first term is the gain on the improved product itself: the merger raises the insider's margin from $\frac{9}{14}$ to $\frac{117}{178}$, so this term rises from $\frac{5}{7}$ to $\frac{65}{89}$. The second term is new: a better product 1 takes customers from product 2, which the merged firm now owns. This loss, $\frac{13}{89}$, outweighs the gain from the higher margin, so the insiders invest less. Federico et al. (2018), studied in a companion chapter, call these two forces

price coordination and the innovation externality, in a model in which uncertain product innovation comes before price competition; they also find that the externality dominates. In a general quality model the paper notes that the net effect is ambiguous; in the hedonic class, which includes logit and linear demands, the equivalence settles it in favor of lower investment by the insiders.

16.9 Seventh Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Simultaneous choice of prices and cost-reducing investments, with investment increasing in output, so the game reduces to one in prices (the first step).
- The merger to monopoly: higher prices and lower investment without efficiency gains, and the gap between the efficiency gains that restore investment and those that restore consumer surplus (the second and third steps).
- With an outsider: lower investment by the insiders, higher investment by the outsider, a fall in the aggregate, total investment, and consumer surplus, and an outsider price that can move in either direction (the fourth step).
- The thresholds for efficiency gains, and the advantage of a network-sharing agreement over the merger (the fifth step), and the equivalence between quality-enhancing and cost-reducing investment for hedonic demands (the sixth step).

What it drops.

- General demand and cost functions, and the aggregative-game proofs for any number of firms. With linear demand, consumer surplus is not a function of the aggregate alone, so the toy model verifies the fall in consumer surplus by computation, as the paper does for its linear example.
- Logit, CES, and circular-city demands. The paper finds that consumers lose in all of them, and that total surplus can rise with a circular city, where total demand does not respond to prices.
- Asymmetric products, involuntary spillovers between firms' investments, and sequential moves, in which firms observe investments before setting prices. The paper shows that its conclusions extend to asymmetric products, that spillovers work like efficiency gains because the merger internalizes them, and that its examples with sequential moves give the same qualitative results.

16.10 Connections to Other Chapters

- Chapter 15 studies quality-enhancing innovation, where an improved product also takes sales from the merging partner. In this chapter, a cost reduction affects the partner only through prices, and the sixth step splits the merged firm's incentive to improve quality into the same two forces.

- Chapter 11 studies how merger control should treat cost savings that the merging firms must prove. This chapter shows that efficiency gains large enough to restore investment can still leave consumers worse off.
- Chapter 12 measures the upward pricing pressure of a merger. In this chapter, higher prices also lower output and the return to cost-reducing investment, which raises marginal cost and adds to the pressure.

References

Anderson, S. P., N. Erkal, and D. Piccinin (2020). "Aggregative Games and Oligopoly Theory: Short-Run and Long-Run Analysis." *RAND Journal of Economics*, 51(2), pp. 470–495.

Deneckere, R., and C. Davidson (1985). "Incentives to Form Coalitions with Bertrand Competition." *RAND Journal of Economics*, 16(4), pp. 473–486.

Federico, G., G. Langus, and T. Valletti (2018). "Horizontal Mergers and Product Innovation." *International Journal of Industrial Organization*, 59, pp. 1–23.

Häckner, J. (2000). "A Note on Price and Quantity Competition in Differentiated Oligopolies." *Journal of Economic Theory*, 93(2), pp. 233–239.

Motta, M., and E. Tarantino (2021). "The Effect of Horizontal Mergers, when Firms Compete in Prices and Investments." *International Journal of Industrial Organization*, 78, 102774. https://doi.org/10.1016/j.ijindorg.2021.102774

Chapter 17

Merging Research Labs: Less Effort or Less Duplication?

Companion note to Federico, G., G. Langus, and T. Valletti (2017), "A Simple Model of Mergers and Innovation," Economics Letters 157: 136–140, and Denicolò, V., and M. Polo (2018), "Duplicative Research, Mergers and Innovation," Economics Letters 166: 56–59.

17.1 Introduction

Several drug makers can search for the same new treatment. Each runs a lab that may fail, and if two labs succeed, their firms compete for the same patients. What happens to research when two of these firms merge? An earlier chapter studies Federico, Langus, and Valletti (2018), where firms invest to improve products they already sell; here firms race to discover a new product, and a firm whose lab fails earns nothing.

Before the merger, a firm whose lab succeeds ignores the profit it takes from a rival that also succeeds. After the merger, a success in one lab adds nothing when the firm's other lab has also succeeded, since the firm sells one version of the product. The merger therefore internalizes this loss, the innovation externality, and lowers the return to effort in each merging lab. At the same time, two labs that search for the same product duplicate each other, and a merged firm can avoid duplication by putting its effort into one lab.

Federico et al. (2017) assume that the merged firm keeps both labs equally busy. Then the merging labs invest less, the rivals invest more, total research effort falls when the industry has few firms, and consumers lose. Denicolò and Polo (2018) show that equal effort is optimal only if the returns to research fall fast enough; otherwise the merged firm concentrates its effort in one lab, and the merger can raise the probability of discovery. For merger policy, a merger between research rivals is not always bad for innovation.

The toy model uses a quadratic research cost and three firms, so that one firm stays out of the merger. The first step solves the race before the merger, the second step adds the merger with

equal effort, the third and fourth steps compute its effects, the fifth step asks whether equal effort is optimal, the sixth step studies a merger to monopoly, and the seventh step lists what the toy model keeps and drops.

17.2 The Toy Model

Labs. There are n firms, each with one research lab. Lab i discovers a new product with probability $x_i \in [0, 1]$, and research spending is

$$R_i = x_i^2, \tag{17.1}$$

so the probability bought with spending R_i is $x_i = \sqrt{R_i}$. This is the iso-elastic case of Denicolò and Polo (2018) with elasticity $\frac{1}{2}$: research has decreasing returns. Outcomes are independent across labs.

Payoffs. A firm whose lab fails earns nothing. If one firm succeeds, it earns $V > 0$. If two firms succeed, each earns δV, with $0 \leq \delta \leq \frac{1}{2}$; the case $\delta = 0$ is Bertrand competition with a homogeneous product, and $\delta = \frac{1}{2}$ means that the two firms split the monopoly profit. If three or more firms succeed, each earns 0. These payoffs come from demand $q = V(2 - p)$ with zero marginal cost, the example of Federico et al. (2017) scaled by V: one seller charges $p = 1$, two sellers charge $p_2 = 1 - \sqrt{1 - 2\delta}$, and three or more charge $p = 0$. With k successful sellers, consumer surplus $CS_k = \frac{V(2-p)^2}{2}$ is

$$CS_0 = 0, \qquad CS_1 = \frac{V}{2}, \qquad CS_2 = V\big(1 + \sqrt{1 - 2\delta} - \delta\big), \qquad CS_3 = 2V, \tag{17.2}$$

so $CS_1 \leq CS_2 \leq CS_3$, with $CS_1 = CS_2$ at $\delta = \frac{1}{2}$ and $CS_2 = CS_3$ at $\delta = 0$.

Merger and timing. Firms 1 and 2 merge into firm M, which owns both labs; the other firms are outsiders. If both of M's labs succeed, M sells one version of the product, so it counts as one successful firm. In stage 1, firms choose the probabilities of their labs simultaneously; in stage 2, outcomes are observed and payoffs are earned. The first to fourth steps use $n = 3$, $V = 1$, and the baseline $\delta = \frac{1}{2}$, so that firm 3 is the outsider.

17.3 First Step: The Race before the Merger

Consider that all rivals of firm i choose x. Firm i earns V if no rival succeeds and δV if one rival succeeds, so its expected profit is

$$\Pi_i = x_i V\Big[(1 - x)^{n-1} + (n - 1)\,\delta\, x\,(1 - x)^{n-2}\Big] - x_i^2.$$

Profit is strictly concave in x_i, and the first-order condition at a symmetric equilibrium is

$$V(1 - x)^{n-2}\big[1 - x + (n - 1)\,\delta\, x\big] = 2x. \tag{17.3}$$

With $n = 3$ and $V = 1$, the condition reads $(1-x)\left[1-(1-2\delta)x\right] = 2x$. Its left side decreases from 1 to 0 as x rises from 0 to 1, and its right side increases, so there is one solution.

Result 1 (the race before the merger). With $n = 3$ and $V = 1$, the symmetric equilibrium is unique, and each lab succeeds with probability

$$x^* = \frac{2-\delta-\sqrt{3-2\delta+\delta^2}}{1-2\delta} \quad \text{for } \delta < \frac{1}{2}, \qquad x^* = \frac{1}{3} \quad \text{for } \delta = \frac{1}{2}. \tag{17.4}$$

- x^* is increasing in δ, from $2-\sqrt{3} = 0.268$ at $\delta = 0$ to $\frac{1}{3}$ at $\delta = \frac{1}{2}$;
- at $\delta = \frac{1}{2}$, each firm spends $\frac{1}{9}$ and earns expected profit $\frac{1}{9}$.

Intuitively, a lab's success is worth more when the firm keeps a share δV of the market after a rival also succeeds, so a higher δ raises the return to effort. At $\delta = \frac{1}{2}$ the expected value of a success, $(1-x)^2 + 2x(1-x)\cdot\frac{1}{2}$, simplifies to $1-x$, so the first-order condition becomes $1-x = 2x$.

17.4 Second Step: The Merger with Equal Effort in Both Labs

Federico et al. (2017) assume that the merged firm sets the same probability x_M in its two labs. Write x_O for the outsider's probability. Consider lab 1 and the four outcomes of lab 2 and lab 3. Lab 1's success is worth V if both fail, δV if only lab 3 succeeds, and nothing if both succeed, before and after the merger. The difference is the outcome in which only lab 2 succeeds: before the merger, firm 1 earns δV from its own success there; after the merger, firm M already earns V from lab 2, so lab 1's success adds nothing. The marginal benefit of a merging lab therefore falls from

$$V\Big[(1-x_M)(1-x_O) + \delta\, x_M(1-x_O) + \delta\,(1-x_M)\,x_O\Big] \quad \text{to} \quad (1-x_M)\,G,$$

where

$$G \equiv V\big[1-(1-\delta)\,x_O\big] \tag{17.5}$$

is the expected value to M of having at least one success. The merged firm's first-order condition, $(1-x_M)\,G = 2x_M$, gives its best response

$$x_M = \frac{G}{2+G}. \tag{17.6}$$

The outsider's success is worth V if M fails and δV if M succeeds, also when both of M's labs succeed, since M then sells one product. Its best response is

$$x_O = \frac{V}{2}\Big[\delta + (1-\delta)(1-x_M)^2\Big], \tag{17.7}$$

which is decreasing in x_M: research efforts are strategic substitutes. At $\delta = \frac{1}{2}$ and $V = 1$, the best responses of the merging labs are $x_M = \frac{2-x_O}{5}$ before the merger and $x_M = \frac{2-x_O}{6-x_O}$ after it, and those of the outsider are $x_O = \frac{1-x_M}{2}$ before and $x_O = \frac{1+(1-x_M)^2}{4}$ after. Write $\widehat{x}_M$ and $\widehat{x}_O$ for the equilibrium values after the merger.

Result 2 (merging labs invest less, the outsider more). With $n = 3$ and $V = 1$:

- for every $\delta \in (0, \frac{1}{2}]$, the post-merger equilibrium is unique and $\widehat{x}_M < x^* < \widehat{x}_O$;
- at $\delta = 0$, the merger does not change efforts, $\widehat{x}_M = \widehat{x}_O = x^* = 2 - \sqrt{3}$;
- at $\delta = \frac{1}{2}$, $\widehat{x}_M = 0.2887$ and $\widehat{x}_O = 0.3765$, compared with $x^* = \frac{1}{3}$.

Figure 17.1 depicts the best responses at the baseline, with the probability of each merging lab on the horizontal axis and the outsider's probability on the vertical axis. Before the merger, the dashed curves cross at $(\frac{1}{3}, \frac{1}{3})$. The merger shifts the best response of the merging labs to the left, from $\frac{1}{3}$ to $\frac{5}{17} = 0.294$ at $x_O = \frac{1}{3}$, and shifts the outsider's best response up, from $\frac{1}{3}$ to $\frac{13}{36} = 0.361$ at $x_M = \frac{1}{3}$. The solid curves cross at $(\widehat{x}_M, \widehat{x}_O) = (0.2887, 0.3765)$.

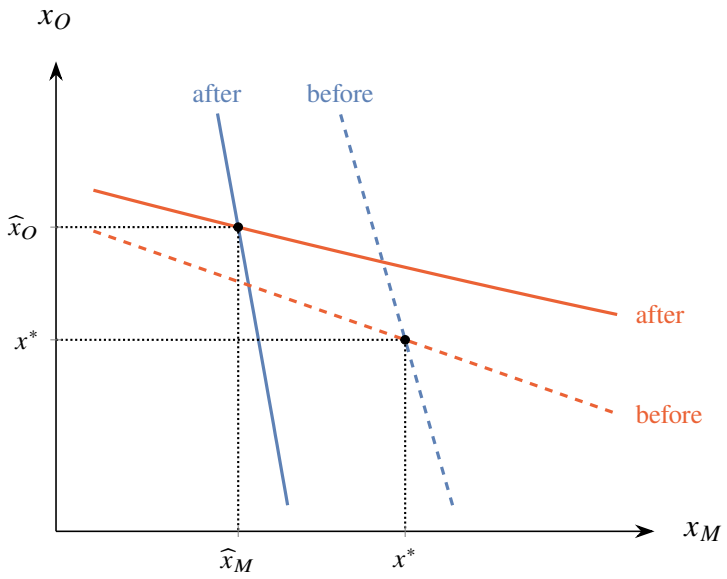


Figure 17.1: Best responses of the merging labs (blue) and of the outsider (red), before the merger (dashed) and after it with equal effort (solid), with $n = 3$, $V = 1$, and $\delta = \frac{1}{2}$.

Intuitively, the merged firm stops counting the profit that lab 1 earns at the expense of lab 2. Evaluated at the pre-merger probabilities, the marginal benefit of a merging lab falls short of its marginal cost by $\delta V x^*(1 - x^*)$, the value δV times the probability that lab 2 succeeds and lab 3 fails; at the baseline this gap is $\frac{1}{2} \cdot \frac{1}{3} \cdot \frac{2}{3} = \frac{1}{9}$, so M cuts effort in each lab. The outsider gains for two reasons. Since efforts are strategic substitutes, it invests more when M invests less. In addition, when all three labs succeed, the outsider now faces one rival instead of two and earns δV instead of zero, which raises its best response by $\frac{\delta V x_M^2}{2}$. At $\delta = 0$ a second success is worth nothing to anyone, so neither change arises.

17.5 Third Step: Total Research Effort and Concentration

Total research spending is $E = \sum_i x_i^2$. At the baseline, the merging labs cut their spending from $2(\frac{1}{3})^2 = 0.2222$ to 0.1667, while the outsider raises its spending from 0.1111 to 0.1417, so total

spending falls from 0.3333 to 0.3084. The probability that no lab succeeds rises from $(\frac{2}{3})^3 = 0.2963$ to $(1-\widehat{x}_M)^2(1-\widehat{x}_O) = 0.3155$. Table 17.1 reports these calculations.

With more firms, more outsiders respond. Federico et al. (2017) compare total effort before and after the merger when δ is close to 0, where efforts do not change. With $V = 1$, write w^* for the pre-merger probability at $\delta = 0$, which solves $(1-w)^{n-1} = 2w$ by (17.3). Using their derivative formulas, with $C'(w) = 2w$ and $C''(w) = 2$ in the toy model,

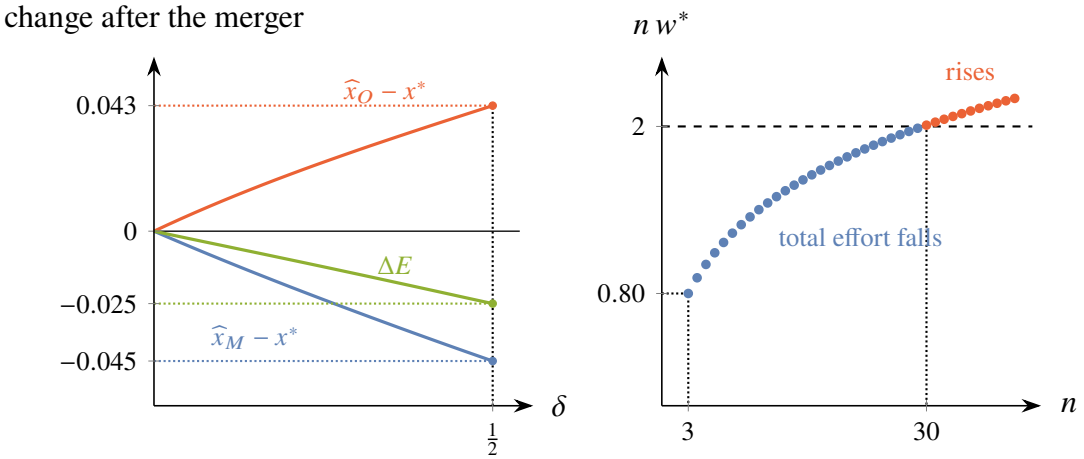
$$\left.\frac{\partial E_{\text{post}}}{\partial \delta}\right|_{\delta=0} - \left.\frac{\partial E_{\text{pre}}}{\partial \delta}\right|_{\delta=0} = \frac{n\,w^* - 2}{1-w^*} \cdot \frac{w^*\,C'(w^*)}{n-1+(1-w^*)^{2-n}\,C''(w^*)}, \tag{17.8}$$

whose second factor is positive. For δ close to 0, total effort therefore falls if and only if $n\,w^* < 2$.

Result 3 (total effort falls in concentrated industries). With $V = 1$:

- with $n = 3$, total spending falls and the probability that no lab succeeds rises for every $\delta \in (0, \frac{1}{2}]$; at $\delta = \frac{1}{2}$, spending falls from 0.3333 to 0.3084 and the probability that no lab succeeds rises from 0.2963 to 0.3155;
- for δ close to 0, total spending falls if $n \le 29$ and rises if $n \ge 30$, since $n\,w^*$ is increasing in n and crosses 2 between $n = 29$ and $n = 30$;
- for δ close to 0, the probability that no lab succeeds moves in the direction opposite to total spending.

Figure 17.2 depicts both parts. Panel (a) shows, for $n = 3$, the change in the probability of each merging lab, $\widehat{x}_M - x^*$, in the outsider's probability, $\widehat{x}_O - x^*$, and in total spending, ΔE, against δ; all three are zero at $\delta = 0$ and grow in size with δ. Panel (b) shows $n\,w^*$ against the number of firms: it equals 0.80 for $n = 3$, 1.987 for $n = 29$, and 2.009 for $n = 30$.



(a) Efforts and spending, $n = 3$ (b) Number of firms, $\delta \to 0$

Figure 17.2: (a) Change in the probability of each merging lab, in the outsider's probability, and in total research spending against δ, with $n = 3$ and $V = 1$; (b) $n\,w^*$ against the number of firms, with $V = 1$: near $\delta = 0$, the merger lowers total effort when $n\,w^* < 2$.

Intuitively, the merger moves effort from the two merging labs to the outsiders. With one outsider, the cut in the merging labs is larger than the outsider's response, so total effort falls. With many outsiders, each responds by a small amount, but the responses add up across all of them. The probability that no lab succeeds is the product of the failure probabilities, so less effort makes this outcome more likely. Total effort is not a welfare measure in this model, since what matters is how effort is spread across the possible outcomes; the fourth step turns to consumers.

Outcome	Before the merger	After the merger
Probability, merging lab	$x^* = \frac{1}{3}$	$\widehat{x}_M = 0.28870$
Probability, outsider	$x^* = \frac{1}{3}$	$\widehat{x}_O = 0.37649$
Spending, merging labs	$2\left(\frac{1}{3}\right)^2 = 0.2222$	$2(0.28870)^2 = 0.1667$
Spending, outsider	$\left(\frac{1}{3}\right)^2 = 0.1111$	$(0.37649)^2 = 0.1417$
Total spending	$0.2222 + 0.1111 = 0.3333$	$0.1667 + 0.1417 = 0.3084$
No lab succeeds	$\left(\frac{2}{3}\right)^3 = 0.2963$	$(0.71130)^2(0.62351) = 0.3155$
Expected CS	$\frac{1}{2} \cdot \frac{18}{27} + 2 \cdot \frac{1}{27} = 0.4074$	$\frac{1}{2}(1 - 0.31546) = 0.3423$
Profit, merging firms	$2 \cdot \frac{1}{9} = 0.2222$	$(0.4941)(0.8118) - 0.1667 = 0.2344$
Profit, outsider	$\frac{1}{3} \cdot \frac{2}{3} - \frac{1}{9} = 0.1111$	$(0.37649)^2 = 0.1417$

Table 17.1: The baseline ($n = 3$, $V = 1$, $\delta = \frac{1}{2}$) before and after the merger with equal effort. The merged firm's profit is $\left[1 - (1 - \widehat{x}_M)^2\right] G - 2\widehat{x}_M^2$ with $G = 1 - \frac{\widehat{x}_O}{2}$, and the outsider's profit equals its spending because its expected revenue is linear in its own probability.

17.6 Fourth Step: Profits and Consumer Surplus

Expected consumer surplus is $\sum_k \Pr(k)\, CS_k$, where k is the number of successful firms and, after the merger, M counts once. Before the merger, k follows a binomial distribution with three trials; after it, $k \le 2$. At the baseline, $CS_1 = CS_2 = \frac{1}{2}$ and $CS_3 = 2$, so expected consumer surplus falls from $\frac{11}{27} = 0.4074$ to 0.3423, by 0.0651. Of this loss, 0.0096 comes from the higher probability that no lab succeeds, $\frac{1}{2}(0.3155 - 0.2963)$, and 0.0556 comes from losing the outcome with three competing sellers: before the merger, it has probability $\frac{1}{27}$ and gives consumers 2; after the merger, every outcome with a success gives consumers $\frac{1}{2}$, so the loss is $\frac{1}{27}\left(2 - \frac{1}{2}\right)$.

As δ approaches 0, efforts do not change, and Federico et al. (2017) show that the change in expected consumer surplus approaches

$$w^{*2}(1 - w^*)^{n-2}\,(CS_1 - CS_2) + (n - 2)\, w^{*3}(1 - w^*)^{n-3}\,(CS_3 - CS_2) < 0. \tag{17.9}$$

With $n = 3$, $CS_2 = CS_3 = 2$ at $\delta = 0$, and $w^* = 2 - \sqrt{3}$, this equals $w^{*2}(1 - w^*)\left(\frac{1}{2} - 2\right) = -0.0788$: when the two merging labs succeed and the outsider fails, consumers face one seller instead of two.

Result 4 (profits and consumers). With $n = 3$ and $V = 1$, for every $\delta \in (0, \frac{1}{2}]$:

- the merger is profitable, $\Pi_M > 2\Pi_i$, with a gain of $x^{*2}(1 - x^*) = 0.0526$ as δ approaches 0 and of $0.2344 - 0.2222 = 0.0121$ at $\delta = \frac{1}{2}$;
- expected consumer surplus falls, by 0.0788 as δ approaches 0 and by 0.0651 at $\delta = \frac{1}{2}$.

Figure 17.3 depicts the merging firms' gain in profit and the change in expected consumer surplus against δ. The profit gain shrinks as δ rises, and the consumer loss stays between 0.0651 and 0.0788.

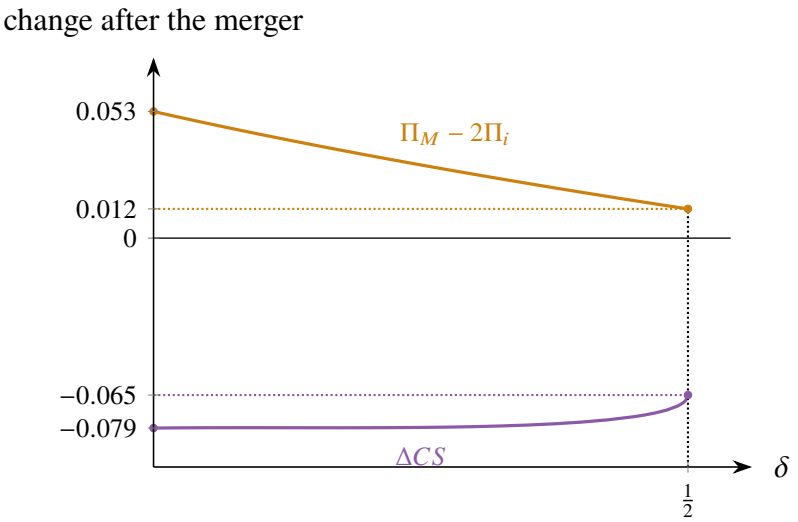


Figure 17.3: The merging firms' gain in expected profit, $\Pi_M - 2\Pi_i$, and the change in expected consumer surplus, ΔCS, against δ, with $n = 3$ and $V = 1$; the values at $\delta = 0$ are the limits as δ approaches 0.

Intuitively, the merger harms consumers through two channels. The first is price coordination: whenever both merging labs succeed, consumers face one seller fewer. The second is the change in efforts, which makes failure of all labs more likely. Near $\delta = 0$ only the first channel operates, and it falls on the outcome in which only the merging labs succeed, where two competing sellers would have charged a price of zero. At $\delta = \frac{1}{2}$, two sellers already charge the monopoly price, so that outcome no longer harms consumers, and the loss comes from the outcome in which all three labs succeed and from the lower probability of discovery. The merging firms gain from the same price coordination, which is why the merger is profitable even though it lowers the return to effort in their labs.

17.7 Fifth Step: Equal Effort or One Lab?

Denicolò and Polo (2018) let the merged firm choose a different probability in each lab. Given the value G of having at least one success, defined in (17.5), M's profit is

$$\Pi_M(x_1, x_2) = \left(x_1 + x_2 - x_1 x_2\right) G - x_1^2 - x_2^2, \tag{17.10}$$

where $x_1 x_2$ is the probability that both labs succeed, the duplication that M does not value. Starting from equal probabilities $\bar{x}$, move to $x_1 = \bar{x} + t$ and $x_2 = \bar{x} - t$. The probability of at least one success

rises by t^2, because duplication falls from $\bar{x}^2$ to $\bar{x}^2 - t^2$, and spending rises by $2t^2$, so

$$\Pi_M(\bar{x}+t, \bar{x}-t) - \Pi_M(\bar{x}, \bar{x}) = (G-2)\, t^2. \tag{17.11}$$

The Hessian of (17.10) has determinant $4 - G^2$, so Π_M is strictly concave if and only if $G < 2$. With equal effort, the best plan is $x_1 = x_2 = \frac{G}{2+G}$ and earns $\frac{G^2}{2+G}$. With one lab, the best plan is $x_1 = \frac{G}{2}$ when $G \leq 2$, earning $\frac{G^2}{4}$, and $x_1 = 1$ when $G \geq 2$, earning $G - 1$.

Result 5 (equal or concentrated effort). Given G:

- if $G < 2$, the merged firm's unique optimal plan is equal effort, $x_1 = x_2 = \frac{G}{2+G} < \frac{1}{2}$, with profit $\frac{G^2}{2+G}$;
- if $G > 2$, the optimal plan uses one lab, $x_1 = 1$ and $x_2 = 0$ (or the reverse), with profit $G - 1$, and the product is discovered for sure;
- at $G = 2$, both plans earn 1.

Equivalently, equal effort is optimal if and only if each lab's equal-effort probability is below $\frac{1}{2}$, that is, if and only if the failure probability $1 - \sqrt{R}$ is log-convex in spending at the equal-effort level $R = x^2$, the condition of Denicolò and Polo (2018). With spending $R = x^k$ for $k > 1$, the same local condition reads $x < 1 - \frac{1}{k}$.

Figure 17.4 depicts the profit of the best equal-effort plan and of the best one-lab plan against G. The two curves cross at $G = 2$, where both earn 1; equal effort earns more to the left and one lab earns more to the right.

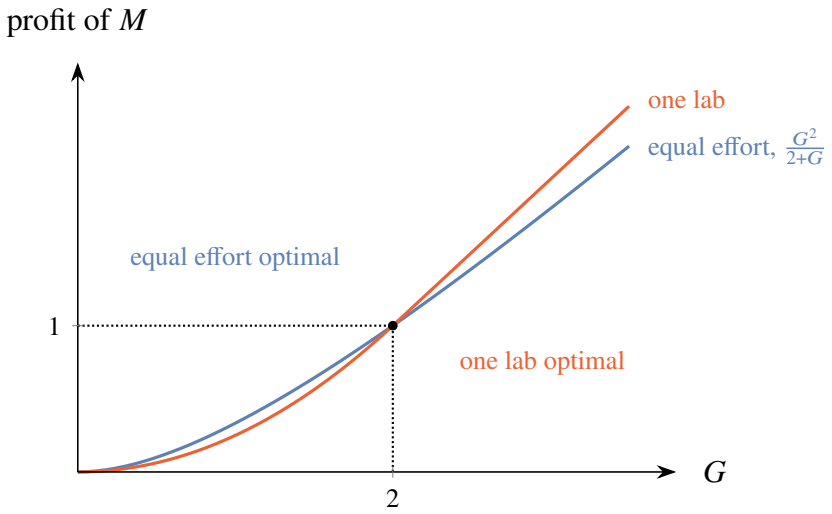


Figure 17.4: Profit of the merged firm with equal effort in both labs (blue) and with one lab (red), $\frac{G^2}{4}$ for $G \leq 2$ and $G - 1$ for $G \geq 2$, against the value G of having at least one success.

Intuitively, the two labs search for the same product, so the second success is wasted. Spreading effort evenly keeps spending low, because the cost of a probability rises fast, while concentrating effort avoids duplication. By (17.11), the first force dominates when the value of success is small

relative to the curvature of the cost, $G < 2$, and the second dominates when it is large. With spending $R = x^k$, a k close to 1 means that the returns to research fall slowly, and the threshold $1 - \frac{1}{k}$ is close to 0, so concentrated effort becomes attractive even at low probabilities. At the baseline of the second to fourth steps, $G = 1 - \frac{\widehat{x}_O}{2} = 0.812 < 2$, and the merged firm's profit $\frac{G^2}{2+G}$ equals the 0.2344 of Table 17.1: equal effort is optimal there, and Results 2 to 4 stand.

17.8 Sixth Step: A Merger to Monopoly

Denicolò and Polo (2018) study the case most likely to harm innovation: the only two firms in the race merge, and if both succeed before the merger, each earns $\frac{V}{2}$. In the toy model, this is $n = 2$ and $\delta = \frac{1}{2}$, and the value of the discovery V is free. Before the merger, condition (17.3) becomes $V\left(1 - \frac{x}{2}\right) = 2x$, so

$$x^* = \frac{2V}{4+V} \quad \text{for } V \le 4, \qquad R^* = x^{*2} = \frac{4V^2}{(4+V)^2}, \tag{17.12}$$

and $x^* = 1$ for $V \ge 4$. After the merger there is no rival, so $G = V$. By Result 5, the merged firm chooses equal effort, $x = \frac{V}{2+V}$ and $R = \frac{V^2}{(2+V)^2}$, if $V < 2$, and one lab with $x_1 = 1$ if $V > 2$, which are the closed forms that Denicolò and Polo (2018) report for this case. The probability that the product is discovered is

$$X_{\text{pre}} = 1 - \left(\frac{4-V}{4+V}\right)^2, \qquad X_{\text{post}} = 1 - \left(\frac{2}{2+V}\right)^2 \text{ if } V < 2, \qquad X_{\text{post}} = 1 \text{ if } V > 2, \tag{17.13}$$

with X_{pre} for $V \le 4$. At $\delta = \frac{1}{2}$, $CS_1 = CS_2 = \frac{V}{2}$, so expected consumer surplus is $\frac{V}{2}X$ before and after the merger.

Result 6 (merger to monopoly). With $n = 2$ and $\delta = \frac{1}{2}$:

- if $V < 2$, the merger lowers each lab's probability, from $\frac{2V}{4+V}$ to $\frac{V}{2+V}$, and lowers the probability of discovery and expected consumer surplus;
- if $2 < V < 4$, the merged firm runs one lab with $x_1 = 1$, and the merger raises the probability of discovery from $X_{\text{pre}} < 1$ to 1 and raises expected consumer surplus;
- the cutoff is $V = 2$, where the merged firm is indifferent between the two plans.

Figure 17.5 depicts the probability of discovery against V. Before the merger it rises smoothly and reaches 1 at $V = 4$. After the merger it is lower for $V < 2$ and jumps to 1 at $V = 2$, from $\frac{3}{4}$ with equal effort, while the pre-merger probability there is $\frac{8}{9} = 0.89$. The dashed curve shows the equal-effort plan for $V > 2$, which the merged firm no longer chooses. Table 17.2 reports three values of V.

Intuitively, the merger removes the business-stealing motive for research, which lowers effort, and it lets the merged firm avoid duplication, which raises the probability of discovery per unit spent. For a small discovery, the first effect dominates and the merged firm runs two slower labs. For a large discovery, duplication is the larger waste, and the merged firm runs one lab at full speed. The single lab then spends 1, while the two rivals together spend $\frac{8V^2}{(4+V)^2}$, which exceeds 1 for

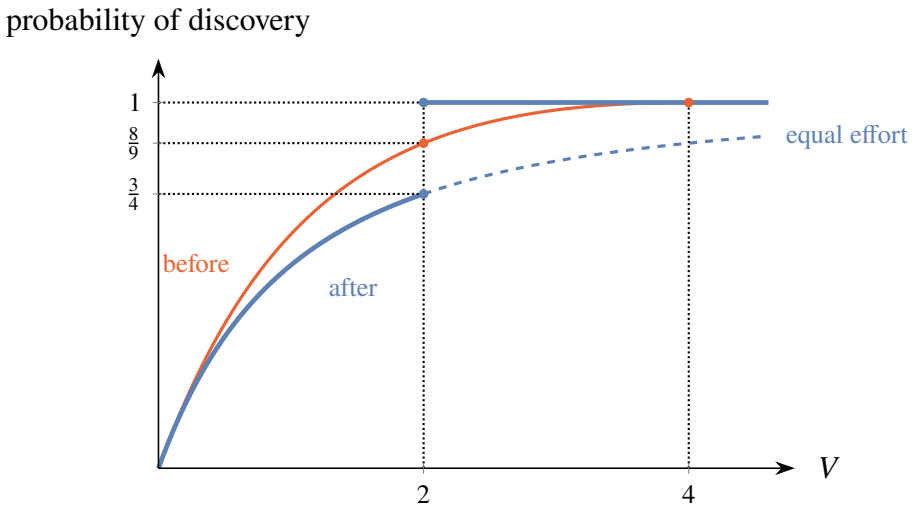


Figure 17.5: Probability that the product is discovered before the merger (red) and after it (blue), against the value of the discovery V, with $n = 2$ and $\delta = \frac{1}{2}$; the dashed curve is the equal-effort plan for $V > 2$.

Value of the discovery	$V = 1$	$V = \frac{3}{2}$	$V = \frac{5}{2}$
Before: $x^* = \frac{2V}{4+V}$	$\frac{2}{5}$	$\frac{6}{11}$	$\frac{10}{13}$
Before: $X_{\text{pre}} = 1 - (1 - x^*)^2$	$1 - \frac{9}{25} = 0.6400$	$1 - \frac{25}{121} = 0.7934$	$1 - \frac{9}{169} = 0.9467$
Equal effort: profit $\frac{V^2}{2+V}$	$\frac{1}{3} = 0.3333$	$\frac{9}{14} = 0.6429$	$\frac{25}{18} = 1.3889$
One lab: profit $\frac{V^2}{4}$ or $V - 1$	$\frac{1}{4} = 0.2500$	$\frac{9}{16} = 0.5625$	$\frac{5}{2} - 1 = 1.5000$
After: plan	$x_1 = x_2 = \frac{1}{3}$	$x_1 = x_2 = \frac{3}{7}$	$x_1 = 1, x_2 = 0$
After: X_{post}	$1 - \frac{4}{9} = 0.5556$	$1 - \frac{16}{49} = 0.6735$	1
$\Delta X = X_{\text{post}} - X_{\text{pre}}$	−0.0844	−0.1199	+0.0533
$\Delta CS = \frac{V}{2}\,\Delta X$	−0.0422	−0.0899	+0.0666

Table 17.2: A merger to monopoly with $n = 2$ and $\delta = \frac{1}{2}$: the merged firm compares the two plans, and the merger lowers the probability of discovery for $V < 2$ and raises it for $V > 2$.

$V > \frac{4}{2\sqrt{2}-1} = 2.188$: at $V = \frac{5}{2}$, for instance, the rivals spend $\frac{200}{169} = 1.18$ and still fail with probability 0.0533, while the merged firm spends less and never fails.

17.9 Seventh Step: What the Toy Model Keeps and Drops

The quadratic research cost, with three firms in the first to fourth steps and two firms in the sixth step, keeps the main conclusions of both papers:

1. With equal effort in both labs, the merging labs invest less, because the merged firm internalizes the innovation externality, and the outsiders invest more (the second step).
2. Total research effort falls after the merger in concentrated industries and can rise in fragmented ones, and the probability that no lab succeeds moves in the opposite direction (the third step).
3. The merger is profitable and lowers expected consumer surplus, through price coordination and lower effort (the fourth step).
4. Equal effort is optimal for the merged firm only if the failure probability is log-convex at the equal-effort spending, which in the toy model means a probability below $\frac{1}{2}$ in each lab; otherwise the merged firm concentrates its effort in one lab (the fifth step).
5. In a merger to monopoly, the merger lowers innovation for small discoveries and raises it for large ones, with the cutoff $V = 2$ (the sixth step).

Federico et al. (2017) prove their results for δ close to 0 and illustrate some of them numerically; the toy model shows Results 2 to 4 for every $\delta \in (0, \frac{1}{2}]$ with three firms.

The toy model drops several features. The research cost has a finite marginal cost at $x = 1$, unlike the cost in Federico et al. (2017), so a firm can buy certainty; this is why the one-lab plan of the fifth step sets $x_1 = 1$, as in the iso-elastic example of Denicolò and Polo (2018). With a quadratic cost, the merged firm never chooses two labs with different positive probabilities, a case that their example with elasticity $\frac{2}{3}$ produces for intermediate values of V. The toy model compares the two plans in a merger to monopoly and at the baseline of the second step, where equal effort is optimal; it does not study whether a merged firm that faces outsiders would choose one lab when the discovery is very valuable. Payoffs are given in reduced form, with one homogeneous product, so the toy model leaves aside differentiated products, research spillovers, efficiencies, and correlated research projects.

17.10 Connections to Other Chapters

- Chapter 15 studies firms that improve products they already sell. In this chapter, firms race to discover a new product, and a merger also lets the merged firm avoid duplicated research.
- Chapter 40 studies an incumbent that buys an entrant with an overlapping project and may shut it down. In this chapter, a merger can instead concentrate research in one lab and raise the probability of a discovery.

- Chapter 39 also models research as a race in which only success pays, so the value of research depends on what rivals achieve.

References

Denicolò, V., and M. Polo (2018). "Duplicative Research, Mergers and Innovation." *Economics Letters*, 166, pp. 56–59. https://doi.org/10.1016/j.econlet.2018.02.021

Federico, G., G. Langus, and T. Valletti (2017). "A Simple Model of Mergers and Innovation." *Economics Letters*, 157, pp. 136–140. https://doi.org/10.1016/j.econlet.2017.06.014

Federico, G., G. Langus, and T. Valletti (2018). "Horizontal Mergers and Product Innovation." *International Journal of Industrial Organization*, 59, pp. 1–23. https://doi.org/10.1016/j.ijindorg.2018.03.001

Chapter 18

Switching Costs and Market Power: Concentration and Mergers in Customer Markets

Companion note to Somaini, P., and L. Einav (2013), "A Model of Market Power in Customer Markets," Journal of Industrial Economics 61(4): 938–986.

18.1 Introduction

Many people keep the same bank for years. They stay with their phone carrier and their insurer too, even when a rival offers a lower price, because moving an account takes time and effort. Markets in which buyers are partly locked in to the seller they used before are called customer markets, and the cost of moving is a switching cost.

Switching costs give a firm two opposing pricing incentives. A firm with many locked-in customers wants to charge them a high price, which the literature calls "milking" its customer base. At the same time, every new customer it wins today is locked in tomorrow, so the firm also wants to charge a low price to build its base, an "investment" incentive. Forward-looking new customers add a third force: they expect that a firm that is low-priced today will raise its price once they are locked in, so they respond less to its price cuts, an "anticipation" effect.

Somaini and Einav (2013) build an oligopoly model of customer markets with any number of firms in which the dynamic equilibrium is unique and has a closed form. In equilibrium, the average markup is proportional to the Herfindahl index, as in the static Cournot model. The model can therefore be used to predict how a merger changes prices and welfare, and the paper finds that a merger of two low-cost firms raises prices the most, while a merger of two high-cost firms can raise welfare.

The toy model has four firms, two cost levels, and the paper's baseline values for the dynamic parameters. The first step solves the static market. The second step derives demand with switching

costs, the third step finds the equilibrium and shows it is unique, the fourth step separates the three forces, the fifth step links markups to concentration, the sixth step studies mergers, and the seventh step lists what the toy model keeps and drops.

18.2 The Toy Model

Products and consumers. There are $J = 4$ products, sold by four single-product firms, at the vertices of a regular tetrahedron. Every pair of products is joined by an edge of length 1, so there are $\frac{J(J-1)}{2} = 6$ edges. Consumers are spread uniformly along the edges, and a consumer on the edge between products i and j only considers these two products. She buys one unit, pays a transport cost equal to her distance to the product she buys, and values the good enough that she always buys. Firm i has a constant marginal cost c_i, and $\overline{c}$ is the average cost. We use two cost configurations: equal costs, $c_i = 0$, and two low-cost firms with $c_L = 0$ (firms 1 and 2) and two high-cost firms with $c_H = \frac{1}{2}$ (firms 3 and 4), so that $\overline{c} = \frac{1}{4}$.

Cohorts and switching. Each consumer lives for two periods. In every period a young cohort arrives, of size g times the old cohort; the toy model sets $g = 1$. An old consumer stays on the same edge but draws a new location on it, and she pays a switching cost s if she buys a product different from the one she bought when young. Consumers discount the next period by δ_c. Firms discount by $\delta_f = r_f\, g$, where r_f is their discount factor and g reflects the growth of the market. The toy model uses the paper's baseline values,

$$s = 0.4, \qquad \delta_c = 0.7, \qquad \delta_f = 0.5, \qquad g = 1. \tag{18.1}$$

Since the transport cost of crossing a whole edge is 1, $s = 0.4$ is the cost of travelling 40 percent of an edge.

Prices and normalization. In every period each firm sets one price p_i for all consumers: it cannot charge young and old consumers different prices, or its own past customers and its rivals' past customers different prices. All demands are measured per old consumer. The average price of firm i's rivals is $\overline{p}_{-i} = \frac{1}{J-1}\sum_{j\neq i} p_j$, and $\overline{p}$ is the average price of all firms.

18.3 First Step: The Static Market

Consider first a single period without switching costs ($s = 0$). On the edge between i and j, a consumer at distance θ from i buys i if $\theta + p_i \le 1 - \theta + p_j$, so firm i sells to the fraction $\frac{1}{2}(1 - p_i + p_j)$ of the edge. Adding over its $J - 1$ edges and dividing by the $\frac{J(J-1)}{2}$ edges, firm i's market share is

$$x_i = \frac{1}{J}\big(1 - p_i + \overline{p}_{-i}\big), \tag{18.2}$$

so the rivals' prices enter only through their average $\overline{p}_{-i}$. Firm i maximizes $(p_i - c_i)\, x_i$, and the first-order condition $x_i - \frac{1}{J}(p_i - c_i) = 0$ gives

$$p_i = \frac{1}{2}\big(1 + c_i + \overline{p}_{-i}\big) \qquad \text{and} \qquad p_i - c_i = J\, x_i. \tag{18.3}$$

The Herfindahl index is $HHI = \sum_i x_i^2$.

Result 1 (static equilibrium; equations (5) and (6) of the paper). The static equilibrium prices are

$$p_i = \bar{c} + 1 + \frac{J-1}{2J-1}\,(c_i - \bar{c}),$$

every markup is proportional to the market share, $p_i - c_i = J\,x_i$, and the average markup is proportional to the Herfindahl index,

$$\sum_i (p_i - c_i)\,x_i = J \cdot HHI.$$

In the toy model with $c_L = 0$ and $c_H = \frac{1}{2}$, a low-cost firm faces the rival average $\frac{1}{3}(p_L + 2p_H)$, since one rival has low cost and two have high cost, and a high-cost firm faces $\frac{1}{3}(2p_L + p_H)$. Substituting into (18.3),

$$p_L = \frac{3 + 2p_H}{5}, \qquad p_H = \frac{9 + 4p_L}{10}, \tag{18.4}$$

so $p_L = \frac{8}{7} \approx 1.143$ and $p_H = \frac{19}{14} \approx 1.357$, with shares $x_L = \frac{2}{7}$ and $x_H = \frac{3}{14}$. The markups are $\frac{8}{7} = 4 \cdot \frac{2}{7}$ and $\frac{19}{14} - \frac{1}{2} = \frac{6}{7} = 4 \cdot \frac{3}{14}$, and

$$HHI = 2\left(\frac{2}{7}\right)^2 + 2\left(\frac{3}{14}\right)^2 = \frac{25}{98}, \qquad \sum_i (p_i - c_i)\,x_i = 2 \cdot \frac{2}{7} \cdot \frac{8}{7} + 2 \cdot \frac{3}{14} \cdot \frac{6}{7} = \frac{50}{49} = 4 \cdot \frac{25}{98}.$$

Figure 18.1 depicts the two reaction curves in (18.4), with the price of the low-cost firms on the horizontal axis. Each curve is the best response of a firm of one type when the other firm of its type charges the same price. They cross at the equilibrium $\left(\frac{8}{7}, \frac{19}{14}\right)$.

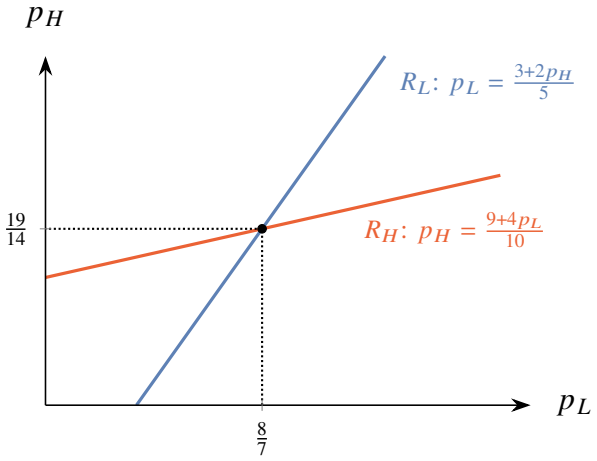


Figure 18.1: Reaction curves of the low-cost firms (R_L) and of the high-cost firms (R_H) in the static market, with $J = 4$, $c_L = 0$, and $c_H = \frac{1}{2}$.

Intuitively, demand (18.2) is linear with the same slope $\frac{1}{J}$ for every firm, so, as for a monopolist facing linear demand, each firm's markup equals its quantity divided by the slope of its demand. A

firm with a cost advantage passes only $\frac{J-1}{2J-1} = \frac{3}{7}$ of it to its price, so it keeps both a larger share and a larger markup. Cost differences therefore raise concentration and the average markup together, and the Herfindahl index summarizes market power in the same way as in the static Cournot model of Cowling and Waterson (1976).

18.4 Second Step: Demand with Switching Costs

Old consumers. Let x_i be firm i's share of the old consumers' past purchases, its *customer base*. On the edge between i and j, an old consumer attached to i buys i if $\theta + p_i \le 1 - \theta + p_j + s$, and an old consumer attached to j buys i if $\theta + p_i + s \le 1 - \theta + p_j$, so these two groups buy i in the fractions $\frac{1}{2}(1 - p_i + p_j + s)$ and $\frac{1}{2}(1 - p_i + p_j - s)$. Adding over edges, demand from old consumers is (equation (10) of the paper)

$$D_i^o = \frac{1}{J}\Big[\, 1 - p_i + \overline{p}_{-i} + (J\,x_i - 1)\,s \,\Big]. \tag{18.5}$$

A base above the average share $\frac{1}{J}$ shifts demand out, which is the source of the milking incentive.

Young consumers. A young consumer who buys i today is attached to i tomorrow. Let p_i' and p_j' be tomorrow's prices. Appendix A of the paper shows that, averaging over tomorrow's location, being attached to i instead of j changes the expected cost tomorrow by $s\,(p_i' - p_j')$: a consumer attached to the product that will be more expensive either pays the higher price or pays the switching cost to leave. The indifferent young consumer is at

$$\theta^* = \frac{1}{2}\Big[\, 1 - p_i + p_j - \delta_c\, s\,\big(p_i' - p_j'\big) \Big]. \tag{18.6}$$

Consider that every firm follows a pricing rule linear in its own base,

$$p_j(x_j) = \alpha_j + \beta\, x_j, \tag{18.7}$$

with a common slope β and its own intercept α_j (the paper calls these *parallel strategies*), and that young consumers know these rules. Tomorrow's base of firm i is its share x_i' of today's young consumers. Write $\overline{\alpha}_{-i}$ for the average intercept of i's rivals and

$$K \equiv J - 1 + \delta_c\, s\, \beta.$$

Result 2 (demand of young consumers; equation (26) of the paper). With rational expectations, firm i's share of the young consumers is

$$x_i' = \frac{1}{J} + \frac{J-1}{J\,K}\Big[\, \overline{p}_{-i} - p_i - \delta_c\, s\,\big(\alpha_i - \overline{\alpha}_{-i}\big) \Big].$$

Its sensitivity to the own price, $\frac{J-1}{JK}$, is below the static value $\frac{1}{J}$ whenever $\delta_c\, s\, \beta > 0$.

To derive it, substitute $p_i' - p_j' = \alpha_i - \alpha_j + \beta(x_i' - x_j')$ in (18.6), add over i's edges using $\sum_{j \ne i} x_j' = 1 - x_i'$, and solve the resulting linear equation for x_i'. At the equilibrium slope of the third step,

$K = 3 + 0.28\,\beta = 3.0979$, so the price sensitivity of young consumers is $\frac{3}{4 \cdot 3.0979} = 0.2421$ instead of $\frac{1}{4}$.

Intuitively, a young consumer who is drawn to a firm by a price cut expects the firm to have a larger base tomorrow and, with $\beta > 0$, a higher price, which she will either pay or escape by paying s. She values the price cut less, and more so the more patient she is and the larger s and β are. This is the anticipation effect: it makes young consumers less price sensitive and softens competition for them.

18.5 Third Step: The Equilibrium and Its Uniqueness

The firm's problem. Firm i's state is its base x_i. Since its rivals' bases add up to $1 - x_i$, rule (18.7) gives $\overline{p}_{-i} = \overline{\alpha}_{-i} + \frac{\beta(1-x_i)}{J-1}$, so a single state variable is enough. Per old consumer, firm i sells $D_i = g\,x_i' + D_i^o$ today, and its value function solves

$$V_i(x_i) = \max_{p_i} \; (p_i - c_i)\,D_i + \delta_f\,V_i(x_i'). \tag{18.8}$$

Demand is linear in p_i and x_i, so the problem is linear-quadratic and $V_i(x) = A_i + B_i\,x + C\,x^2$. Four slopes describe it: today's demand D_i falls by $e = \frac{g(J-1)}{JK} + \frac{1}{J}$ per unit of price and rises by $h = s - \frac{g\beta}{JK} - \frac{\beta}{J(J-1)}$ per unit of base, and tomorrow's base x_i' falls by $a = \frac{J-1}{JK}$ per unit of price and by $b = \frac{\beta}{JK}$ per unit of base. The two negative terms in h appear because a larger own base means smaller rival bases and, therefore, lower rival prices. The first-order condition is

$$D_i - e\,(p_i - c_i) - \delta_f\,a\,\left(B_i + 2C\,x_i'\right) = 0. \tag{18.9}$$

The last term is the investment incentive: a lower price raises tomorrow's base, which is worth $B_i + 2C\,x_i'$ per unit.

The fixed point. If firm i itself follows the slope β, tomorrow's base moves with today's base by $-(a\beta + b) = -\frac{\beta}{K}$, and matching the coefficients of x^2 in (18.8) gives

$$C(\beta) = \frac{\beta\left(s - \frac{\beta}{J-1} - \frac{g\,\beta}{K}\right)}{1 - \delta_f\,\frac{\beta^2}{K^2}}. \tag{18.10}$$

Solving (18.9) for p_i, the slope of firm i's best response to rivals (and future selves) who use the slope β is[1]

$$f(\beta) = \frac{h + 2\,\delta_f\,C(\beta)\,a\,b}{2\left(e - \delta_f\,C(\beta)\,a^2\right)}, \tag{18.11}$$

[1]The paper's equation (65), in its Appendix D, prints the middle term of the numerator of $C(\beta)$ as $(J-1)(1+g\delta_c s^2)\beta$. Matching the coefficients of x^2 gives $(J-1)(\delta_c s^2 - 1 - g)\beta$, that is, $C(\beta) = \frac{\beta K[(J-1)^2 s + (J-1)(\delta_c s^2 - 1 - g)\beta - \delta_c s \beta^2]}{(J-1)(K^2 - \delta_f \beta^2)}$, which equals (18.10). With this expression the model gives the steady-state price of the paper's Figure 1 at its baseline ($J = 5$), 0.935; the printed expression gives 0.921.

and an equilibrium slope solves $\beta = f(\beta)$. The base then evolves as $x_i' = \text{constant} - \frac{\beta}{K} x_i$, so shares do not diverge only if $\left|\frac{\beta}{K}\right| \leq 1$, that is, if β lies in

$$\left[-\frac{J-1}{1+\delta_c s},\ \frac{J-1}{1-\delta_c s} \right] = [-2.3438,\ 4.1667]$$

at the baseline (18.1).

Result 3 (existence and uniqueness; Theorem 1 of the paper). Let $s < 1$, $\delta_c < 1$, $\delta_f < 1$, and $g > 0$. Then:

- there is a unique Markov perfect equilibrium in parallel strategies, that is, a unique solution of $\beta = f(\beta)$ in the interval of non-divergent shares;
- the slope satisfies $0 \leq \beta < \frac{J-1}{1-\delta_c s}$ and the value function is convex, $C > 0$;
- the intercepts are $\alpha_i = \overline{c} + \mu_0 + \mu_1 (c_i - \overline{c})$, with $\mu_0 > 0$ and $0 < \mu_1 < 1$ depending only on $(J, g, s, \delta_c, \delta_f)$.

Figure 18.2 depicts the best-response slope $f(\beta)$ over the whole interval at the baseline. It is decreasing and crosses the 45° line once, at $\beta = 0.3498$, where $C = 0.0600$. The second-order condition holds, since the coefficient of p_i^2 in the maximand is $-(e - \delta_f C a^2) = -0.4903 < 0$, and shares converge, with $\frac{\beta}{K} = 0.1129$.

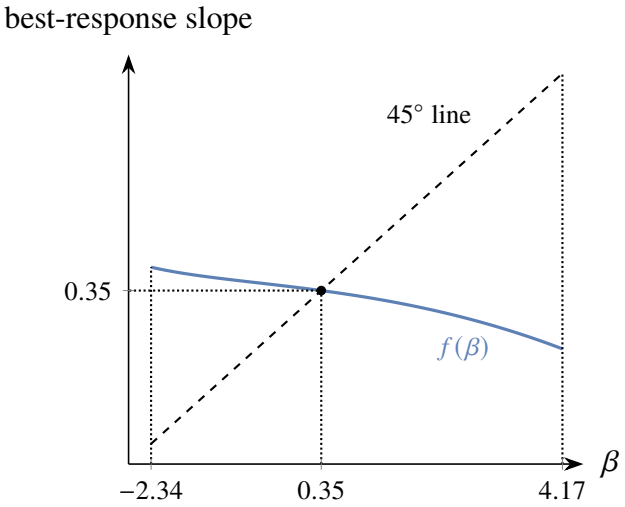


Figure 18.2: Best-response slope $f(\beta)$ against the slope β used by rivals and future selves, over the interval of non-divergent shares, with $J = 4$, $g = 1$, $s = 0.4$, $\delta_c = 0.7$, and $\delta_f = 0.5$.

Intuitively, a firm wants to raise its price when its base is large, because old consumers attached to it are less likely to leave, which is the term s in h. When rivals respond more strongly to their own bases, a firm with a large base faces rivals with small bases and low prices, which takes demand away from it and weakens its own incentive to raise its price. This force makes f fall with β in Figure 18.2, so that f crosses the 45° line once. When firms and consumers are both myopic,

$\delta_c = \delta_f = 0$, the fixed point has the closed form

$$\beta = \frac{h}{2e} = \frac{s\,J(J-1)}{(1+g)(2J-1)} = \frac{6s}{7}, \qquad C = \frac{18\,s^2}{49},$$

which gives $\beta = 0.3429$ at $s = 0.4$; the forward-looking terms move the slope only slightly, to 0.3498.

18.6 Fourth Step: The Steady State and the Three Forces

Under rule (18.7), prices and bases settle at a steady state in which each firm's base equals its share of the young, $x_i' = x_i$. Let $\kappa_X \equiv 1 + \delta_c s$ be the long-run price sensitivity of young consumers.

Result 4 (steady state; Corollary 2 of the paper). In the steady state:

- the average price is $\overline{p} = \overline{c} + \mu_0 + \frac{\beta}{J}$;
- prices move with costs by the pass-through γ^e,

$$p_i = \overline{p} + \gamma^e\,(c_i - \overline{c}), \qquad \gamma^e = \frac{(J-1)\,\mu_1}{J - 1 + \beta\,\kappa_X};$$

- firm i's share of the young is $x_i = \frac{1}{J} - \frac{\kappa_X}{J-1}\,(p_i - \overline{p})$.

Table 18.1 reports the steady state at $s = 0.4$ when the three forces are switched on one at a time. With only milking, $\delta_c = \delta_f = 0$, the average markup is $\overline{p} - \overline{c} = 1$, as in the static market. Adding the investment incentive ($\delta_f = 0.5$) lowers it to 0.9221, adding the anticipation effect ($\delta_c = 0.7$) raises it to 1.0160, and with all three forces it is 0.9381. In every dynamic case the pass-through γ^e is below the static value $\frac{3}{7}$.

Forces	(δ_c, δ_f)	β	μ_0	$\overline{p} - \overline{c} = \mu_0 + \frac{\beta}{4}$	μ_1	$\gamma^e = \frac{3\mu_1}{3+\kappa_X\beta}$
None (static, $s = 0$)	any	0	1	1	0.4286	0.4286
Milking	$(0, 0)$	0.3429	0.9143	1.0000	0.4286	0.3846
Milking, investment	$(0, 0.5)$	0.3443	0.8361	0.9221	0.4503	0.4040
Milking, anticipation	$(0.7, 0)$	0.3483	0.9289	1.0160	0.3973	0.3459
All three (baseline)	$(0.7, 0.5)$	0.3498	0.8507	0.9381	0.4185	0.3642

Table 18.1: Steady state with $J = 4$, $g = 1$, and $s = 0.4$, for different discount factors, with $\kappa_X = 1 + \delta_c s$; entries are computed from unrounded values and rounded to four decimals.

Figure 18.3 depicts the average markup $\overline{p} - \overline{c}$ against the switching cost s for the same four cases. With milking only the curve is flat at 1, with investment it falls, with anticipation it rises, and with all three forces it is U-shaped, with its minimum, 0.921, at $s = 0.756$, as in panel (e) of the paper's Figure 1.

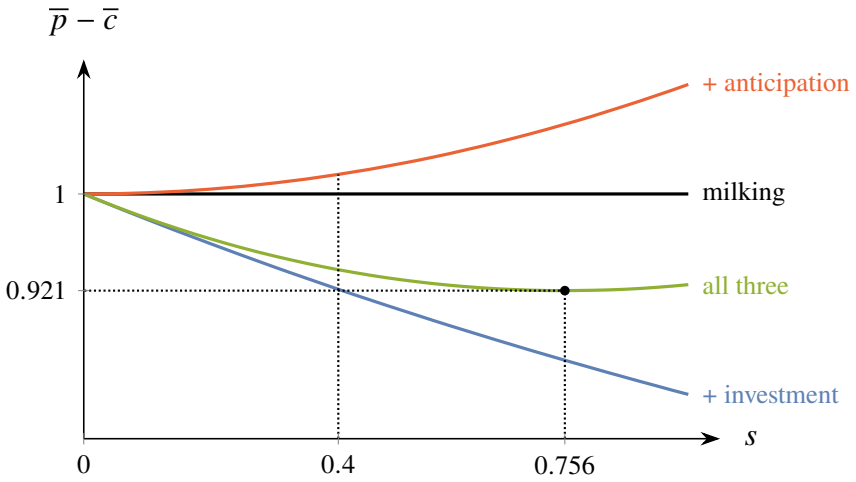


Figure 18.3: Steady-state average markup against the switching cost s, with $J = 4$ and $g = 1$: milking only ($\delta_c = \delta_f = 0$), with investment ($\delta_f = 0.5$), with anticipation ($\delta_c = 0.7$), and all three ($\delta_c = 0.7$, $\delta_f = 0.5$); the dotted line at $s = 0.4$ marks the baseline.

Intuitively, milking alone does not change the average price: at the symmetric steady state every firm's base is $\frac{1}{J}$, so in (18.5) the term $(Jx_i - 1)s$ is zero, and switching costs only make a firm's price respond to its base. That response lowers the pass-through, since a low-cost firm gains share and its larger base pushes its price back up. Patient firms compete for tomorrow's base and cut prices, the effect found by Doganoglu (2010) for small s; patient young consumers anticipate future milking, respond less to price cuts, and allow higher prices. For small switching costs the investment incentive dominates and prices fall below the static level, while for large switching costs the anticipation effect grows faster and prices rise again. Switching costs can therefore make the market either more or less competitive, depending on the discount factors.

18.7 Fifth Step: Markups and Concentration

Firm i's share of all purchases in a period, from both cohorts, is $m_i = \frac{g\,x_i + D_i^o}{1+g}$. In the steady state (equations (14) and (21) of the paper),

$$m_i = \frac{1}{J} - \frac{\kappa_M}{J-1}\,(p_i - \overline{p}), \qquad \kappa_M \equiv \frac{1 + (1+\delta_c s)(g+s)}{1+g}, \qquad HHI = \sum_i m_i^2 = \frac{1}{J} + \frac{\kappa_M^2}{J-1}\,\sigma_p^2, \tag{18.12}$$

where $\sigma_p^2 = \frac{1}{J-1}\sum_i (p_i - \overline{p})^2$ is the variance of prices. At the baseline, $\kappa_M = \frac{1+1.28\cdot 1.4}{2} = 1.396$.

Result 5 (markups and the Herfindahl index; equation (35) of the paper). In the steady state,

$$\sum_i (p_i - c_i)\, m_i = \Lambda \cdot HHI, \qquad \Lambda \equiv \frac{(J-1)\,(1-\gamma^e)}{\kappa_M\,\gamma^e}.$$

With $s = 0$, $\kappa_M = 1$ and $\gamma^e = \frac{J-1}{2J-1}$, so $\Lambda = J$, which is Result 1.

In the toy model the proportionality holds firm by firm, $p_i - c_i = \Lambda\, m_i$, in every case we computed. Table 18.2 reports the two cost types in the static market and at the baseline. At the baseline, $\gamma^e = 0.3642$, so $\Lambda = \frac{3\,(1-0.3642)}{1.396 \cdot 0.3642} = 3.752$, below the static value 4. Then

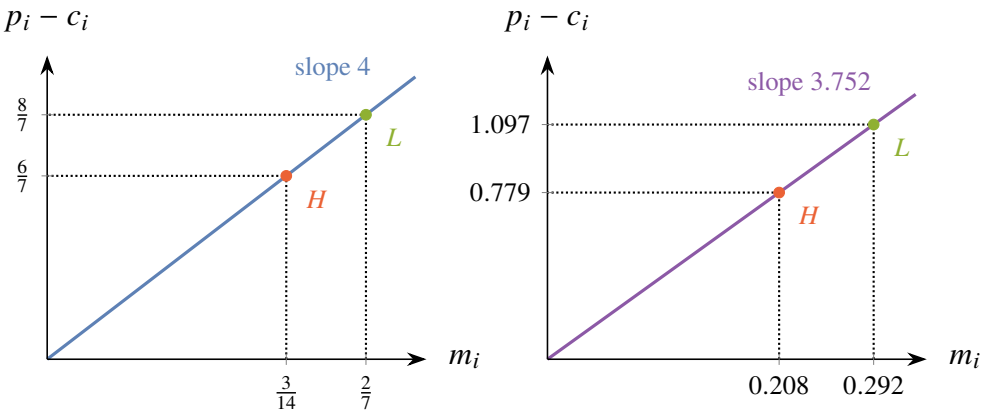
$$HHI = 2\,(0.2924)^2 + 2\,(0.2076)^2 = 0.2572, \qquad \sum_i (p_i - c_i)\, m_i = 3.752 \cdot 0.2572 = 0.9650,$$

against $\frac{25}{98} = 0.2551$ and $\frac{50}{49} = 1.0204$ in the static market.

Market and firm type	c_i	p_i	$p_i - c_i$	m_i	$\frac{p_i - c_i}{m_i}$
Static, low cost	0	$\frac{8}{7} = 1.1429$	1.1429	$\frac{2}{7} = 0.2857$	4
Static, high cost	$\frac{1}{2}$	$\frac{19}{14} = 1.3571$	0.8571	$\frac{3}{14} = 0.2143$	4
Dynamic, low cost	0	1.0971	1.0971	0.2924	3.752
Dynamic, high cost	$\frac{1}{2}$	1.2791	0.7791	0.2076	3.752

Table 18.2: Steady-state prices, markups, and shares of all purchases, with $J = 4$, $g = 1$, $c_L = 0$, and $c_H = \frac{1}{2}$; the dynamic market uses $s = 0.4$, $\delta_c = 0.7$, and $\delta_f = 0.5$.

Figure 18.4 depicts the markup against the share for the two cost types, in the static market in panel (a) and at the baseline in panel (b). In both panels the two firms lie on a ray from the origin, with slope Λ.



(a) Static market (b) Switching costs, baseline

Figure 18.4: Markup against the share of all purchases for the low-cost firms (L) and the high-cost firms (H), with $J = 4$, $g = 1$, $c_L = 0$, and $c_H = \frac{1}{2}$; panel (b) uses $s = 0.4$, $\delta_c = 0.7$, and $\delta_f = 0.5$.

Intuitively, demand stays linear with switching costs, so, as in the first step, a firm's markup is its quantity divided by an effective slope of demand, and a firm with zero share would charge its cost. Switching costs change that effective slope, and therefore the factor Λ, but not the proportionality. As a result, the Herfindahl index remains a sufficient statistic for the average markup among markets with the same number of products and the same dynamic parameters. Across markets with different switching costs or discount factors it is not: at the baseline the Herfindahl index is higher than in the static market, 0.2572 against 0.2551, while the average markup is lower, 0.9650 against 1.0204.

18.8 Sixth Step: Mergers

Consider a merger of the owners of products $J-1$ and J in the static market. The merged firm sets one price p^m for both products, the remaining $J-2$ single-product firms set p_i^s, and the consumers on the edge between the two merged products buy from the merged firm, since the market is covered. Let $\overline{c}_s$ be the average cost of the single-product firms and $\overline{c}_m$ the average cost of the merged products. The merged firm sells to the whole captured edge and competes on $2(J-2)$ edges, so it maximizes a profit proportional to $(p^m - \overline{c}_m)\big[1 + (J-2)(1 - p^m + \overline{p}_s)\big]$, where $\overline{p}_s$ is the average price of the single-product firms.

Result 6 (static mergers; equations (37) to (39) of the paper). After the merger,

$$p_i^s = 1 + \frac{1}{J(J-2)} + \overline{c} + \frac{J-1}{2J-1}(c_i - \overline{c}_s) + \frac{1}{J}(\overline{c}_s - \overline{c}_m),$$
$$p^m = 1 + \frac{J+1}{2J(J-2)} + \overline{c} - \frac{J-3}{2J}(\overline{c}_s - \overline{c}_m),$$

and the average price rises from $1 + \overline{c}$ by

$$\Delta\overline{p} = \frac{2J-1}{J^2(J-2)} + \frac{1}{J^2}(\overline{c}_s - \overline{c}_m).$$

- With equal costs, the merged firm charges more than the single-product firms, since $\frac{J+1}{2J(J-2)} > \frac{1}{J(J-2)}$.
- The average price rises more when the merging firms have low costs, $\overline{c}_s > \overline{c}_m$.

In the toy model, a merger of the two low-cost firms has $\overline{c}_s - \overline{c}_m = \frac{1}{2}$, and every price becomes $\frac{3}{2}$: the single-product firms charge $1 + \frac{1}{8} + \frac{1}{4} + \frac{1}{8}$ and the merged firm $1 + \frac{5}{16} + \frac{1}{4} - \frac{1}{16}$. A merger of the two high-cost firms has $\overline{c}_s - \overline{c}_m = -\frac{1}{2}$, so the single-product firms charge $1 + \frac{1}{8} + \frac{1}{4} - \frac{1}{8} = \frac{5}{4}$ and the merged firm $1 + \frac{5}{16} + \frac{1}{4} + \frac{1}{16} = \frac{13}{8}$. The average price rises from $\frac{5}{4}$ by $\frac{7}{32} + \frac{1}{32} = \frac{1}{4}$ in the first case and by $\frac{7}{32} - \frac{1}{32} = \frac{3}{16}$ in the second.

Table 18.3 reports the effect on welfare. Per consumer, revenue is $R = \sum_i m_i p_i$, production cost is $PC = \sum_i m_i c_i$, and the transport cost is $T = \frac{1}{4} + \frac{1}{2}\sigma_p^2$. Consumers pay $R + T$, the total cost of serving the market is $PC + T$, and profit is $R - PC$. Since every consumer buys one unit, total welfare falls when the total cost rises.

Both mergers raise the consumer cost, by 0.2577 and 0.1717, and raise profit. The low-cost merger raises the total cost by 0.0281, so welfare falls. The high-cost merger lowers it by 0.0110: the production cost falls by $\frac{3}{14} - \frac{3}{16} = 0.0268$, because sales move to the low-cost firms, while the transport cost rises by $\frac{35}{128} - \frac{101}{392} = 0.0158$. Figure 18.5 depicts the change in the average price, $\Delta\overline{p}$, against the number of products J, for the three values of $\overline{c}_s - \overline{c}_m$, with the toy model at $J = 4$.

Intuitively, the merged firm no longer competes for the consumers on the edge between its two products, so it raises its price, and since prices are strategic complements its rivals follow. When the two low-cost firms merge, the remaining rivals have high costs and set high prices, which leaves the merged firm more room to raise its own; in the toy model every price ends at $\frac{3}{2}$, and output

Outcome	Before the merger	Low-cost merger	High-cost merger
Prices (p_1, p_2, p_3, p_4)	$\left(\frac{8}{7}, \frac{8}{7}, \frac{19}{14}, \frac{19}{14}\right)$	$\left(\frac{3}{2}, \frac{3}{2}, \frac{3}{2}, \frac{3}{2}\right)$	$\left(\frac{5}{4}, \frac{5}{4}, \frac{13}{8}, \frac{13}{8}\right)$
Shares (m_1, m_2, m_3, m_4)	$\left(\frac{2}{7}, \frac{2}{7}, \frac{3}{14}, \frac{3}{14}\right)$	$\left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right)$	$\left(\frac{5}{16}, \frac{5}{16}, \frac{3}{16}, \frac{3}{16}\right)$
Average price $\overline{p}$	$\frac{5}{4}$	$\frac{3}{2}$	$\frac{23}{16}$
Revenue R	$\frac{121}{98}$	$\frac{3}{2}$	$\frac{89}{64}$
Production cost PC	$\frac{3}{14}$	$\frac{1}{4}$	$\frac{3}{16}$
Transport cost T	$\frac{101}{392}$	$\frac{1}{4}$	$\frac{35}{128}$
Consumer cost $R + T$	$\frac{585}{392} = 1.4923$	$\frac{7}{4} = 1.7500$	$\frac{213}{128} = 1.6641$
Total cost $PC + T$	$\frac{185}{392} = 0.4719$	$\frac{1}{2} = 0.5000$	$\frac{59}{128} = 0.4609$
Profit $R - PC$	$\frac{50}{49} = 1.0204$	$\frac{5}{4} = 1.2500$	$\frac{77}{64} = 1.2031$

Table 18.3: Static merger of the two low-cost firms (1 and 2) or of the two high-cost firms (3 and 4), with $J = 4$, $c_L = 0$, and $c_H = \frac{1}{2}$; all values per consumer.

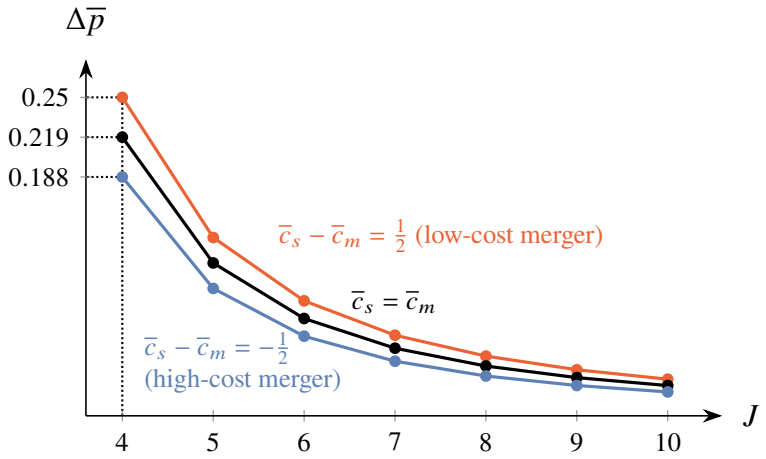


Figure 18.5: Increase in the average price after a static merger of two products against the number of products J, for three differences between the average cost of the single-product firms and of the merged products.

shifts toward the high-cost firms. When the two high-cost firms merge, they face low-cost rivals, so prices rise less, and the price increase of the merged firm shifts sales to the low-cost firms, which saves on production costs without any cost synergy. The effect shrinks quickly as J grows, from $\frac{7}{32}$ with equal costs at $J = 4$ to $\frac{19}{800}$ at $J = 10$, because the merged firm captures only one of $\frac{J(J-1)}{2}$ edges.

18.9 Seventh Step: What the Toy Model Keeps and Drops

The tetrahedron with two cost levels keeps the main conclusions of the paper:

1. With linear transport costs and products equally distant from one another, each firm's demand depends on its rivals only through their average price, and the static markup is proportional to the share, so the average markup is J times the Herfindahl index (the first step, equations (5) and (6)).
2. Forward-looking young consumers respond less to prices when firms milk their bases (the second step, equation (26)).
3. The Markov perfect equilibrium in parallel strategies exists, is unique, and has prices linear in the base and in costs (the third step, Theorem 1).
4. Milking lowers pass-through, investment lowers prices, anticipation raises them, and prices are U-shaped in the switching cost (the fourth step, Corollary 2 and Figure 1 of the paper).
5. The average markup is proportional to the Herfindahl index, with a factor that depends on the dynamic parameters (the fifth step, equation (35)).
6. Mergers raise prices, more so between low-cost firms, and a merger of high-cost firms can raise welfare (the sixth step, equations (37) to (41)).

Relative to Somaini and Einav (2013), the toy model fixes $J = 4$, $g = 1$, and the paper's baseline discount factors and switching cost, and it computes β numerically as the unique root in the interval of non-divergent shares; the paper shows that β solves a quartic equation, so it has a closed form. The toy model drops the dynamic merger, which the paper solves numerically with two types of firms: the effects go in the same direction as in the static market, with magnitudes that the dynamic parameters reduce or amplify, and with fewer than eight symmetric firms the number of firms ranks the price effects of a merger for any value of the dynamic parameters. The toy model keeps the paper's assumptions that consumers consider only the two products at the ends of their edge, redraw their location when old, and always buy, and that firms cannot price discriminate. The spatial structure follows von Ungern-Sternberg (1991) and Chen and Riordan (2007), and the dynamics follow Klemperer (1987), Beggs and Klemperer (1992), and Doganoglu (2010).

18.10 Connections to Other Chapters

- Chapter 12 uses pass-through to turn the pricing pressure created by a merger into a predicted price change. In this chapter, switching costs lower the pass-through of cost differences to prices.
- Chapter 11 studies mergers whose cost savings must be proven to the authority. In this chapter, a merger of two high-cost firms lowers the total cost of serving the market without any cost savings, because sales move to their low-cost rivals.
- Chapter 13 also studies mergers in a dynamic market, but the dynamics come from mergers that arrive in sequence rather than from consumers who are locked in.

References

Beggs, A., and P. Klemperer (1992). "Multi-Period Competition with Switching Costs." *Econometrica*, 60(3), pp. 651–666.

Chen, Y., and M. H. Riordan (2007). "Price and Variety in the Spokes Model." *Economic Journal*, 117(522), pp. 897–921.

Cowling, K., and M. Waterson (1976). "Price-Cost Margins and Market Structure." *Economica*, 43(171), pp. 267–274.

Doganoglu, T. (2010). "Switching Costs, Experience Goods and Dynamic Price Competition." *Quantitative Marketing and Economics*, 8(2), pp. 167–205.

Klemperer, P. (1987). "Markets with Consumer Switching Costs." *Quarterly Journal of Economics*, 102(2), pp. 375–394.

Somaini, P., and L. Einav (2013). "A Model of Market Power in Customer Markets." *Journal of Industrial Economics*, 61(4), pp. 938–986. https://doi.org/10.1111/joie.12039

von Ungern-Sternberg, T. (1991). "Monopolistic Competition on the Pyramid." *Journal of Industrial Economics*, 39(4), pp. 355–368.

Part IV

Collusion and Cartels

The three chapters in this part study when firms that meet repeatedly can sustain collusion. In each chapter, collusion is stable when the gain from deviating today is smaller than the loss from the punishment that follows, and each chapter adds an ingredient that changes one side of this comparison. Chapter 19 adds forward contracts. When firms meet only once, forward sales make competition tougher; when they meet again and again, a deviating firm can only take the part of the market that has not been contracted in advance, so forward contracts make the monopoly outcome easier to sustain. Chapter 20 adds the competition authority: audits, leniency programs, and rewards for whistleblowers lower the payoff from colluding or raise the payoff from reporting the cartel, and rewards to the firm that reports deter collusion more than reduced fines. Chapter 21 drops the assumption that costs are common knowledge. Each member privately learns its cost, and the cartel must induce truthful reports without side payments, which it does by promising a larger future market share to a member that admits high costs today.

Chapters 19 and 20 therefore study the constraint that keeps each firm from undercutting the cartel. Chapter 21 adds a second constraint, which keeps each firm from misreporting its cost, and shows that patient firms can satisfy both without giving up any profit.

Chapter 19

Forward Trading and Collusion

Companion note to Liski, M., and J.-P. Montero (2006), "Forward Trading and Collusion in Oligopoly," Journal of Economic Theory 131(1): 212–230.

19.1 Introduction

Many producers sell part of their output before it is made. An electricity generator, for example, may sign a forward contract today to deliver power next month at a price fixed today. Allaz and Vila (1993) show that, when firms meet only once, such contracts make the spot market more competitive. A firm that sells forward commits to producing more, so both firms sell forward and earn less, as in a prisoner's dilemma. For this reason, forward contracts have been proposed as a way to reduce market power in electricity markets. Does this conclusion survive when the same firms meet again and again?

Liski and Montero (2006) show that it does not. When firms interact repeatedly, they can sustain the monopoly outcome (collude) only if the gain from cheating today is smaller than the loss from the punishment that follows. Forward contracts change both sides of this comparison. A firm that cheats can only steal the part of the market that has not been sold forward, so cheating is less profitable. With quantity competition, forward trading also makes the punishment harsher, since after a deviation each firm keeps selling forward to commit to a larger output.

Therefore, the possibility of trading forward makes collusion easier: it lowers the critical discount factor, the smallest degree of patience at which firms can sustain the monopoly outcome. With quantity competition, this holds even when no forwards are traded along the collusive path, because the harsher punishment alone deters cheating. Forward contracts are therefore not, on their own, a remedy for market power in repeated markets.

The toy model uses linear demand and zero costs. The first step finds the critical discount factors without forward trading, and the second step solves the one-shot game of Allaz and Vila. The third and fourth steps add forward trading under price and under quantity competition, the fifth

step checks deviations in the forward market, and the sixth step lists what the toy model keeps and drops.

19.2 The Toy Model

Firms and demand. Two firms produce a homogeneous good at zero marginal cost. In every spot market, the inverse demand is

$$p = 1 - Q, \tag{19.1}$$

where $Q = q_1 + q_2$ is total output. As in the paper, (19.1) holds for every Q, so the price is negative when total output exceeds one. The paper uses $p = a - Q$ and marginal cost c; every profit in the paper equals the corresponding profit below multiplied by $(a-c)^2$, so every critical discount factor is the same as in the paper.

Monopoly benchmark. A monopolist produces $q_m = \frac{1}{2}$, charges $p_m = \frac{1}{2}$, and earns $\pi_m = \frac{1}{4}$. If the two firms collude on this outcome and split it equally, each firm earns $\frac{\pi_m}{2} = \frac{1}{8}$ per spot market.

Forward contracts. Before each spot market, a forward market opens in which each firm can sell ($f_i > 0$, a short position) or buy ($f_i < 0$, a long position) contracts for delivery in that spot market. Forward positions are observable, and competitive speculators buy or sell contracts at a forward price equal to the spot price they expect, so there are no arbitrage profits. Once firm i holds a forward position f_i, it sells only $q_i - f_i$ units in the spot market, so its spot payoff is

$$\pi_i^s = p\,(q_i - f_i), \tag{19.2}$$

while the revenue from its forward sales was collected when it sold them. Since forwards sell at the expected spot price, firm i's total profit from a spot market, counting its forward revenue, is $p\,q_i$.

Repeated interaction. The two firms meet in an infinite sequence of spot markets, each preceded by a forward market. The discount factor between two consecutive spot markets is $\delta \in (0,1)$. A firm that deviates from a collusive plan is punished by its rival in all later periods. The question is how small δ can be for the firms to sustain the monopoly outcome, and how forward trading changes this critical discount factor.

19.3 First Step: Collusion without Forward Trading

Quantity competition. In a one-shot Cournot game, firm i maximizes $(1 - q_i - q_j)\,q_i$, so its best response is $q_i = \frac{1-q_j}{2}$. The equilibrium is $q_1 = q_2 = \frac{1}{3}$, with price $\frac{1}{3}$ and profit $\frac{1}{9}$ per firm. If the rival produces its collusive output $\frac{q_m}{2} = \frac{1}{4}$, a deviating firm produces

$$\frac{1 - \frac{1}{4}}{2} = \frac{3}{8},$$

the price falls to $\frac{3}{8}$, and the deviator earns $\frac{9}{64}$. With grim-trigger strategies that revert to the Cournot equilibrium forever, collusion is sustainable if

$$\frac{1/8}{1-\delta} \;\geq\; \frac{9}{64} + \frac{\delta}{1-\delta}\cdot\frac{1}{9} \quad\Longleftrightarrow\quad \delta \;\geq\; \frac{9}{17} \approx 0.529. \tag{19.3}$$

Price competition. In a one-shot Bertrand game, both firms price at marginal cost and earn zero. A firm that undercuts the monopoly price slightly takes the whole market and earns $\pi_m = \frac{1}{4}$ in the deviation period, so collusion is sustainable if

$$\frac{1/8}{1-\delta} \;\geq\; \frac{1}{4},$$

that is, $\delta \geq \frac{1}{2}$.

Result 1 (no forward trading). The monopoly outcome can be sustained:

- with quantity competition and Cournot reversion, if and only if $\delta \geq \frac{9}{17} \approx 0.529$;
- with price competition and Bertrand reversion, if and only if $\delta \geq \frac{1}{2}$.

Reverting to the Cournot equilibrium is not the harshest credible punishment. With an optimal penal code (Abreu, 1986), firms punish a deviation with one period of high output, $z = \frac{5}{12}$ each, and then return to collusion; the critical discount factor falls to $\frac{9}{32} \approx 0.281$. This cutoff uses (19.1) also at negative prices, which makes a harsh punishment costly for the punishing firms. If the price cannot fall below zero, producing one unit each is a stage Nash equilibrium with zero profit, and reverting to it forever sustains collusion if and only if $\delta \geq \frac{1}{9}$. We return to this benchmark in the sixth step.

19.4 Second Step: Forward Trading in a One-Shot Market

The spot market. Given forward positions f_i and f_j, firm i chooses q_i to maximize its spot payoff (19.2), $(1 - q_i - q_j)(q_i - f_i)$. The first-order condition,

$$1 - 2q_i - q_j + f_i = 0,$$

gives the best response

$$q_i = \frac{1 - q_j + f_i}{2}. \tag{19.4}$$

A forward sale shifts firm i's best response outward: the firm no longer cares about the price it receives on the f_i units it has already sold, so it behaves as if it were larger. Solving the two best responses,

$$q_i = \frac{1 + 2f_i - f_j}{3}, \qquad p = \frac{1 - f_i - f_j}{3}. \tag{19.5}$$

Figure 19.1 depicts the best responses (19.4) without forward positions (solid) and when both firms sell $f = \frac{1}{5}$ forward (dashed). Both best responses shift outward by $\frac{f}{2}$, and the equilibrium moves from the Cournot point C, with $q_i = \frac{1}{3}$, to point AV, with $q_i = \frac{2}{5}$.

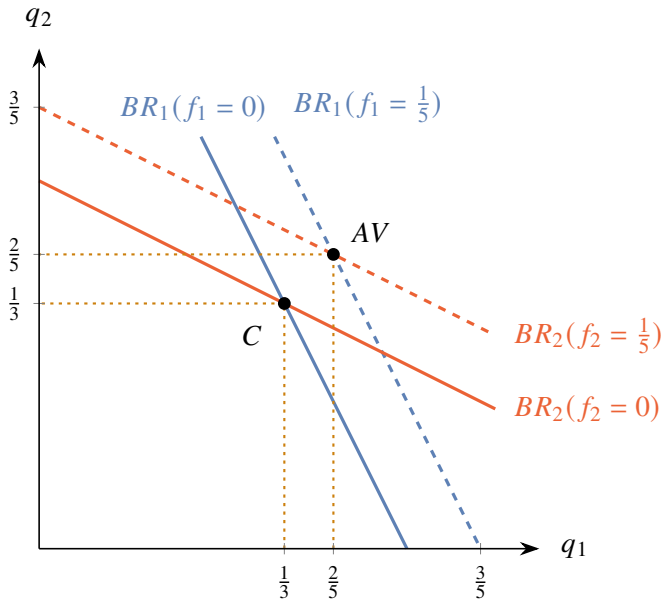


Figure 19.1: Spot-market best responses without forward positions (solid) and with forward sales $f_1 = f_2 = \frac{1}{5}$ (dashed).

The forward market. Anticipating (19.5), and since forwards sell at the expected spot price, firm i chooses f_i to maximize its total profit,

$$p\, q_i = \frac{(1 - f_i - f_j)(1 + 2f_i - f_j)}{9}. \tag{19.6}$$

The first-order condition,

$$-(1 + 2f_i - f_j) + 2(1 - f_i - f_j) = 0,$$

gives $f_i = \frac{1-f_j}{4}$, so in the symmetric equilibrium $f_i = \frac{1}{5}$. Substituting into (19.5), $q_i = \frac{2}{5}$ and $p = \frac{1}{5}$, and each firm earns

$$\frac{2}{25} = 0.08 < \frac{1}{9} \approx 0.111.$$

Result 2 (Allaz and Vila, one forward round). When a single forward market precedes the spot market, each firm sells $\frac{1}{5}$ forward and produces $\frac{2}{5}$; the price falls from $\frac{1}{3}$ to $\frac{1}{5}$, and each firm's profit falls from $\frac{1}{9}$ to $\frac{2}{25}$.

Intuitively, selling forward works as a commitment to produce more, which leads the rival to produce less, as in a Stackelberg game. Each firm therefore wants to sell forward, but when both do it, both produce more and both earn less than without a forward market, a prisoner's dilemma.

Several forward rounds. Consider now that N forward markets open before the spot market. In the last round, a firm that already holds P_i from earlier rounds chooses its final position F_i, and its payoff is $p\,(q_i - P_i)$, since only the units it sells in this round earn the price it can still affect. The first-order condition gives the best response $F_i = \frac{1-F_j+3P_i}{4}$, so the last round's equilibrium is

$$F_i = \frac{1 + 4P_i - P_j}{5}.$$

In the earlier round, firm i anticipates this response and maximizes

$$p\,q_i = \frac{(1 - P_i - P_j)(2 + 3P_i - 2P_j)}{25},$$

whose first-order condition gives $P_i = \frac{1}{7}$ in a symmetric equilibrium. With $N = 2$ rounds, each firm therefore ends up selling $\frac{2}{7}$ forward, produces $\frac{3}{7}$, and the price is $\frac{1}{7}$. Repeating these steps for any N,

$$F(N) = \frac{N}{3+2N}, \qquad q(N) = \frac{1+N}{3+2N}, \qquad p(N) = \frac{1}{3+2N}, \qquad \pi(N) = \frac{1+N}{(3+2N)^2}. \tag{19.7}$$

Result 3 (Allaz and Vila, N forward rounds). Each firm's profit $\pi(N)$ in (19.7) decreases in the number of forward rounds and tends to zero, the competitive outcome, as N grows.

Figure 19.2 depicts $\pi(N)$ for $N = 1, \ldots, 12$, together with the Cournot profit $\frac{1}{9}$ and the collusive profit $\frac{1}{8}$.

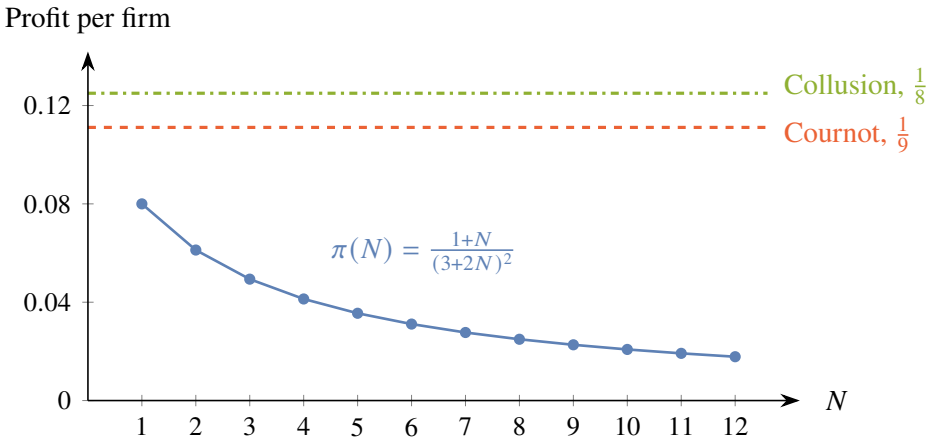


Figure 19.2: Profit per firm when N forward rounds precede the spot market.

Intuitively, every additional forward round gives each firm another chance to commit to a larger output, so the market approaches the competitive outcome. This path will serve as the punishment in the fourth step.

19.5 Third Step: Collusion with Price Competition in the Spot Market

The collusive plan. Firms now meet repeatedly. Along the collusive path, each firm sells $x\frac{q_m}{2} = \frac{x}{4}$ units forward one period ahead, with $0 \le x \le 1$, and both firms charge $p_m = \frac{1}{2}$ in the spot market. A share x of the monopoly output is thus contracted in advance, and each firm sells the rest, $(1-x)\frac{q_m}{2}$, in the spot market. If any firm deviates, both price at marginal cost forever after.

Deviations. A deviation in the forward market is never profitable: speculators see it immediately and anticipate competitive pricing in the next spot market, so they pay at most the marginal cost, zero, for the deviator's contracts. In the spot market, the contracted output $x\,q_m = \frac{x}{2}$ is already committed, so a deviating firm faces the residual demand $1 - \frac{x}{2} - p$. Undercutting slightly is no longer optimal; the deviator sets

$$p_d = \arg\max_p \; p\left(1 - \frac{x}{2} - p\right) = \frac{2-x}{4}, \qquad \text{earning} \qquad \pi_d = \frac{(2-x)^2}{16}. \tag{19.8}$$

Critical discount factor. At the opening of a spot market, a firm that follows the plan earns its uncontracted share of monopoly profit in that market, $(1-x)\frac{1}{8}$, plus $\frac{1}{8}$ in every later market. The firm does not deviate if

$$\frac{1-x}{8} + \frac{\delta}{1-\delta}\cdot\frac{1}{8} \;\geq\; \frac{(2-x)^2}{16}, \tag{19.9}$$

which holds if and only if

$$\delta \;\geq\; \delta_B(x) \equiv \frac{x^2 - 2x + 2}{x^2 - 2x + 4}. \tag{19.10}$$

Result 4 (Proposition 1 of the paper). With price competition in the spot market, the critical discount factor $\delta_B(x)$ decreases in the share x of output sold forward, from $\delta_B(0) = \frac{1}{2}$ to $\delta_B(1) = \frac{1}{3}$.

Figure 19.3 depicts $\delta_B(x)$. Collusion can be sustained above the curve. For $x < 0$, when firms buy forwards (long positions), the curve is flat at $\frac{1}{2}$: a deviator can undercut and take the whole spot market, and speculators close their positions at the same price whether a deviation occurs or not (Proposition 2 of the paper).

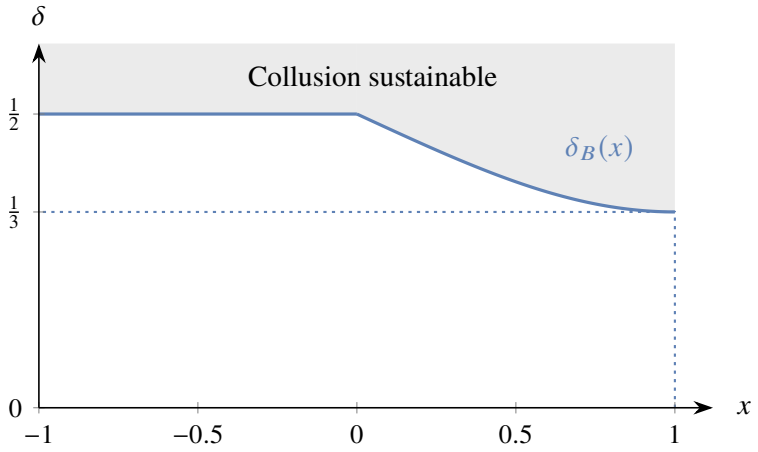


Figure 19.3: Critical discount factor with price competition in the spot market, as a function of the share x of monopoly output sold forward one period ahead ($x < 0$: long positions).

Intuitively, forward trading has two opposing effects. On the one hand, it lowers the payoff from following the plan in the current period, since the contracted share x of current sales has already been paid for; this is the pro-competitive effect. On the other hand, it lowers the payoff from deviating, since the contracted units already serve part of the market and the deviator can only capture the residual demand; this is the pro-collusive effect. With price competition, the second effect dominates for every $x \in [0, 1]$.

19.6 Fourth Step: Collusion with Quantity Competition in the Spot Market

The collusive plan. Along the collusive path, each firm sells $\frac{x}{4}$ units forward one period ahead, now with $-1 \le x \le 1$, and supplies another $(1-x)\frac{1}{4}$ in the spot market, so that each firm's total output is $\frac{1}{4}$ and the price is $\frac{1}{2}$. If any firm deviates, both play the Allaz–Vila equilibrium thereafter. Since firms can trade in every forward market for delivery in any later spot market, the spot market that opens N periods after a deviation is preceded by N forward openings, and each firm earns $\pi(N)$ in (19.7) there.

Deviations. In the spot market, the rival produces $\frac{1}{4}$ in total, of which $\frac{x}{4}$ was sold forward. A deviating firm chooses its total output Q_i to maximize

$$\left(1 - \frac{1}{4} - Q_i\right)\left(Q_i - \frac{x}{4}\right),$$

which gives

$$Q_i = \frac{3+x}{8}, \qquad p_d = \frac{3-x}{8}, \qquad \pi_d = \frac{(3-x)^2}{64}. \tag{19.11}$$

Since π_d decreases in x, forward sales again reduce the profit from a deviation. With $x = 0$, $\pi_d = \frac{9}{64}$, as in the first step. For $x \ge 0$, a deviation in the forward market is not profitable: against a rival who sells $\frac{x}{4}$, the best forward sale is $\frac{1}{4}\left(1 - \frac{x}{4}\right)$, which triggers the punishment and earns

$$\frac{1}{8}\left(1 - \frac{x}{4}\right)^2 \;\le\; \frac{1}{8}$$

in the next spot market.

A first look: one forward round per period. Consider, only for this paragraph, that firms could trade forwards only in the period right before delivery, so the punishment is the one-round equilibrium of Result 2 forever, with profit $\frac{2}{25}$ per period. With $x = 0$, collusion is sustainable if

$$\frac{1/8}{1-\delta} \;\ge\; \frac{9}{64} + \frac{\delta}{1-\delta}\cdot\frac{2}{25} \quad\Longleftrightarrow\quad \delta \;\ge\; \frac{25}{97} \approx 0.258. \tag{19.12}$$

The deviation profit is the same as without forward trading, $\frac{9}{64}$; the critical discount factor falls from $\frac{9}{17}$ to $\frac{25}{97}$ only because the punishment profit falls from $\frac{1}{9}$ to $\frac{2}{25}$.

The full punishment. When forward positions accumulate after a deviation, the punishment profit falls to zero over time (Figure 19.2). At the opening of a spot market, the firm does not deviate if

$$\frac{1-x}{8} + \frac{\delta}{1-\delta}\cdot\frac{1}{8} \;\ge\; \frac{(3-x)^2}{64} + \sum_{N=1}^{\infty}\delta^N\,\frac{1+N}{(3+2N)^2}. \tag{19.13}$$

The left-hand side increases in δ faster than the right-hand side, so (19.13) holds if and only if $\delta \ge \delta_C(x)$, where $\delta_C(x)$ solves (19.13) with equality. The sum has no simple closed form, so we solve for $\delta_C(x)$ numerically.

Result 5 (Proposition 3 of the paper, short positions). With quantity competition in the spot market and $0 \le x \le 1$:

- $\delta_C(0) = 0.237$, $\delta_C(\frac{1}{2}) = 0.391$, and $\delta_C(1) = 0.512$;
- $\delta_C(x)$ increases in x, and $\delta_C(x) < \frac{9}{17}$ for every $x \in [0, 1]$.

The paper reports $\delta_C(0) = 0.238$; solving (19.13) to more digits gives 0.2372. Figure 19.4 depicts $\delta_C(x)$ for $x \in [0, 1]$ (solid), together with the pure-spot cutoff $\frac{9}{17}$ with Cournot reversion and the cutoff $\frac{9}{32}$ with an optimal penal code. Every point between the solid curve and the line at $\frac{9}{17}$ is a discount factor at which collusion can be sustained with forward trading but not without it. The part of the figure for $x < 0$ is the subject of the fifth step.

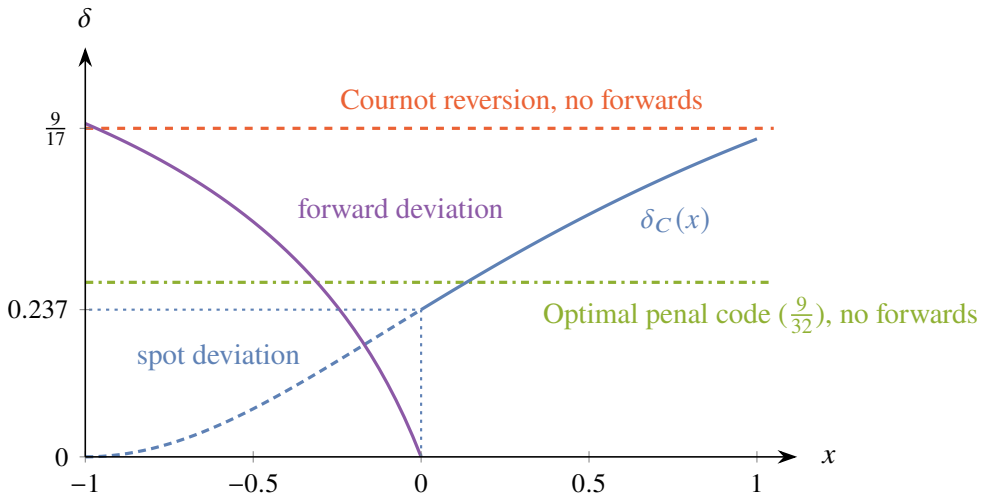


Figure 19.4: Critical discount factor with quantity competition in the spot market. For $x < 0$, the dashed curve is the cutoff against a spot-market deviation and the purple curve the cutoff against a forward-market deviation (see the fifth step).

Intuitively, the two effects of the third step are still present, but a third one appears: the punishment is harsher than Cournot reversion, because after a deviation each firm keeps selling forward to commit to a larger output. This third effect does not depend on x, and it is the reason why $\delta_C(0) = 0.237$ lies below $\frac{9}{17}$ even though nobody sells forward along the collusive path: the threat of falling into a situation of heavy forward contracting is what deters a deviation. In contrast with price competition, a larger x now raises the critical discount factor, because with quantity competition forward sales lower the payoff from following the plan, $\frac{1-x}{8}$, by more than they lower the deviation profit, $\frac{(3-x)^2}{64}$.

Figure 19.5 compares the two sides of the no-deviation conditions (19.9) and (19.13) at $\delta = 0.4$, as functions of x. With price competition (panel (a)), the payoff from deviating falls faster than the payoff from following the plan, so collusion requires enough forward sales,

$$x \;\geq\; 1 - \frac{1}{\sqrt{3}} \approx 0.423.$$

With quantity competition (panel (b)), the payoff from following the plan falls faster, so collusion requires few forward sales, $x \leq 0.534$.

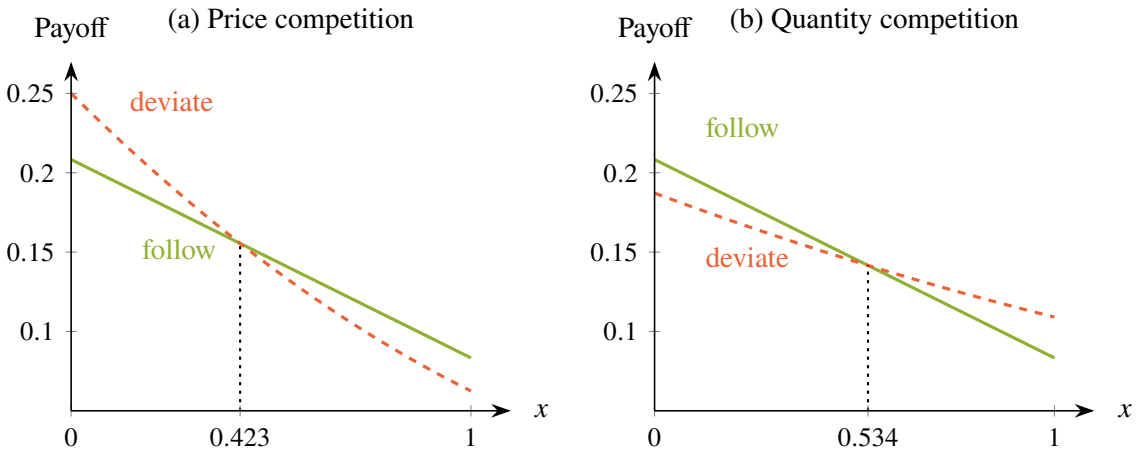


Figure 19.5: Payoff from following the collusive plan and from deviating in the spot market, as a function of the share x sold forward, for $\delta = 0.4$.

19.7 Fifth Step: Long Positions and Deviations in the Forward Market

When firms buy forwards along the collusive path ($x < 0$), the spot-market deviation becomes less attractive still: with $x = -1$, each firm buys forward its entire collusive output, $\frac{1}{4}$, and $\pi_d = \frac{1}{4}$ equals the payoff from following the plan in that period, so condition (19.13) holds for every δ (the dashed part of the blue curve in Figure 19.4). Long positions are not exotic: Mahenc and Salanié (2004) show that, with price competition in a one-shot market, firms buy part of their own output forward to soften competition. The paper concludes that $\delta_C(-1) = 0$. Its argument that a deviation in the forward market is never profitable, however, relies on the deviation profit $\frac{1}{8}\left(1 - \frac{x}{4}\right)^2$ being at most $\frac{1}{8}$, which holds only for $x \geq 0$.

With long positions, a firm can deviate in the forward market by selling forwards instead of buying them. Its rival has committed to buying $\frac{|x|}{4}$, which makes the rival less aggressive in the next spot market, and the deviator earns $\frac{1}{8}\left(1 - \frac{x}{4}\right)^2 > \frac{1}{8}$ there, followed by the punishment path. The firm does not deviate in this way if

$$\frac{1}{1-\delta} \cdot \frac{1}{8} \;\geq\; \frac{1}{8}\left(1 - \frac{x}{4}\right)^2 + \sum_{N=1}^{\infty} \delta^N \, \frac{1+N}{(3+2N)^2}. \tag{19.14}$$

Result 6 (long positions). With quantity competition in the spot market and $x < 0$:

- condition (19.14) fails for small δ, so the critical discount factor is at least the cutoff of (19.14), the purple curve in Figure 19.4;
- at $x = -1$, this cutoff is 0.537, above the pure-spot cutoff $\frac{9}{17}$;
- taking both deviations into account, the lowest cutoff, about 0.181, is reached with a small long position, $x \approx -0.17$.

The cutoff in Result 6 is a lower bound, since it considers only forward sales for the next spot market; a deviator could also sell forwards for later markets. The qualitative message of the paper survives against the deviations considered here: forward trading still lowers the critical discount factor below $\frac{9}{17}$ for every x above about -0.97. But firms cannot rely on large long positions to sustain collusion at any discount factor, because a firm that is expected to buy forwards has an incentive to sell them instead. The general lesson is that a collusive plan with two markets must be checked against deviations in both of them.

19.8 Sixth Step: What the Toy Model Keeps and Drops

Table 19.1 reports the critical discount factors of the first to fifth steps.

Table 19.1: Critical discount factors for sustaining the monopoly outcome ($p = 1 - Q$, zero costs).

Setting	Critical δ
Price competition, no forwards	$\frac{1}{2}$
Price competition, all output sold forward one period ahead ($x = 1$)	$\frac{1}{3}$
Quantity competition, no forwards, Cournot reversion	$\frac{9}{17} \approx 0.529$
Quantity competition, no forwards, optimal penal code	$\frac{9}{32} \approx 0.281$
Quantity competition, one forward round per period, $x = 0$	$\frac{25}{97} \approx 0.258$
Quantity competition, forwards, $x = 0$	0.237
Quantity competition, forwards, $x = 1$	0.512
Quantity competition, forwards, $x = -1$ (forward-market deviation)	at least 0.537

What the toy model keeps.

- The two effects of forward trading on collusion: fewer uncontracted sales along the collusive path (pro-competitive) and a smaller market for a deviator plus, with quantity competition, a harsher punishment (pro-collusive).
- The contrast between price and quantity competition: collusion calls for selling forward with price competition, but for selling little or nothing with quantity competition.
- The result that forward trading can help collusion even when no forwards are traded along the equilibrium path, because of the punishment that forward trading makes possible.

What it drops.

- Optimal punishments with forward trading. Without forward trading, an optimal penal code lowers the cutoff to $\frac{9}{32}$, which is still above $\delta_C(0) = 0.237$; the paper does not derive the optimal penal code with forward trading.
- Price setting in the forward market. If firms set the prices of their forward contracts, the

punishment becomes competitive pricing, and with $x = 0$ the cutoff falls to $\frac{1}{9}$, since

$$\frac{1/8}{1-\delta} \geq \frac{9}{64}$$

if and only if $\delta \geq \frac{1}{9}$.

- Contracting more than one period ahead, hedging demand from risk-averse firms or buyers, and more than two firms.
- Regulation. If a regulator requires firms to keep large forward sales, the pro-competitive effect can dominate. Forward contracting is therefore not, on its own, a remedy for market power in repeated markets such as electricity.

19.9 Connections to Other Chapters

- Chapter 20 studies policies that make collusion harder to sustain: antitrust audits, leniency programs, and whistleblower rewards.
- Chapter 5 studies another contract signed before competition takes place, a rebate conditional on exclusivity, and shows that, unlike an exclusive dealing contract, it does not commit the buyer.
- Chapter 21 studies collusion when each firm privately observes its cost, and shows how patient firms keep production efficient by promising future market share.

References

Abreu, D. (1986). Extremal equilibria of oligopolistic supergames. *Journal of Economic Theory*, 39(1), 191–225.

Allaz, B., and Vila, J.-L. (1993). Cournot competition, forward markets and efficiency. *Journal of Economic Theory*, 59(1), 1–16.

Liski, M., and Montero, J.-P. (2006). Forward trading and collusion in oligopoly. *Journal of Economic Theory*, 131(1), 212–230.

Mahenc, P., and Salanié, F. (2004). Softening competition through forward trading. *Journal of Economic Theory*, 116(2), 282–293.

Chapter 20

Collusion, Leniency Programs, and Whistleblower Rewards

Companion note to Aubert, C., P. Rey, and W.E. Kovacic (2006), "The Impact of Leniency and Whistle-Blowing Programs on Cartels," International Journal of Industrial Organization 24(6): 1241–1266.

20.1 Introduction

Price-fixing cartels are illegal, yet they keep forming. Consider two cement producers that agree to charge high prices instead of competing. The agreement lasts only if each firm prefers the stream of cartel profits to the one-time gain from undercutting its partner. How can a competition authority make this agreement harder to sustain?

Authorities use three tools. Audits detect some cartels and fine their members. Leniency programs reduce the fine of a cartel member that reports the cartel. Whistleblower programs pay a bounty to individuals, such as employees, who report it. Each tool works by changing the comparison above. A colluding firm faces the expected fine in every period but a deviating firm only once, so audits make collusion less attractive. A leniency reward goes only to the firm that betrays the cartel, so it makes deviation more attractive. A bounty for employees forces the cartel to pay its informed staff to keep quiet, period after period, which lowers the return from colluding.

Aubert et al. (2006) show that each tool raises the minimal discount factor, the smallest degree of patience at which firms can sustain collusion. Positive rewards to a reporting firm deter collusion more than reduced fines alone. Bounties for employees deter it even more once the total bounty is large enough, since the cartel must buy the silence of every informed employee. These results support paying informants, a policy that some competition authorities already follow.

The toy model is a Bertrand duopoly with linear demand and zero costs, repeated forever. The first step finds the minimal discount factor without enforcement, the second step adds audits, the third step adds a leniency program, and the fourth step adds a whistleblower program and ranks the four

regimes. A short section then describes how competition authorities use these programs, and the fifth step lists what the toy model keeps and drops.

20.2 The Toy Model

Firms and demand. Consider two firms selling a homogeneous product and competing à la Bertrand. They face a linear demand function $Q(p) = 1 - p$, where Q denotes aggregate output and p represents the market price. Both firms have the same marginal cost, normalized to $c = 0$.

Repeated interaction. Firms interact in an infinitely repeated game with common discount factor $\delta \in (0, 1)$ and use grim-trigger strategies: any deviation from a collusive agreement triggers permanent reversion to the static Bertrand–Nash equilibrium.

Profits. We use the following notation throughout:

- π^m: monopoly profit. Solving $\max_{p \geq 0} \; p(1 - p)$ yields $p^m = \frac{1}{2}$ and $\pi^m = \frac{1}{4}$.
- $\pi^c = \frac{\pi^m}{2} = \frac{1}{8}$: each firm's collusive profit (equal split).
- $\pi^d = \pi^m = \frac{1}{4}$: deviation profit (undercut the collusive price by ε and capture all sales).
- $\pi^n = 0$: Nash equilibrium profit (Bertrand with symmetric costs).

The question is how small δ can be for the firms to sustain collusion, and how audits, leniency programs, and whistleblower programs run by a competition authority (CA) change this minimal discount factor.

20.3 First Step: No Fines

When the CA does not monitor collusion, each firm compares the present value of colluding forever with the payoff from a one-shot deviation followed by permanent Nash punishment. Collusion is sustainable if

$$\frac{\pi^c}{1 - \delta} \;\geq\; \pi^d + \frac{\delta}{1 - \delta}\,\pi^n.$$

Substituting $\pi^c = \frac{1}{8}$, $\pi^d = \frac{1}{4}$, and $\pi^n = 0$, and rearranging:

$$\frac{1}{8} \;\geq\; (1 - \delta)\,\frac{1}{4} \qquad \Longrightarrow \qquad \delta \;\geq\; \overline{\delta} \;=\; \frac{1}{2}.$$

Result 1 (no enforcement). Without enforcement, collusion is sustainable if and only if $\delta \geq \overline{\delta} = \frac{1}{2}$.

Intuitively, when firms are patient enough, $\delta \geq \frac{1}{2}$, the long-run gain from collusion outweighs the short-run temptation to deviate. This threshold, $\overline{\delta} = \frac{1}{2}$, is the benchmark we compare against when the enforcement tools of the second to fourth steps are introduced.

20.4 Second Step: Antitrust Audits

Consider that the CA audits the industry in every period with probability $\alpha \in [0, 1]$. If collusion is detected, each firm pays a fine $F > 0$. Both colluding and deviating firms face the expected fine αF in the current period (the cartel was active that period). The sustainability condition becomes

$$\frac{1}{1-\delta}\,(\pi^c - \alpha F) \;\geq\; \Big(\pi^d - \alpha F\Big) + \frac{\delta}{1-\delta}\,\pi^n.$$

Since $\pi^c = \frac{1}{8}$, $\pi^d = \frac{1}{4}$, and $\pi^n = 0$:

$$\frac{1}{8} - \alpha F \;\geq\; (1-\delta)\left(\frac{1}{4} - \alpha F\right).$$

Solving for δ:

$$\delta \;\geq\; \overline{\delta}_A(\alpha, F) \;\equiv\; \frac{1}{2(1-4\alpha F)}\,.$$

For $\alpha F < \frac{1}{8}$, this threshold lies between $\frac{1}{2}$ and 1, so collusion can be sustained when

$$F \;<\; \frac{1}{8\alpha}.$$

If instead $F \geq \frac{1}{8\alpha}$, the no-deviation condition above fails for every $\delta \in (0, 1)$, so no discount factor sustains collusion: the expected fine alone deters the cartel.

Comparative statics. Differentiating $\overline{\delta}_A$,

$$\frac{\partial \overline{\delta}_A}{\partial \alpha} = \frac{2F}{(1-4\alpha F)^2} > 0, \qquad \frac{\partial \overline{\delta}_A}{\partial F} = \frac{2\alpha}{(1-4\alpha F)^2} > 0,$$

and $\overline{\delta}_A \geq \overline{\delta} = \frac{1}{2}$ for all $\alpha F \geq 0$.

Result 2 (antitrust audits). With audits:

- if $\alpha F < \frac{1}{8}$, collusion is sustainable if and only if $\delta \geq \overline{\delta}_A(\alpha, F)$, where $\frac{1}{2} \leq \overline{\delta}_A < 1$ and $\overline{\delta}_A$ increases in α and in F;
- if $\alpha F \geq \frac{1}{8}$, no discount factor $\delta \in (0, 1)$ sustains collusion.

Figure 20.1 depicts $\overline{\delta}_A$ as a function of the audit probability α, for $F = 0.08$ and $F = 0.15$. With $F = 0.15$, the threshold reaches one at $\alpha = \frac{5}{6}$, where $\alpha F = \frac{1}{8}$.

Intuitively, a colluding firm faces the expected fine αF in every period, while a deviating firm faces it only in the period of its deviation, since the cartel collapses afterwards. Both a more severe fine (F) and a more likely audit (α) therefore raise the minimal discount factor, making collusion harder to sustain. Since $\overline{\delta}_A \geq \overline{\delta} = \frac{1}{2}$, audits unambiguously hinder collusion.

When $F \geq \frac{1}{8\alpha}$, the expected fine alone deters the cartel, leaving nothing for a leniency or whistleblower program to add. We therefore maintain

$$F \;<\; \frac{1}{8\alpha}$$

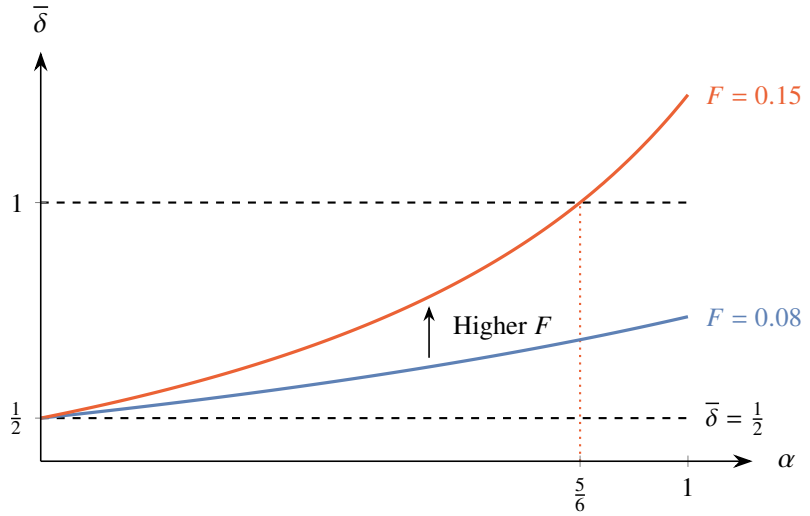


Figure 20.1: Minimal discount factor under audits, $\overline{\delta}_A(\alpha, F)$, for $F = 0.08$ and $F = 0.15$. Collusion is sustainable for discount factors above each curve, and an increase in F shifts the curve upward.

for the rest of this chapter, so that collusion is sustainable in this audit benchmark and the enforcement tools below have a cartel to break.

20.5 Third Step: Leniency Programs

The CA now offers a leniency program: a firm that provides evidence about collusion receives a reward R, where $R > -\alpha F$. The reward can be negative (a reduced fine, $0 > R > -\alpha F$) or positive (full immunity plus a cash reward, $R > 0$). A deviating firm can simultaneously deviate and report the cartel, collecting R instead of facing the expected fine. Recall from the second step that we maintain $F < \frac{1}{8\alpha}$, so that collusion is sustainable to begin with and the program has a cartel to deter.

The sustainability condition becomes

$$\frac{1}{1-\delta}\left(\pi^c - \alpha F\right) \;\geq\; \left(\pi^d + R\right) + \frac{\delta}{1-\delta}\,\pi^n.$$

The colluding firm still faces the expected fine αF, but the deviating firm now receives R instead. Substituting and simplifying:

$$\frac{1}{8} - \alpha F \;\geq\; (1-\delta)\left(\frac{1}{4} + R\right).$$

Solving for δ:

$$\delta \;\geq\; \overline{\delta}_L(\alpha, F, R) \;\equiv\; \frac{1 + 8(R + \alpha F)}{2(1 + 4R)}.$$

Comparison with audits. We have $\overline{\delta}_L \geq \overline{\delta}_A$ if and only if

$$\frac{1 + 8(R + \alpha F)}{2(1 + 4R)} \;\geq\; \frac{1}{2(1 - 4\alpha F)}.$$

Cross-multiplying:

$$[1 + 8(R + \alpha F)](1 - 4\alpha F) \;\geq\; 1 + 4R.$$

Expanding and simplifying yields

$$4(R + \alpha F)(1 - 8\alpha F) \;\geq\; 0,$$

which holds since $R + \alpha F > 0$ and, because $\alpha F < \frac{1}{8}$ from the second step, $1 - 8\alpha F > 0$.

Result 3 (leniency program). With a leniency program, collusion is sustainable if and only if $\delta \geq \overline{\delta}_L(\alpha, F, R)$, and $\overline{\delta}_L \geq \overline{\delta}_A$ for every reward $R > -\alpha F$.

Figure 20.2 depicts $\overline{\delta}_A$ and $\overline{\delta}_L$ as functions of α, for $F = 0.12$ and two reward levels, $R = 0.02$ and $R = 0.05$.

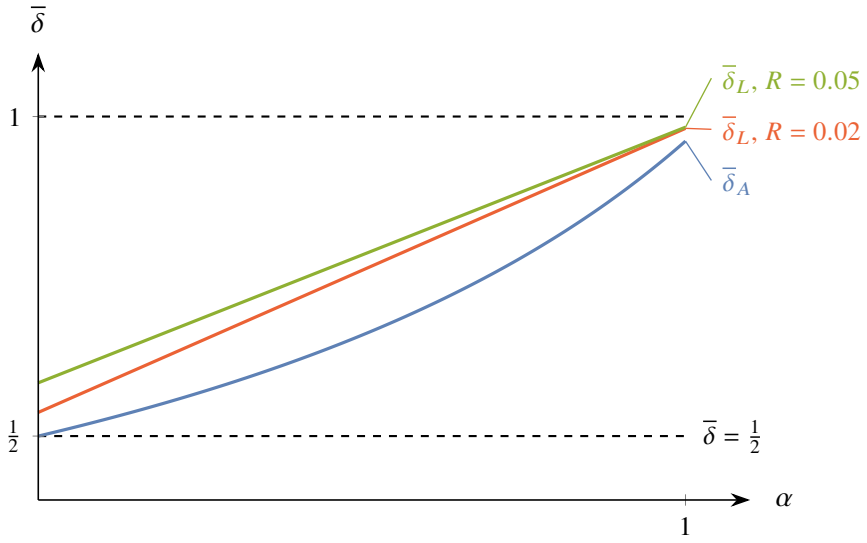


Figure 20.2: Minimal discount factor under audits and under leniency, with $F = 0.12$. The leniency curve lies above the audit curve for all α, and an increase in R shifts it further upward.

Intuitively, the payoff from staying in the cartel is the same as under audits, but a firm that deviates and reports collects R instead of facing the expected fine αF. Since $R > -\alpha F$, reporting raises the payoff from deviating, so leniency programs hinder collusion beyond audits alone.

Comparative statics. Differentiating $\overline{\delta}_L$,

$$\frac{\partial \overline{\delta}_L}{\partial R} = \frac{2(1 - 8\alpha F)}{(1 + 4R)^2}, \qquad \frac{\partial \overline{\delta}_L}{\partial \alpha} = \frac{4F}{1 + 4R} > 0, \qquad \frac{\partial \overline{\delta}_L}{\partial F} = \frac{4\alpha}{1 + 4R} > 0.$$

Since $F < \frac{1}{8\alpha}$ from the second step, we have $1 - 8\alpha F > 0$, so the first derivative is also positive. In addition,

$$1 - \overline{\delta}_L = \frac{1 - 8\alpha F}{2(1 + 4R)} > 0.$$

Result 4 (size of the leniency reward). With a leniency program:

- $\overline{\delta}_L$ increases in R, in α, and in F;
- $\overline{\delta}_L = \overline{\delta}_A$ at $R = -\alpha F$, and $\overline{\delta}_L$ approaches one as R grows but $\overline{\delta}_L < 1$ for every reward R.

Figure 20.3 depicts $\overline{\delta}_L$ as a function of the reward R, for $\alpha = \frac{1}{2}$ and $F = \frac{1}{8}$, so that $\alpha F = \frac{1}{16}$ and $\overline{\delta}_A = \frac{2}{3}$.

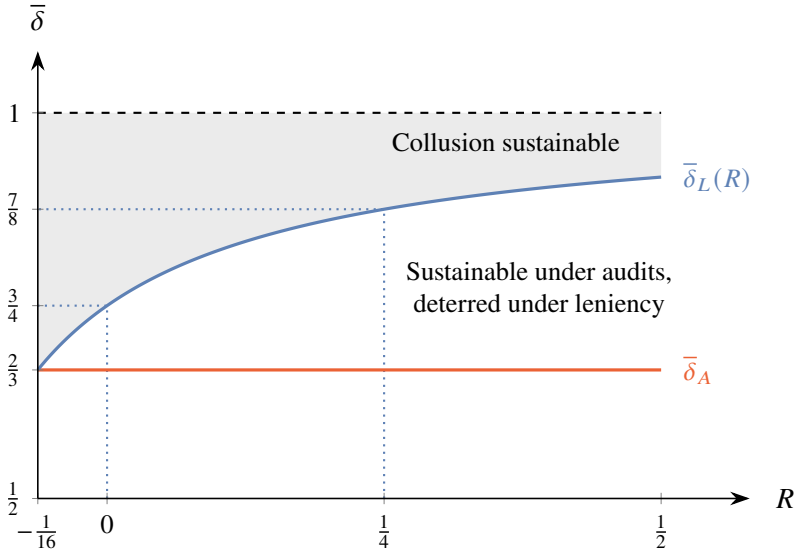


Figure 20.3: Threshold $\overline{\delta}_L$ as a function of the leniency reward R, with $\alpha = \frac{1}{2}$ and $F = \frac{1}{8}$. The shaded area contains the discount factors that sustain collusion under the leniency program.

Intuitively, a larger reward R raises the payoff of a firm that deviates and reports, while the payoff from staying in the cartel is unchanged, so only more patient firms keep colluding. A more severe or more likely fine lowers the payoff from staying in the cartel, which also hinders collusion. At $R = -\alpha F$, the reporting firm receives $-\alpha F$, the same payoff as a deviating firm under audits, and $\overline{\delta}_L$ coincides with $\overline{\delta}_A$. As R grows, $\overline{\delta}_L$ approaches one but never reaches it: for every reward R, firms with $\delta \in [\overline{\delta}_L, 1)$ still collude.

20.6 Fourth Step: Whistleblower Programs

This step extends the analysis following Aubert et al. (2006). That paper distinguishes between leniency programs (which reduce the fine of, or even reward, a cartel *firm* that self-reports) and whistleblower programs (which pay a bounty to *employees* who report the cartel to the CA).

Informed employees. Consider that each firm employs $n \geq 1$ workers who observe whether the firm participates in the cartel and who stay with the firm for one period only. The CA offers a whistleblower bounty $B > 0$ to any employee who reports evidence of collusion. If an employee reports, the CA obtains hard evidence, convicts the cartel, and the employee collects B.

To prevent its n employees from reporting, each colluding firm must pay each worker a "bribe" of at least B per period. The total cost of silencing all employees is therefore nB per period per firm,

as in Aubert et al. (2006). A deviating firm must also pay the bribe in the current period, since its current employees observed the collusive agreement and could still report to the CA. After the cartel collapses (Nash reversion), there is no collusive evidence to report, so no bribes are needed.

Collusion condition. With the whistleblower program in place (alongside audits), per-period collusive profits become $\pi^c - \alpha F - nB$, deviation profits become $\pi^d - \alpha F - nB$, and Nash-reversion profits remain $\pi^n = 0$. The sustainability condition is

$$\frac{1}{1-\delta}\,(\pi^c - \alpha F - nB) \;\geq\; \Big(\pi^d - \alpha F - nB\Big) + \frac{\delta}{1-\delta}\,\pi^n.$$

Substituting and simplifying:

$$\frac{1}{8} - \alpha F - nB \;\geq\; (1-\delta)\left(\frac{1}{4} - \alpha F - nB\right).$$

Solving for δ:

$$\delta \;\geq\; \overline{\delta}_W(\alpha, F, n, B) \;\equiv\; \frac{1}{2\big(1 - 4\alpha F - 4nB\big)}\,.$$

This expression has the same structure as the audit threshold $\overline{\delta}_A$ but with αF replaced by $\alpha F + nB$. It lies below one, so that some $\delta \in (0,1)$ sustains collusion, only when $\alpha F + nB < \frac{1}{8}$.

Comparative statics. Differentiating $\overline{\delta}_W$,

$$\frac{\partial \overline{\delta}_W}{\partial B} = \frac{2n}{\big(1 - 4\alpha F - 4nB\big)^2} > 0, \qquad \frac{\partial \overline{\delta}_W}{\partial n} = \frac{2B}{\big(1 - 4\alpha F - 4nB\big)^2} > 0.$$

Comparison with audits. Since $nB > 0$, the denominator in $\overline{\delta}_W$ is smaller than that in $\overline{\delta}_A$, so $\overline{\delta}_W > \overline{\delta}_A$ whenever $\alpha F + nB < \frac{1}{4}$, and for larger nB no $\delta \in (0,1)$ sustains collusion under the whistleblower program.

Result 5 (whistleblower program). With a whistleblower program and audits:

- if $\alpha F + nB < \frac{1}{8}$, collusion is sustainable if and only if $\delta \geq \overline{\delta}_W(\alpha, F, n, B)$, where $\overline{\delta}_A < \overline{\delta}_W < 1$ and $\overline{\delta}_W$ increases in B and in n;
- if $\alpha F + nB \geq \frac{1}{8}$, no discount factor $\delta \in (0,1)$ sustains collusion.

Intuitively, the program operates through the cost of keeping employees silent: a larger bounty (B) or more informed employees (n) force each firm to spend more on bribes, which erodes the per-period return from collusion. A colluding firm pays these bribes in every period, while a deviating firm pays them only in the period of its deviation, so whistleblower programs provide additional deterrence beyond audits alone.

Comparison with leniency. The ranking between $\overline{\delta}_W$ and $\overline{\delta}_L$ depends on the magnitude of the total bribe nB relative to the leniency reward R. Comparing:

$$\overline{\delta}_W = \frac{1}{2(1 - 4\alpha F - 4nB)} \qquad \text{vs.} \qquad \overline{\delta}_L = \frac{1 + 8(R + \alpha F)}{2(1 + 4R)}\,.$$

When $\alpha F + nB < \frac{1}{4}$, so that both denominators are positive, $\overline{\delta}_W \geq \overline{\delta}_L$ if and only if

$$1 + 4R \;\geq\; [1 + 8(R + \alpha F)](1 - 4\alpha F - 4nB),$$

that is, if and only if

$$nB \;\geq\; \frac{\left(\frac{1}{8} - \alpha F\right)(\alpha F + R)}{\frac{1}{8} + \alpha F + R}.$$

Under the maintained assumption $\alpha F < \frac{1}{8}$, the right-hand side is positive for every $R > -\alpha F$. As a function of R, it equals zero at $R = -\alpha F$, it increases in R, since its derivative with respect to R is

$$\frac{1 - 8\alpha F}{(1 + 8\alpha F + 8R)^2} > 0,$$

and it approaches $\frac{1}{8} - \alpha F$ as R grows.

Result 6 (whistleblower program versus leniency). When $\alpha F + nB < \frac{1}{4}$:

- $\overline{\delta}_W \geq \overline{\delta}_L$ if and only if
$$nB \;\geq\; \frac{\left(\frac{1}{8} - \alpha F\right)(\alpha F + R)}{\frac{1}{8} + \alpha F + R};$$
- this threshold equals zero at $R = -\alpha F$, increases in R, and approaches $\frac{1}{8} - \alpha F$ as R grows.

Figure 20.4 depicts this threshold as a function of the leniency reward R, for $\alpha = \frac{1}{2}$ and $F = \frac{1}{8}$, where it simplifies to $\frac{1+16R}{16(3+16R)}$.

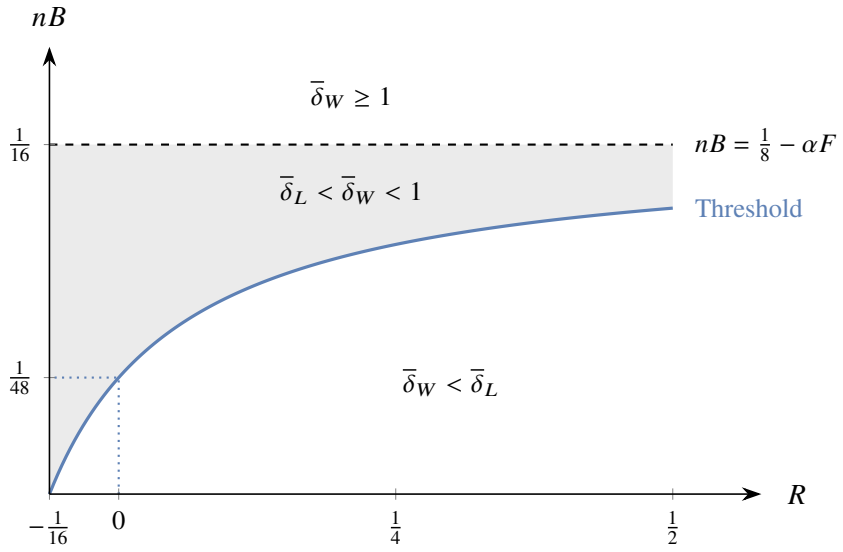


Figure 20.4: Ranking of the whistleblower and leniency programs in the (R, nB) plane, with $\alpha = \frac{1}{2}$ and $F = \frac{1}{8}$. In the shaded area, the whistleblower program deters more cartels than leniency but not all of them.

Intuitively, a more generous leniency reward makes deviation more attractive, so the whistleblower program needs a larger silence cost nB to deter more than leniency does, and the threshold is increasing in R. When $R = -\alpha F$, leniency adds nothing to audits, and any positive nB makes the whistleblower program the stronger tool. As R grows, the threshold approaches $\frac{1}{8} - \alpha F$, the total bribe at which the whistleblower program alone deters collusion for every $\delta \in (0, 1)$.

A numerical illustration. Set $F = 0.12$, $R = 0.04$, $n = 2$, $B = 0.02$ (so $nB = 0.04$). At $\alpha = 0$, the threshold equals

$$\frac{\frac{1}{8}(0.04)}{\frac{1}{8} + 0.04} = \frac{0.005}{0.165} \approx 0.030,$$

which is less than $nB = 0.04$. As α increases, the threshold first rises, reaching its maximum of about 0.034 at $\alpha \approx 0.21$, and then falls (to about 0.003 at $\alpha = 1$), so it stays below $nB = 0.04$. Hence $\overline{\delta}_W > \overline{\delta}_L$ for all α under these parameters.

Full ranking. Combining Results 1 to 6 gives a ranking of the four regimes.

Result 7 (full ranking). When $\alpha F + nB < \frac{1}{4}$ and nB is at least the threshold of Result 6,

$$\overline{\delta}_W \;\geq\; \overline{\delta}_L \;\geq\; \overline{\delta}_A \;\geq\; \overline{\delta} \;=\; \frac{1}{2}.$$

Figure 20.5 depicts the four thresholds as functions of α, for $F = 0.12$, $R = 0.04$, $n = 2$, and $B = 0.02$.

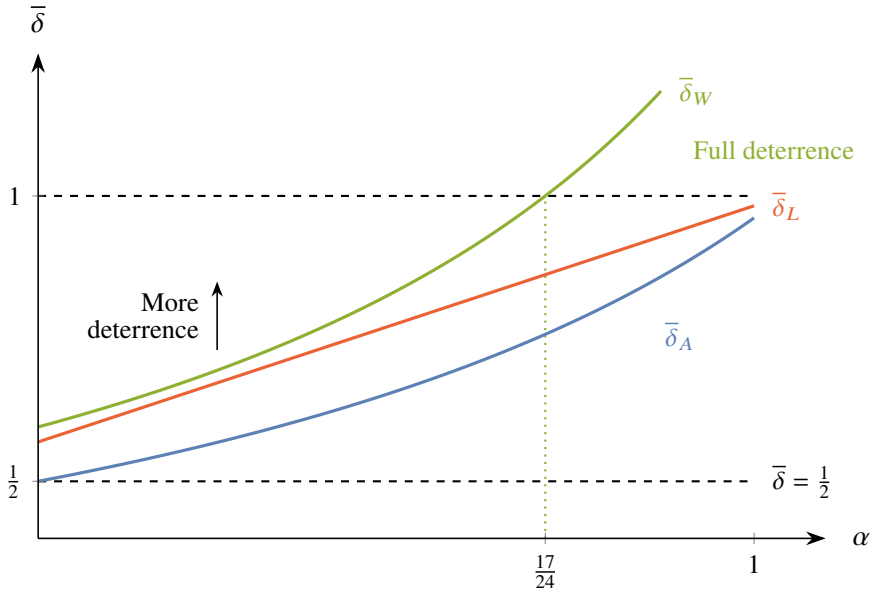


Figure 20.5: Minimal discount factor under the four regimes ($F = 0.12$, $R = 0.04$, $n = 2$, $B = 0.02$). For a given α, collusion is sustained for discount factors above the regime's curve; above $\overline{\delta} = 1$, collusion is deterred for all discount factors. An increase in F, R, n, or B shifts the respective curve upward.

Intuitively, each enforcement tool shrinks the set of discount factors $[\overline{\delta}, 1)$ that sustain collusion. The whistleblower program provides the strongest deterrence (when nB is large enough), followed by leniency, audits, and no enforcement.

Aubert et al. (2006) emphasize that whistleblower programs attack cartel stability from a different margin than leniency: rather than rewarding a firm for betraying its partner, they reward an *individual* for betraying the firm. This forces the cartel to internalize the cost of keeping employees silent (nB per period), which reduces the per-period return from collusion even when the firms themselves would never self-report. The paper stresses that the two instruments are complements, because leniency changes what a deviating firm owes its employees:

- without leniency, a firm that deviates is still exposed to prosecution, so it must keep paying its informed employees for their silence; with employees who stay in the firm forever, the bounty then costs a deviating firm as much as a colluding one, and it does not affect whether collusion can be sustained;
- with leniency, a deviating firm reports the cartel instead of paying its employees, so the bounty raises the cost of colluding but not the cost of deviating; a bounty and a reduced fine can then deter collusion even when neither instrument alone would.

20.7 These Programs in Practice

The leniency and whistleblower programs of this exercise are used by competition authorities around the world. The modern leniency program began in the United States, where the Department of Justice adopted a corporate leniency policy in 1978 and, after it went almost unused, rewrote it in 1993 to grant automatic amnesty to the first firm that reports a cartel (a reward $R = 0$ in the third step). The European Commission issued its first Leniency Notice in 1996, revising it in 2002 and 2006. Adoption then spread to developing and emerging economies: Brazil enacted a leniency program in 2000, South Africa in 2004, and India in 2009, each granting the first cooperating firm full or near-full immunity. The whistleblower bounty B of the fourth step is also real. South Korea's Fair Trade Commission pays informants who report cartels: until 2026 the reward was capped at 3 billion won, and it is now set at up to 10 percent of the fine, and the United Kingdom's Competition and Markets Authority raised its cartel whistleblower reward to £250,000 in 2023.

Two cases show the mechanism at work. In the international vitamins cartel, Rhône-Poulenc (later Aventis) escaped criminal fines in 1999 by being the first member to cooperate under the U.S. program, while its partners F. Hoffmann-La Roche and BASF pleaded guilty and paid $500 million and $225 million. In the air-cargo cartel, Lufthansa's leniency application set off a worldwide investigation that led the European Commission to fine 11 carriers about 799 million euros in 2010.

20.8 Fifth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The deterrence effect of an expected fine: a colluding firm faces it in every period and a deviating firm only once, so audits raise the minimal discount factor.
- The leniency effect: a reward to a firm that reports raises the payoff from deviating and leaves the payoff from colluding unchanged, so larger rewards, positive ones included, deter more cartels than reduced fines.
- The whistleblower effect: a bounty forces the cartel to pay its informed employees for their silence in every period, which lowers the return from collusion, and more so the more employees are informed.

What it drops.

- The communication stage. In the paper, collusion requires communication that leaves hard evidence for one period, which firms or employees can bring to the CA, and an industry whose cartel was reported stays under close scrutiny forever after.
- The combination of leniency and bounties. In the toy model, a deviating firm still pays its employees, so the complementarity between the two instruments, described at the end of the fourth step, is not modeled.
- Implementation issues of large rewards, such as the CA's limited budget, the credibility of its promises, and the public acceptance of paying wrongdoers.
- Side effects of whistleblower programs on firms' behavior, such as deterring desirable cooperation between firms, a more rigid employment structure with longer tenure of informed employees, and changes in firms' investment decisions; and the question of why firms keep incriminating evidence at all.

20.9 Connections to Other Chapters

- Chapter 19 shows that forward contracts can make collusion easier to sustain when the same firms meet again and again.
- Chapter 31 studies how an agency should choose its examination resources when applicants with bad projects respond by making their applications harder to evaluate.
- Chapter 21 studies collusion when each firm privately observes its cost, so that the cartel must give a low-cost firm the incentive to report its cost truthfully.

References

Aubert, C., Rey, P., and Kovacic, W.E. (2006). The impact of leniency and whistle-blowing programs on cartels. *International Journal of Industrial Organization*, 24(6), 1241–1266.

European Commission (2006). *Notice on Immunity from Fines and Reduction of Fines in Cartel Cases*, OJ C 298, 8 December 2006.

U.S. Department of Justice, Antitrust Division (1993). *Corporate Leniency Policy*.

Chapter 21

Optimal Collusion with Private Costs: Market-Share Favors

Companion note to Athey, S., and K. Bagwell (2001), "Optimal Collusion with Private Information," RAND Journal of Economics 32(3): 428–465.

21.1 Introduction

A cartel must agree on more than a high price. It must also decide who produces. Consider two firms that agree to charge a high price. In each period, the cost of each firm is low or high, and only the firm itself knows its cost. A single firm owning both plants would let the low-cost plant serve the market. Can a secret cartel do the same?

The difficulty is that market share is valuable even to a high-cost firm, since the collusive price is well above its cost. A high-cost firm therefore wants to claim low costs to keep its sales. A legal cartel could pay this firm to step aside, but a secret cartel cannot easily make side payments. Without them, the cartel must either lower its price, which hurts both firms, or give up productive efficiency, letting a high-cost firm produce even when its rival has low cost.

Athey and Bagwell (2001) show that promises about future market shares can replace side payments. A firm that admits high costs today receives a larger share of the market in future periods in which both firms have the same cost. Moving sales between two equally efficient firms costs the cartel nothing, so patient firms earn the same profit as a single firm, the first best. Less patient firms give up some productive efficiency before they give up the high price. An efficient cartel thus produces unstable market shares, so stable shares are not a reliable sign of collusion.

The toy model has two firms whose unit cost is zero or one half, with equal probability, and buyers who pay at most one. The first step solves the one-shot game, which serves as the punishment. The second and third steps show why the first best fails without side payments and study rigid pricing. The fourth step shows how market-share favors deliver the first best, the fifth step studies

less patient firms, the sixth step discusses communication and antitrust policy, and the seventh step lists what the toy model keeps and drops.

21.2 The Toy Model

Firms and buyers. Two firms sell a homogeneous good to a unit mass of buyers, each of whom buys one unit if the price is at most $r = 1$. In every period, firm i's unit cost is $\theta_i \in \{\theta_L, \theta_H\} = \{0, \frac{1}{2}\}$, each value with probability $\eta_L = \eta_H = \frac{1}{2}$, independently across firms and over time. Each firm observes its own cost, but not its rival's. We write (j, k) for the state in which firm 1 has cost θ_j and firm 2 has cost θ_k. The paper allows any $r > \theta_H > \theta_L$ and assumes $\eta_L > \frac{1}{2}$ to simplify part of its exposition; equal probabilities are the boundary of that case, and every formula of the paper used below gives the values we derive directly.

The stage game. In each period: (i) each firm learns its cost; (ii) each firm makes a costless, nonbinding announcement, "low" or "high"; (iii) each firm sets a price and proposes a market share. The firm with the lower price serves the whole market. At equal prices, the firms split the market as their proposals indicate. Announcements, prices, and market shares are public; costs are never observed by the rival.

Repeated interaction. The stage game is repeated forever with discount factor $\delta \in (0, 1)$. We report every value as a per-period average, that is, the expected discounted sum of profits multiplied by $1 - \delta$; the paper reports discounted sums. Firms condition only on public information (perfect public equilibria, as in Fudenberg et al., 1994). A firm can deviate in two ways. An *on-schedule* deviation mimics the behavior prescribed for the other cost, for instance announcing a low cost when its cost is high; the rival cannot detect it. An *off-schedule* deviation chooses a price that no cost type is supposed to choose, for instance undercutting the collusive price; the rival detects it, and, as in the paper's computations, firms then play the one-shot equilibrium of the first step forever.

First best. A single firm owning both plants would charge $r = 1$ and produce with the plant that has the lower cost. The expected minimum cost is

$$E[\min\{\theta_1, \theta_2\}] = \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{8},$$

since the minimum is $\frac{1}{2}$ only when both costs are high. The first-best profit is therefore

$$2\pi^{FB} = r - E[\min\{\theta_1, \theta_2\}] = \frac{7}{8} \quad \text{in total, or} \quad \pi^{FB} = \frac{7}{16} \text{ per firm.} \tag{21.1}$$

It requires *pricing efficiency* (price r in every state) and *productive efficiency* (in states (L, H) and (H, L), only the low-cost firm produces).

21.3 First Step: The One-Shot Pricing Game

Without communication or repetition, firms play a Bertrand game with private costs. A high-cost firm prices at its cost, $\frac{1}{2}$. A low-cost firm undercuts this price, but if it always chose the same price,

its low-cost rival would undercut it, so it mixes. Let $F(p)$ be the probability that a low-cost firm charges less than p. A low-cost firm that charges p sells when its rival has high cost (probability $\frac{1}{2}$) or has low cost but charges more (probability $\frac{1}{2}(1 - F(p))$). Indifference with charging just below $\frac{1}{2}$, which earns $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$, requires

$$p\left[\frac{1}{2} + \frac{1}{2}\big(1 - F(p)\big)\right] = \frac{1}{4} \quad \Longleftrightarrow \quad F(p) = 2 - \frac{1}{2p}, \qquad \frac{1}{4} \le p \le \frac{1}{2}. \tag{21.2}$$

Result 1 (one-shot equilibrium). In the one-shot game, there is an equilibrium in which:

- a high-cost firm charges $\frac{1}{2}$ and earns zero;
- a low-cost firm draws its price from $F(p)$ in (21.2) on $[\frac{1}{4}, \frac{1}{2}]$ and earns $\frac{1}{4}$;
- each firm's ex ante expected profit is

$$\pi^{NE} = (\theta_H - \theta_L)\eta_H\eta_L = \frac{1}{8},$$

compared with $\pi^{FB} = \frac{7}{16}$.

Figure 21.1 depicts the price distribution of a low-cost firm. The median price is $\frac{1}{3}$, and no low-cost firm charges more than the high-cost price $\frac{1}{2}$.

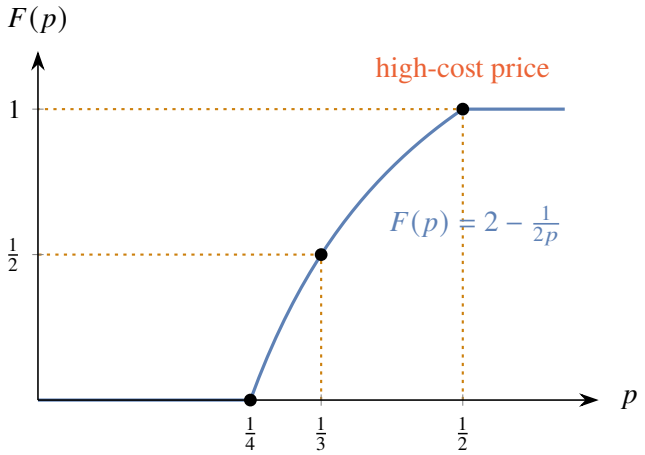


Figure 21.1: Price distribution of a low-cost firm in the one-shot game. A high-cost firm charges $\frac{1}{2}$.

Intuitively, private information softens price competition. A low-cost firm cannot tell whether its rival has high cost, in which case it can charge almost $\frac{1}{2}$ and still sell, or low cost, in which case it must undercut. Each low-cost firm therefore keeps an expected profit of $\frac{1}{4}$, what it earns by pricing just below $\frac{1}{2}$ and selling only to the buyers of a high-cost rival. The one-shot equilibrium is productively efficient, since a low-cost firm always undercuts a high-cost firm, but prices are far below $r = 1$. This equilibrium is the punishment that follows an off-schedule deviation.

21.4 Second Step: First-Best Collusion without Transfers

The first-best scheme. Consider that firms announce their costs, charge $r = 1$, give the whole market to the low-cost firm when costs differ, and split it equally when costs are equal. A high-cost firm that tells the truth sells only if its rival also has high cost, so it expects $\frac{1}{8}$. If it announces a low cost instead, it gets half the market against a low-cost rival and the whole market against a high-cost rival:

$$\underbrace{\frac{1}{2}\cdot\frac{1}{2}\cdot\frac{1}{2}}_{\text{truth}} = \frac{1}{8} < \frac{3}{8} = \underbrace{\frac{1}{2}\cdot\frac{1}{2}\cdot\frac{1}{2}+\frac{1}{2}\cdot 1\cdot\frac{1}{2}}_{\text{announce low cost}}. \tag{21.3}$$

A low-cost firm, by contrast, prefers the truth: it expects $\frac{3}{4}$ by announcing a low cost and $\frac{1}{4}$ by announcing a high cost.

Instruments without transfers. Consider now that what happens after this period cannot reward one firm at the expense of the other, as in a single meeting, or when future values can only be lowered for both firms. The paper's Lemma 4 studies this case, ignoring off-schedule deviations. Two schemes are candidates. In the first, a low-cost firm charges a lower price $\hat{p}$, which makes announcing a low cost less attractive. Truth-telling by a high-cost firm requires

$$\frac{1}{8} = \frac{1}{2}\cdot\frac{1}{2}\left(\hat{p}-\frac{1}{2}\right)+\frac{1}{2}\left(\hat{p}-\frac{1}{2}\right),$$

that is, $\hat{p} = \frac{2}{3}$. Production is efficient, and the total profit is

$$\frac{3}{4}\cdot\frac{2}{3}+\frac{1}{4}\cdot\frac{1}{2} = \frac{5}{8}.$$

In the second scheme, *pooling*, both firms charge 1 and split the market in every state, regardless of their costs. Nobody has a reason to lie, and the total profit is

$$r - E[\theta] = 1 - \frac{1}{4} = \frac{3}{4}.$$

Result 2 (no transfers; Lemma 4 of the paper). When future values cannot be shifted from one firm to the other:

- the first-best scheme is not incentive compatible: a high-cost firm earns $\frac{3}{8}$ by announcing a low cost and $\frac{1}{8}$ by telling the truth;
- the most profitable scheme is pooling, with total profit $\frac{3}{4}$, above the $\frac{5}{8}$ of the productively efficient scheme with $\hat{p} = \frac{2}{3}$ and below the first-best $\frac{7}{8}$.

The comparison depends on the paper's condition (3),

$$\kappa \equiv \frac{r-\theta_H}{\theta_H-\theta_L} > \eta_H,$$

which holds in the toy model since $\kappa = 1 > \frac{1}{2}$. Figure 21.2 depicts the total profit of each scheme as a function of the reservation price r, holding costs and probabilities fixed. Pooling yields $r - \frac{1}{4}$, the first best yields $r - \frac{1}{8}$, and productive efficiency yields

$$\eta_H(r - \theta_H) + \eta_L(1 + \eta_H)(\theta_H - \theta_L) = \frac{r}{2} + \frac{1}{8}.$$

Pooling is better if and only if $r > \frac{3}{4}$, which is condition (3).

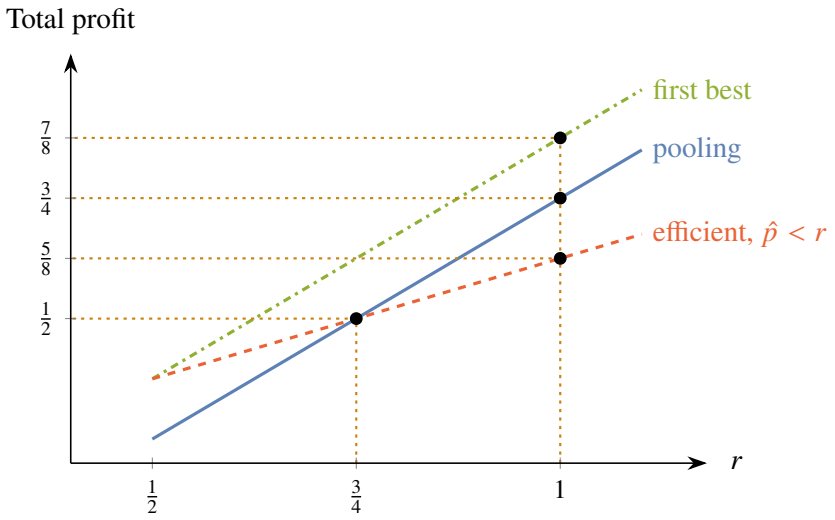


Figure 21.2: Total collusive profit per period without transfers, as a function of the reservation price r ($\theta_L = 0$, $\theta_H = \frac{1}{2}$, equal probabilities).

Intuitively, the first best fails because market share is valuable to a high-cost firm too: at price 1, a high-cost firm earns a margin of $\frac{1}{2}$ per unit. Without a way to compensate the firm that gives up its sales, the only tool that deters false announcements is a lower price for low-cost firms, which reduces the profit of both firms. The gain from productive efficiency, an efficiency rent of $\theta_H - \theta_L = \frac{1}{2}$ in half of the states, is smaller than this cost when the high-cost margin $r - \theta_H$ is large, so the cartel prefers to give up productive efficiency.

21.5 Third Step: Rigid Pricing

Pooling requires no communication: both firms charge 1 and share equally every period, and nothing that happens on the equilibrium path changes future play. Each firm then earns $\frac{3}{8}$ per period. The firm most tempted to undercut is a low-cost firm: undercutting slightly takes the whole market, worth $r - \theta_L = 1$, instead of half of it, and triggers the one-shot equilibrium, worth $\frac{1}{8}$ per period. A low-cost firm does not undercut if

$$(1 - \delta)\frac{1}{2} + \delta\frac{3}{8} \;\geq\; (1 - \delta) \cdot 1 + \delta\frac{1}{8} \quad\Longleftrightarrow\quad \delta \;\geq\; \frac{2}{3}. \tag{21.4}$$

For a high-cost firm, the same comparison, with margin $\frac{1}{2}$ instead of 1, gives $\delta \geq \frac{1}{2}$.

Result 3 (rigid pricing). With reversion to the one-shot equilibrium after an off-schedule deviation, pooling at price 1 is an equilibrium if and only if $\delta \geq \frac{2}{3}$. It gives each firm $\frac{3}{8}$ per period, below $\pi^{FB} = \frac{7}{16}$.

Condition (21.4) is the toy-model version of the paper's condition (5), which requires that a low-cost firm with half of the market at price r does not undercut. Intuitively, rigid pricing is the natural first scheme for a cartel that cannot or will not communicate: market shares do not respond to costs, prices do not move, and the only incentive problem is undercutting. Athey et al. (2004) show that when firms are restricted to symmetric equilibria, so that no firm can be favored over the other, this kind of price rigidity is optimal for patient firms when the cost distribution is log-concave. Its cost is productive inefficiency: half of the time costs differ, and half of the market is then served at cost $\frac{1}{2}$ instead of 0.

21.6 Fourth Step: Market-Share Favors and First-Best Collusion

Continuation values as transfers. Let q^1_{jk} denote firm 1's market share and w^1_{jk} its continuation value (a per-period average) after state (j,k); firm 2 gets $1-q^1_{jk}$ and w^2_{jk}. To earn first-best profits forever, firms must charge 1, set $q^1_{LH}=1$ and $q^1_{HL}=0$, and keep total continuation values at the first best, $w^1_{jk}+w^2_{jk}=\frac{7}{8}$. Only the shares in ties, q^1_{LL} and q^1_{HH}, and the split of future profits remain free. A high-cost firm 1 does not announce a low cost if

$$(1-\delta)\,\frac{1}{2}\left[0+q^1_{HH}\,\frac{1}{2}\right]+\delta\,\frac{w^1_{HL}+w^1_{HH}}{2}\;\geq\;(1-\delta)\,\frac{1}{2}\left[q^1_{LL}\,\frac{1}{2}+\frac{1}{2}\right]+\delta\,\frac{w^1_{LL}+w^1_{LH}}{2}, \tag{21.5}$$

where the brackets list the profits against a low-cost and a high-cost rival. Writing the same condition for firm 2, substituting $w^2_{jk}=\frac{7}{8}-w^1_{jk}$, and adding the two conditions, the shares in ties cancel and we obtain

$$w^1_{HL}-w^1_{LH}\;\geq\;\frac{(1-\delta)(r-\theta_H)}{\delta}=\frac{1-\delta}{2\delta}. \tag{21.6}$$

A firm must expect a better future after admitting high costs against a low-cost rival than after serving the market as the only low-cost firm. This condition, which holds with equality in the paper's Lemma 3, shows that continuation values play the role of the side payments of a legal cartel.

Delivering favors without waste. The cartel can raise firm 1's future value without losing profit only by giving firm 1 more of the market in future ties, states (L,L) and (H,H), where both firms are equally efficient. This source is limited. Let $\bar{v}$ be the highest first-best value firm 1 can obtain in equilibrium. Using (21.5) and the fact that no continuation value exceeds $\bar{v}$, one finds

$$\bar{v}\;\leq\;\frac{1}{8}+\frac{q^1_{LL}}{8}+\frac{q^1_{HH}}{4}\;\leq\;\frac{1}{2}=\frac{r-\theta_L}{2}.$$

By symmetry, no firm's first-best value falls below $\frac{7}{8}-\frac{1}{2}=\frac{3}{8}$. All first-best values thus lie on a

segment of width at most $\frac{1}{8}$, and (21.6) requires

$$\frac{1-\delta}{2\delta} \;\le\; \frac{1}{8},$$

that is, $\delta \ge \frac{4}{5}$.

The scheme. Consider the endpoint at which firm 2 is favored, with values $(\frac{3}{8}, \frac{1}{2})$. Firm 1 gets nothing in ties, $q^1_{LL} = q^1_{HH} = 0$, and its continuation value stays at $\frac{3}{8}$ except after state (H, L), when it rises to $\frac{3}{8} + \frac{1-\delta}{2\delta}$. Firm 1 then serves the market only in state (L, H), and its value is

$$(1-\delta)\,\frac{1}{4} + \delta\left[\frac{3}{8} + \frac{1}{4}\cdot\frac{1-\delta}{2\delta}\right] = \frac{3}{8} \quad \text{for every } \delta, \tag{21.7}$$

so the promise is kept. Condition (21.5) and the corresponding condition for firm 2 hold with equality, and the continuation values stay in $[\frac{3}{8}, \frac{1}{2}]$ if and only if $\delta \ge \frac{4}{5}$. The binding off-schedule temptation is firm 1's in state (L, L), where it has low cost but gets no sales; it does not undercut if

$$\delta\left(\frac{3}{8} - \frac{1}{8}\right) \;\ge\; (1-\delta)\cdot 1,$$

that is, $\delta \ge \frac{4}{5}$. Play starts at the symmetric point $(\frac{7}{16}, \frac{7}{16})$, an equal mixture of the two endpoints: ties are split equally, and after state (H, L) firm 1's value rises to $\frac{7}{16} + \frac{1-\delta}{4\delta}$ while firm 2's falls by the same amount.

Result 4 (first-best collusion; Proposition 1 of the paper). First-best profits, $\pi^{FB} = \frac{7}{16}$ per firm, can be sustained in equilibrium if and only if $\delta \ge \frac{4}{5}$. For $\delta \ge \frac{4}{5}$:

- every equilibrium value used lies on the segment between $(\frac{3}{8}, \frac{1}{2})$ and $(\frac{1}{2}, \frac{3}{8})$;
- a firm that announces high cost against a low-cost rival receives a continuation value higher by $\frac{1-\delta}{2\delta}$ than in the reverse state, as in (21.6);
- the cutoff equals the paper's on-schedule cutoff

$$\delta^{Fon} = \frac{\kappa}{\eta_L^2 + 2\kappa(1-\eta_L)} = \frac{4}{5}$$

  and its bound $\delta^{FB} = \frac{4}{5}$.

The "only if" part holds for any punishment, since it uses only on-schedule constraints; the "if" part uses reversion to the one-shot equilibrium. Figure 21.3 depicts the set of values in Result 4, together with the one-shot payoffs and pooling. At $\delta = \frac{4}{5}$, play moves among three states: the symmetric start S, and the endpoints A (firm 2 favored) and B (firm 1 favored). From S, state (L, H) leads to A and state (H, L) leads to B. At A, firm 2 serves every tie, and play stays at A until firm 1 announces high cost against a low-cost firm 2, which moves play to B.

Intuitively, the cartel replaces side payments with promises about future market shares. A firm that steps aside today is compensated in future ties, where reallocating sales between two equally

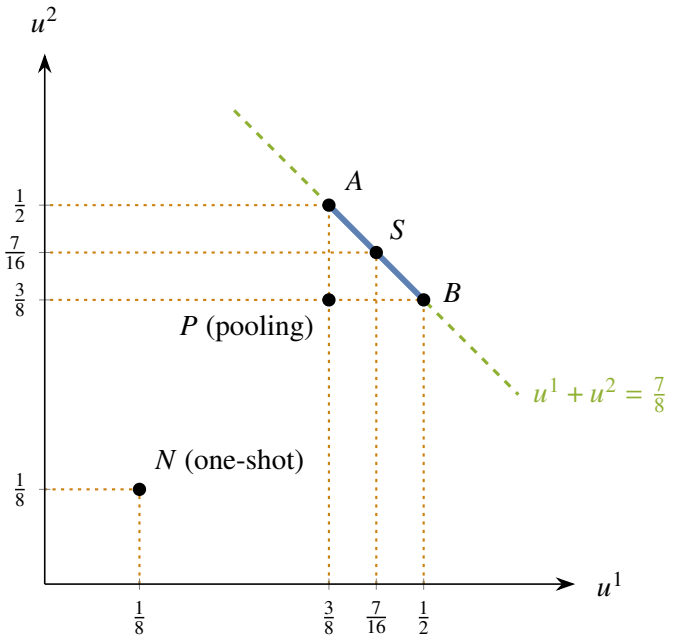


Figure 21.3: Equilibrium values (per-period averages). The thick segment AB supports first-best profits for $\delta \geq \frac{4}{5}$; S is the symmetric start, P pooling, and N the one-shot equilibrium.

efficient firms costs the cartel nothing. Productive efficiency today then coexists with productive efficiency tomorrow, so the cartel earns the first-best profit in every period. The promises are limited, however: a firm can be favored at most in every tie, and a firm that receives no sales in a tie is tempted to undercut. Both limits bind at $\delta = \frac{4}{5}$. The paper's on-schedule cutoff δ^{Fon} increases in κ: the larger the margin of a high-cost firm relative to the efficiency rent, the larger the favor needed to make it admit high cost. The scheme also makes market shares move against past sales: at $\delta = \frac{4}{5}$, a firm that serves the whole market today as the only low-cost firm expects a share of $\frac{1}{4}$ tomorrow, while a firm that admits high cost against a low-cost rival expects a share of $\frac{3}{4}$.

21.7 Fifth Step: Less Patient Firms

For $\frac{2}{3} \leq \delta < \frac{4}{5}$, pooling is sustainable but first best is not. Following the paper's first computation for this case, we look for the most profitable scheme in which firms always charge 1 and the total value is the same, T, in every period, so that equilibrium values lie on a segment $u^1 + u^2 = T$. At the endpoint where firm 2 is favored, let $e = q^1_{LH} - q^1_{HL}$ measure productive efficiency: with $e = 1$, the low-cost firm serves the whole market whenever costs differ; with $e = 0$, as under pooling, the split of the market when costs differ does not depend on which firm has low cost. Shares in ties do not affect total profit, so

$$T = \frac{3}{4} + \frac{e}{8}, \tag{21.8}$$

and at the symmetric start the low-cost firm serves a share $\frac{1+e}{2}$ of the market when costs differ. Maximizing T subject to the on-schedule and off-schedule constraints of both firms, with continu-

ation values on the segment, is a linear program built on the self-generation method of Abreu et al. (1990). Its solution has two regimes, separated by

$$\delta_1 = \frac{2\sqrt{321} - 10}{37} \approx 0.698,$$

with

$$e_A(\delta) = \frac{12\delta(3\delta - 2)}{(4 - \delta)(4 - 5\delta)}, \qquad e_B(\delta) = \frac{12\delta^2}{32 - 12\delta - 23\delta^2}. \tag{21.9}$$

Result 5 (less patient firms). Among schemes with price 1 in every state, total value constant over time, and reversion to the one-shot equilibrium:

- for $\delta < \frac{2}{3}$, no scheme is an equilibrium;
- for $\frac{2}{3} \le \delta < \frac{4}{5}$, the highest total profit is

  $$T^*(\delta) = \frac{3}{4} + \frac{e^*(\delta)}{8}, \qquad \text{where} \qquad e^*(\delta) = \min\{e_A(\delta), e_B(\delta)\}$$

  increases from $e^*(\frac{2}{3}) = 0$ toward $e^*(\frac{4}{5}) = 1$, so $\frac{3}{4} \le T^*(\delta) < \frac{7}{8}$;
- for $\frac{2}{3} \le \delta \le \delta_1$, the disfavored firm 1 serves a share $a(\delta)$ of the market when it announces low cost and $b(\delta)$ when it announces high cost, regardless of its rival's announcement, where

  $$a(\delta) = \frac{8 - 7\delta}{2(4 - \delta)}, \qquad b(\delta) = \frac{32 - 20\delta - 37\delta^2}{2(16 - 24\delta + 5\delta^2)},$$

  and $e^* = a - b = e_A$;
- for $\delta_1 \le \delta < \frac{4}{5}$, the disfavored firm 1 serves nothing when it has high cost, $q^1_{HL} = q^1_{HH} = 0$, a share

  $$q^1_{LL} = \frac{32 - 20\delta - 25\delta^2}{32 - 12\delta - 23\delta^2}$$

  in state (L, L), and only $q^1_{LH} = e_B(\delta) < 1$ when it alone has low cost.

Figure 21.4 depicts $T^*(\delta)$, together with pooling and the first best. At $\delta = \frac{3}{4}$, for instance, $e^* = \frac{108}{161} \approx 0.671$, $T^* \approx 0.834$, and at the symmetric start the low-cost firm serves about 0.835 of the market when costs differ.

Intuitively, the obstacle is the off-schedule constraint of the disfavored firm in state (L, L): a low-cost firm that is assigned few sales undercuts the collusive price unless it values the future enough. When firms are less patient, the cartel must leave this firm a larger share of ties, which uses up the favors available in ties. The rest of the transfer then comes from productive inefficiency: the disfavored firm gives up part of the market even when it alone has low cost (regime $\delta \ge \delta_1$), or market shares respond only partially to costs (regime $\delta \le \delta_1$). Anticipating these future inefficiencies, the cartel also accepts less productive efficiency today. Total profit falls continuously from the first best to the pooling level as δ falls from $\frac{4}{5}$ to $\frac{2}{3}$, while the price stays at 1. The toy model imposes price 1, as the paper's computations do; the paper's second general finding is that, for a wide range of parameters, firms give up productive efficiency before lowering prices, and its numerical example shows that equilibria whose values form a curved frontier, rather than a segment, do slightly better.

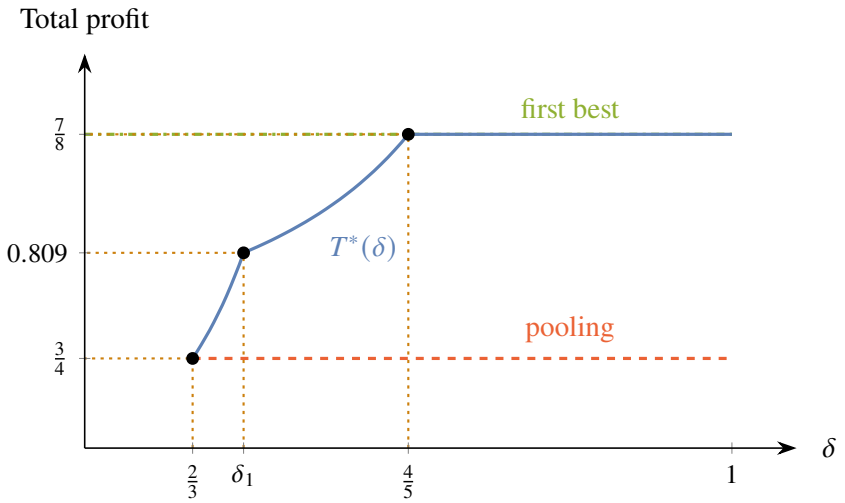


Figure 21.4: Highest total profit per period with price 1 and a constant total value, as a function of the discount factor. Below $\frac{2}{3}$, no such scheme is an equilibrium.

21.8 Sixth Step: Communication and Antitrust Policy

Communication. The scheme of the fourth step can also be run without announcements: firms use slightly different prices for each cost, as in the paper's Proposition 8, so that prices reveal costs after the fact. Without announcements, however, a firm decides whether to undercut before learning its rival's cost. At the endpoint A, a disfavored low-cost firm 1 then expects the whole market with probability $\frac{1}{2}$ (a high-cost rival) and nothing otherwise, and it does not undercut if

$$(1-\delta)\,\frac{1}{2}+\delta\,\frac{3}{8}\;\geq\;(1-\delta)\cdot 1+\delta\,\frac{1}{8}\quad\Longleftrightarrow\quad\delta\;\geq\;\frac{2}{3},\tag{21.10}$$

the same condition as (21.4) for pooling. With announcements, it knows when it faces a tie in which it gets nothing, and the constraint becomes $\delta \geq \frac{4}{5}$, as in the fourth step.

Result 6 (communication; Section 5 of the paper). At the endpoint where firm 1 is disfavored:

- with announcements, firm 1 does not undercut in state (L, L) if and only if $\delta \geq \frac{4}{5}$;
- without announcements, a low-cost firm 1 does not undercut if and only if $\delta \geq \frac{2}{3}$, a high-cost firm 1 if and only if $\delta \geq \frac{1}{2}$, and firm 2 of either cost if and only if $\delta \geq \frac{2}{5}$;
- with or without announcements, first-best profits require $\delta \geq \frac{4}{5}$.

Intuitively, communication helps the cartel allocate the market in rich, state-contingent ways, but it also tells each firm when it is being asked to step aside, which makes undercutting more attractive. In the toy model, the on-schedule constraint and the undercutting constraint with communication bind at the same discount factor, so dropping communication does not lower the cutoff. In the paper's numerical example ($r = 2.5$, $\theta_H = 2$, $\theta_L = 1$, $\eta_L = 0.6$), the cutoff falls from 0.769 with communication to 0.704 without it, and firms may prefer to communicate in some periods but not in others.

Implications. First, market-share favors should appear in the records of cartels that allowed members to exchange cost information without paying each other. The paper cites such arrangements in Japanese construction, the dango system described by McMillan (1991), and in the citric-acid cartel, where a firm that sold above its budget in one year bought from firms below budget in the next. Second, stable market shares are not a reliable sign of collusion: an efficient cartel produces unstable shares, and a firm that sells everything today because it alone has low cost expects to sell less tomorrow. Third, antitrust policy can have unexpected effects. Banning communication, or making side payments risky, may reduce the cartel's productive efficiency without lowering its price: in the paper, bribes that risk detection never fully replace market-share favors (Proposition 9), and bribes that are efficient enough let the cartel produce efficiently (Proposition 10). Such policies raise consumer welfare if they push the critical discount factor above the firms' actual patience, so that collusion at high prices collapses.

21.9 Seventh Step: What the Toy Model Keeps and Drops

Table 21.1 reports the total profit per period and the discount factors of the first to fifth steps.

Table 21.1: Total profit per period and critical discount factors ($\theta_L = 0$, $\theta_H = \frac{1}{2}$, $r = 1$, equal probabilities).

Scheme	Total profit	Discount factors
One-shot equilibrium, the first step	$\frac{1}{4}$	all δ
Productive efficiency with $\hat{p} = \frac{2}{3}$, no transfers, the second step	$\frac{5}{8}$	ignoring undercutting
Pooling at price 1, the third step	$\frac{3}{4}$	$\delta \geq \frac{2}{3}$
Partial productive efficiency, the fifth step	$\frac{3}{4} + \frac{e^*(\delta)}{8}$	$\frac{2}{3} \leq \delta < \frac{4}{5}$
First best with market-share favors, the fourth step	$\frac{7}{8}$	$\delta \geq \frac{4}{5}$

What the toy model keeps.

- The incentive problem of a cartel with private costs: productive efficiency requires a high-cost firm to give up market share, and it does so only if compensated.
- Continuation values as side payments: a firm that admits high cost receives future market-share favors in ties, the paper's Lemma 3.
- A discount factor strictly below one, here $\frac{4}{5}$, above which first-best profits are attained (Proposition 1), and the two limits on favors, the number of ties and the undercutting temptation of the disfavored firm.
- Less patient firms give up productive efficiency, today and in the future, while keeping the price at the reservation value, and the cost of communication through undercutting.

What it drops.

- More than two cost types, more than two firms, and the case $\eta_L > \frac{1}{2}$ assumed in the paper, whose first-best result assumes condition (4), $\eta_L^2 > \kappa(2\eta_L - 1)$; here it reduces to $\frac{1}{4} > 0$.
- The full set of equilibria for less patient firms. The fifth step restricts values to a segment and prices to 1; the paper characterizes the curved Pareto frontier (Propositions 2 to 7) and computes equilibria that do better.
- Harsher punishments than the one-shot equilibrium, which lower the pooling cutoff but not the first-best cutoff $\frac{4}{5}$ of the toy model, and bribes (Section 6 of the paper).
- Persistent costs. Athey and Bagwell (2008) show that when costs are persistent, first-best collusion requires firms to be patient relative to the persistence of costs.

21.10 Connections to Other Chapters

- Chapter 19 studies collusion with complete information, where forward contracts can make the monopoly outcome easier to sustain.
- Chapter 20 studies collusion in prices when costs are known, and the policies that make it harder to sustain: antitrust audits, leniency programs, and whistleblower rewards.
- Chapter 22 studies a regulator that faces firms with private cost information and leaves information rents to the low-cost firms.

References

Abreu, D., Pearce, D., and Stacchetti, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5), 1041–1063.

Athey, S., and Bagwell, K. (2001). Optimal collusion with private information. *RAND Journal of Economics*, 32(3), 428–465.

Athey, S., and Bagwell, K. (2008). Collusion with persistent cost shocks. *Econometrica*, 76(3), 493–540.

Athey, S., Bagwell, K., and Sanchirico, C. (2004). Collusion and price rigidity. *Review of Economic Studies*, 71(2), 317–349.

Fudenberg, D., Levine, D., and Maskin, E. (1994). The folk theorem with imperfect public information. *Econometrica*, 62(5), 997–1039.

McMillan, J. (1991). Dango: Japan's price-fixing conspiracies. *Economics and Politics*, 3(3), 201–218.

Part V

Regulation under Asymmetric Information

The two chapters in this part study a government that intervenes in a market where firms know more than it does. In both chapters, the government must take into account how firms with each kind of private information respond to its intervention, and it pays a price for not knowing what they know. In Chapter 22, the private information is a firm's cost. The regulator leaves an information rent to the low-cost firm and lowers the output of the high-cost firm below the efficient level to reduce this rent; with several firms, competition among them reduces the rents further. In Chapter 23, the private information is the quality of the assets that firms must sell to finance new projects, as in Akerlof's market for lemons. A government buyback removes the worst assets and revives the market, but firms that expect the market to rebound reject a modest offer, so the government always overpays.

The two chapters show the same trade-off in different markets. The regulator of Chapter 22 accepts a distortion in output to limit the rents it leaves to firms, and the government of Chapter 23 accepts paying more than the assets are worth to restore trade.

Chapter 22

Regulating Firms with Unknown Costs: From Monopoly to Oligopoly

Companion note to Baron, D. P., and R. B. Myerson (1982), "Regulating a Monopolist with Unknown Costs," Econometrica 50(4): 911–930, and its oligopoly extension, Wang, G. H. (2000), "Regulating an Oligopoly with Unknown Costs," International Journal of Industrial Organization 18(5): 813–825.

22.1 Introduction

A regulator often sets how much a utility produces and how much it is paid. The regulator, however, rarely knows the utility's costs as well as the utility does. A firm with low costs can then claim high costs to receive a larger payment. How should the regulator choose output and payments, and does it help to regulate several firms rather than one?

Baron and Myerson (1982) answer the first question. To make the low-cost firm report truthfully, the regulator must leave it an information rent, that is, an extra profit that makes lying unattractive. This rent grows with the output asked of the high-cost firm, since a low-cost firm that claims high costs produces that output more cheaply. The regulator therefore keeps the low-cost firm's output at its efficient level ("no distortion at the top") and lowers the high-cost firm's output below the efficient level, giving up some surplus when costs are high to pay a smaller rent when costs are low.

Wang (2000) extends the problem to several firms whose costs are independent. Competition cuts rents: a firm that claims high costs produces only if every rival also reports high costs. When all firms report high costs, the regulator gives production to the firm that is ex ante less efficient, since the firm most likely to have low costs is then left with no rent. As the number of firms grows, the regulator's welfare approaches the level it would reach if it knew that costs are low.

The toy model uses linear demand and two cost types, low and high. The first step solves the complete-information benchmark, the second and third steps solve the monopoly and duopoly problems, the fourth step extends the analysis to many firms, and the fifth step lists what the toy

model keeps and drops.

22.2 The Toy Model

Demand and costs. Inverse demand is $P(q) = 1 - q$, so the gross consumer surplus from output q is $V(q) = q - \frac{q^2}{2}$. A firm's marginal cost is either low, $\theta_l = 0$, with probability p, or high, $\theta_h = \frac{1}{2}$, with probability $1 - p$, and only the firm knows it.

Regulation. By the revelation principle, the regulator asks each firm to report its cost and assigns output and a payment to each report. A firm's rent is its payment minus its production cost. Each firm must be willing to participate (nonnegative rent) and to report truthfully (incentive compatibility). The regulator maximizes

$$W = E\,[V(Q) - \text{payments}] + \alpha\, E[\text{rents}], \tag{22.1}$$

where Q is total output and $\alpha \in [0, 1]$ is the weight on the firms' rents. Wang (2000) sets $\alpha = 0$, which we use from the third step on.

22.3 First Step: Complete Information

If the regulator observes costs, it sets price equal to marginal cost and pays the production cost. Output is $q_l^* = 1$ for the low-cost firm and $q_h^* = \frac{1}{2}$ for the high-cost firm, with surplus

$$V(1) = \frac{1}{2} \qquad \text{and} \qquad V\left(\frac{1}{2}\right) - \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{8}.$$

Expected welfare is

$$\frac{p}{2} + \frac{1 - p}{8},$$

which equals $\frac{7}{32}$ at $p = \frac{1}{4}$.

22.4 Second Step: Monopoly

With one firm, the binding constraints are the high-cost firm's participation constraint and the low-cost firm's incentive constraint. The high-cost firm earns no rent, and the low-cost firm, which could claim high costs and produce q_h at cost θ_l, must earn

$$U_l = (\theta_h - \theta_l)\, q_h = \frac{q_h}{2}. \tag{22.2}$$

Substituting the payments into (22.1), the regulator maximizes

$$p\,[V(q_l) - \theta_l q_l] + (1 - p)\,[V(q_h) - \theta_h q_h] - (1 - \alpha) p\,\frac{q_h}{2},$$

so the high-cost output solves

$$V'(q_h) = \theta_h + (1-\alpha)\frac{p}{1-p}(\theta_h - \theta_l), \tag{22.3}$$

marginal surplus equal to a virtual cost that adds the marginal rent to the true cost.

Result 1 (optimal monopoly regulation, Baron and Myerson, 1982). The regulator:

- asks the low-cost firm to produce the efficient output, $q_l = 1$ (no distortion at the top);
- asks the high-cost firm to produce less than the efficient $\frac{1}{2}$; with $\alpha = 0$,

$$q_h = q^{BM}(p) \equiv \frac{1-2p}{2(1-p)},$$

  and shuts it down if $p \geq \frac{1}{2}$;
- leaves the low-cost firm the rent $\frac{q^{BM}(p)}{2}$ and the high-cost firm none;
- distorts less when it values rents more; with $\alpha = 1$, output is efficient.

Figure 22.1 depicts the high-cost firm's output as a function of p. At $p = \frac{1}{4}$, the high-cost firm produces $\frac{1}{3}$ instead of $\frac{1}{2}$, the low-cost firm earns a rent of $\frac{1}{6}$, and welfare is $W_1\left(\frac{1}{4}\right) = \frac{1}{6}$, so the regulator loses

$$\frac{7}{32} - \frac{1}{6} = \frac{5}{96}$$

from not knowing the cost.

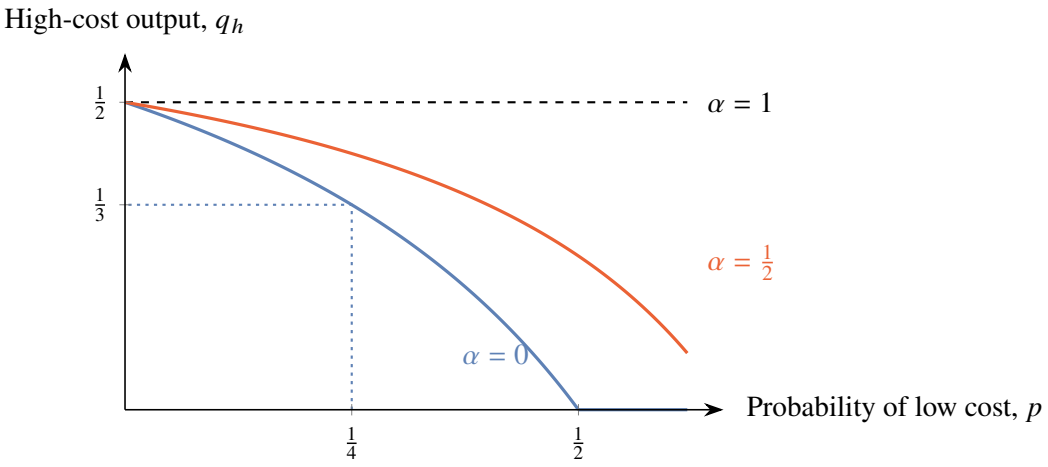


Figure 22.1: Output of the high-cost firm as a function of the probability of low cost, with weight on rents $\alpha = 0$ (blue), $\alpha = \frac{1}{2}$ (red), and $\alpha = 1$ (dashed, efficient output).

Intuitively, every unit the high-cost firm produces raises the rent the low-cost firm can claim by $\theta_h - \theta_l$. This rent is paid whenever the cost is low, which happens with probability p, so the regulator trades a smaller output when costs are high for a smaller rent when costs are low. The more likely the low cost, the more rents weigh, and the larger the distortion.

22.5 Third Step: Duopoly

Consider that two firms produce a homogeneous good, with independent costs and probabilities of low cost $p_1 = \frac{1}{2}$ and $p_2 = \frac{1}{4}$, so firm 1 is ex ante more efficient. Let $\alpha = 0$. As in the monopoly, each high-cost type earns no rent and each low-cost type's rent is $\frac{1}{2}$ times the output it would produce on average if it claimed high costs. The regulator's problem separates into the four cost profiles.

Result 2 (optimal duopoly regulation, Proposition 1 of Wang, 2000). The optimal mechanism:

- if at least one firm reports low cost, lets the low-cost firms produce the efficient output 1 and the high-cost firms nothing;
- if both firms report high costs, lets firm 2, the ex ante less efficient firm, produce $q^{BM}(p_2) = \frac{1}{3}$ and firm 1 nothing;
- yields welfare

$$\begin{aligned} W_2 &= p_1 W_1(1) + (1 - p_1) W_1(p_2) \\ &= \frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot \frac{1}{6} = \frac{1}{3}, \end{aligned}$$

  a convex combination of two monopoly problems.

A brute-force numerical optimization over all quantities and payments that satisfy the incentive and participation constraints confirms this mechanism and the value $\frac{1}{3}$. Favoring firm 1 when both report high costs would yield only

$$p_2 W_1(1) + (1 - p_2) W_1(p_1) = \frac{1}{8} + \frac{3}{4} \cdot \frac{1}{4} = \frac{5}{16},$$

since with $p_1 = \frac{1}{2}$ the Baron-Myerson rule shuts the high-cost firm 1 down.

Figure 22.2 depicts welfare under the two assignments as a function of p_2, with $p_1 = \frac{1}{2}$. Since $W_1(p) = \frac{1}{8(1-p)}$ for $p \leq \frac{1}{2}$, welfare is

$$\frac{1}{4} + \frac{1}{16(1 - p_2)} \text{ when firm 2 produces} \qquad \text{and} \qquad \frac{1 + p_2}{4} \text{ when firm 1 produces.}$$

The two coincide only at $p_2 = p_1$.

Intuitively, competition lets the regulator cut rents in a way a monopoly does not allow. A low-cost firm that claims high costs produces only when the rival also reports high costs. By giving the output in that event to firm 2, the regulator leaves firm 1, the firm most likely to have low costs, with no rent at all, and pays a rent only to firm 2's low-cost type, which is less likely. Moreover, the output distortion depends on the firm's probability of low cost, and $q^{BM}(p)$ is larger for smaller p, so assigning production to firm 2 also produces more when both costs are high.

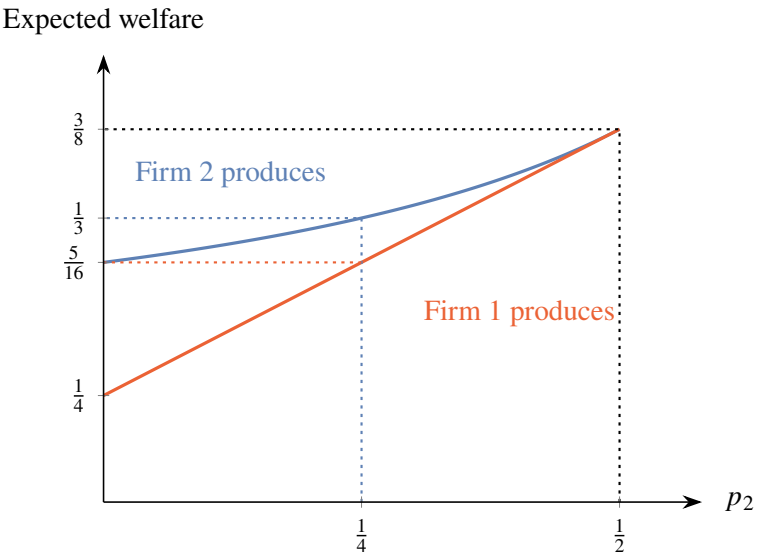


Figure 22.2: Expected welfare when production after two high-cost reports goes to firm 2 (blue) or to firm 1 (red), with $p_1 = \frac{1}{2}$.

22.6 Fourth Step: Extension to Many Firms

With n firms and probabilities $p_1 > p_2 > \cdots > p_n$, the same mechanism applies: firms that report low costs produce the efficient output, and when all report high costs, firm n produces $q^{BM}(p_n)$. When firms are identical, with probability p of low cost, they share $q^{BM}(p)$ when all report high costs, and welfare is

$$W_n(p) = \left[1 - (1-p)^{n-1}\right] W_1(1) + (1-p)^{n-1} W_1(p). \tag{22.4}$$

Result 3 (many firms, Proposition 2 and Corollary 2 of Wang, 2000). With identical firms and $p = \frac{1}{4}$:

- welfare is
$$W_n = \frac{1}{2} - \frac{1}{3}\left(\frac{3}{4}\right)^{n-1},$$
that is, $\frac{1}{6}$, $\frac{1}{4}$, $\frac{5}{16}$, $\frac{23}{64}$, and $\frac{101}{256}$ for one to five firms;
- as n grows, welfare converges to $W_1(1) = \frac{1}{2}$, its level when the regulator knows that costs are low;
- with different probabilities, convergence holds if and only if $\sum_i p_i = \infty$.

Figure 22.3 depicts welfare under the optimal mechanism and under complete information as the number of firms grows. The gap, the cost of not knowing firms' costs, shrinks to zero.

Intuitively, with many firms it becomes almost certain that at least one firm has low costs, so the efficient output is produced. Rents also vanish: a low-cost firm that claims high costs produces only if every other firm also reports high costs, an event whose probability goes to zero.

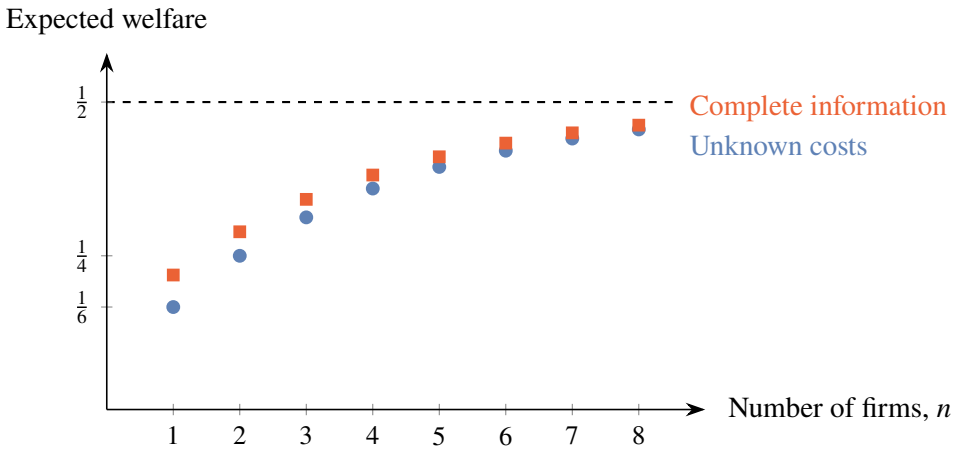


Figure 22.3: Expected welfare with unknown costs (blue) and under complete information (red) as a function of the number of identical firms, with $p = \frac{1}{4}$.

22.7 Fifth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The rent-efficiency trade-off of Baron and Myerson: no distortion at the top, a downward distortion for the high-cost firm, and shutdown when the low cost is likely enough.
- Wang's mechanism, which favors the ex ante less efficient firm, and its welfare as a convex combination of monopoly problems.
- Convergence to the complete-information welfare as the number of firms grows.

What it drops.

- A continuum of cost types and the monotonicity (ironing) issues of Baron and Myerson, as well as unknown fixed costs.
- Convex production costs, under which Wang notes that the regulator may spread output across firms.
- Correlated costs, which let the regulator extract all rents (Crémer and McLean, 1985).

22.8 Connections to Other Chapters

- Chapter 23 studies another government intervention under asymmetric information: a buyback of assets whose quality only the sellers know.
- In Chapter 11, society also commits in advance to a rule that conditions on what the firms reveal, here the evidence of a merger's efficiencies.
- In Chapter 21, the firms themselves must elicit private cost information, and a cartel uses future market shares instead of transfers to do so.

References

Baron, D. P., and Myerson, R. B. (1982). Regulating a monopolist with unknown costs. *Econometrica*, 50(4), 911–930.

Crémer, J., and McLean, R. P. (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2), 345–361.

Wang, G. H. (2000). Regulating an oligopoly with unknown costs. *International Journal of Industrial Organization*, 18(5), 813–825.

Chapter 23

Jump-Starting a Frozen Market: Government Buybacks under Adverse Selection

Companion note to Tirole, J. (2012), "Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning," American Economic Review 102(1): 29–59.

23.1 Introduction

Some markets stop working because sellers know more than buyers. In 2008, the markets for banks' legacy assets, the assets banks already held on their books, froze, and governments offered to buy these assets back. How should a government design such a buyback to restart trade, and how much should it pay?

Tirole (2012) studies this question in an Akerlof (1970) market for lemons. Firms need cash for new projects and must sell legacy assets whose quality only they know. At a low price, only firms with weak assets sell, so buyers are willing to pay little, and small bad news about quality can freeze the market. A government buyback removes the worst assets, which raises the average quality of the assets left in the market and, with it, the market price. Firms anticipate this rebound and reject the government's offer unless it matches the market price. The government thus creates its own competition, as in Coase's (1972) durable-good monopoly.

In contrast with Coase, this competition costs the government nothing, since private buyers break even on what they buy. The optimal intervention buys only the weakest assets, always overpays for them, and leaves the better assets to a market that rebounds. When the market is frozen, a small buyback can restart it, so small bad news can switch the optimal policy from doing nothing (laissez-faire) to intervening.

The toy model assumes that asset quality is uniformly distributed. The first step solves the market without intervention, the second step describes the equilibrium after a buyback offer, the third step

finds the optimal buyback, the fourth step studies market freezes, and the fifth step lists what the toy model keeps and drops.

23.2 The Toy Model

Sellers. A unit mass of cashless firms each own a legacy asset that pays 1 with probability θ and 0 otherwise. A firm knows its θ; outsiders know only that θ is uniform on $[\alpha, 1]$, with $\alpha \in \left[0, \frac{1}{2}\right]$. Each firm also has a new project with net surplus $S = \frac{1}{10}$, which can be financed only if the firm raises at least $k = \frac{3}{10}$ by selling its legacy asset, since the project's own income cannot be pledged to investors.

Trade. A firm that sells its asset at price $p \geq k$ obtains $p + S$; a firm that keeps it obtains θ. Competitive buyers pay the expected value of the assets offered to them. The government can first offer to buy any asset at price p_g; firms that decline can then sell in the market.

Welfare. Public funds have a shadow cost $\lambda > 0$. With buyers breaking even, the government maximizes

$$W = (1 + \lambda)\, E\,[\theta + x(\theta)S] - \lambda\, E[U(\theta)], \tag{23.1}$$

where $x(\theta) = 1$ if type θ's project is financed and $U(\theta)$ is type θ's payoff: total surplus counts fully, and every unit of payoff given to firms costs λ through the deficit.

23.3 First Step: The Market without Intervention

At price p, firms with $\theta \leq p + S$ sell. Buyers break even if p equals the average quality of these assets,

$$p = E[\theta \mid \theta \leq p + S] = \frac{\alpha + p + S}{2}, \qquad \text{so} \qquad p = \alpha + S, \tag{23.2}$$

and types up to $\theta^* = \alpha + 2S$ sell and invest. Trade is possible only if $p \geq k$, that is, if $\alpha \geq \frac{1}{5}$.

Result 1 (laissez-faire, Section II.A of the paper). Without intervention:

- if $\alpha \geq \frac{1}{5}$, assets sell at $p = \alpha + \frac{1}{10}$ and types up to $\alpha + \frac{1}{5}$ invest;
- if $\alpha < \frac{1}{5}$, the market freezes: no asset is sold and no project is financed.

Intuitively, only sellers with assets worth less than $p + S$ sell, so buyers who pay p get assets worth $\alpha + S$ on average. Small bad news, a lower α, lowers the price below the equity needed to invest, and trade collapses entirely, as in Akerlof's market for lemons.

23.4 Second Step: The Market after a Buyback Offer

Consider that the government offers $p_g \geq k$, above the laissez-faire price. If the market paid less than p_g, all types up to $p_g + S$ would sell to the government, and buyers could profitably bid slightly

more than p_g for the assets that remain after the government's purchases, all worth more than $p_g + S$; if it paid more, no firm would sell to the government, and the market would be back at laissez-faire. Therefore, the market price equals p_g.

Result 2 (equilibrium with a buyback, Proposition 1 of the paper). With a buyback at price p_g:

- the market price is p_g, and all types $\theta \le \theta^* = p_g + S$ sell and invest;
- the government buys the worst assets, $\theta < \theta_g = p_g - S$, and the market buys $[\theta_g, \theta^*]$, whose average quality equals p_g;
- the government's deficit equals what it would lose by buying all of $[\alpha, \theta^*]$ itself, since the market breaks even on its part.

Figure 23.1 depicts who sells to whom at $\alpha = \frac{1}{4}$. Without intervention, types up to $\frac{9}{20}$ sell in the market at $\frac{7}{20}$. With a buyback at $p_g = \frac{9}{20}$, the government buys $\left[\frac{1}{4}, \frac{7}{20}\right]$, the market buys $\left[\frac{7}{20}, \frac{11}{20}\right]$, and types above $\frac{11}{20}$ keep their assets.

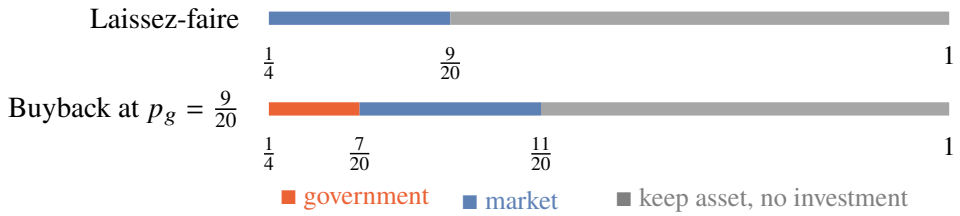


Figure 23.1: Who sells to whom at $\alpha = \frac{1}{4}$: to the government (red), to the market (blue), or nobody (gray), without intervention and with a buyback at $p_g = \frac{9}{20}$.

Intuitively, by removing the worst assets the government raises the quality of the assets left in the market, which raises the market price. Firms anticipate this rebound and accept the government's offer only if it matches the market price, so the government cannot buy at a low price: it creates its own competition.

23.5 Third Step: The Optimal Buyback

The government chooses the cutoff θ^*, or equivalently the price $p_g = \theta^* - S$. Raising θ^* finances one more type, worth $(1 + \lambda)S$ per unit of density, but raises the payoff of every financed type, which costs $\lambda F(\theta^*)$. From (23.1),

$$\frac{\partial W}{\partial \theta^*} = (1 + \lambda)\, S\, f(\theta^*) - \lambda F(\theta^*) = 0 \quad \Longrightarrow \quad \theta^* = \alpha + \frac{(1 + \lambda)S}{\lambda}, \tag{23.3}$$

where F and f are the uniform distribution and density.

Result 3 (optimal intervention, Proposition 2 of the paper). When the market works, $\alpha \ge \frac{1}{5}$:

- the government intervenes if and only if $\lambda < 1$, since the optimal cutoff
$$\alpha + \frac{(1+\lambda)S}{\lambda}$$
exceeds the laissez-faire cutoff $\alpha + 2S$ only then;
- the intervention is never comprehensive: the market keeps the better financed assets;
- the government overpays: it pays p_g for assets worth at most $p_g - S$;
- the presence of the market does not lower welfare relative to a government that could close it but not expropriate firms.

Figure 23.2 depicts welfare as a function of the cutoff, with $\alpha = \frac{1}{4}$ and $\lambda = \frac{1}{2}$. The optimal cutoff is $\frac{11}{20}$, so the government pays $\frac{9}{20}$ for assets worth $\frac{3}{10}$ on average. The deficit is $\frac{1}{50}$, the share of firms that invest rises from $\frac{4}{15}$ to $\frac{2}{5}$, and welfare rises from $\frac{391}{600}$ to $\frac{131}{200}$.

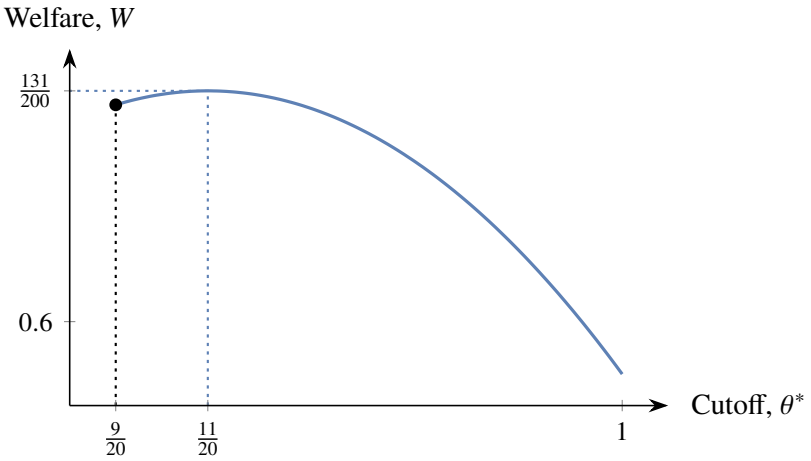


Figure 23.2: Welfare as a function of the cutoff type θ^*, with $\alpha = \frac{1}{4}$ and $\lambda = \frac{1}{2}$. A buyback cannot bring the cutoff below the laissez-faire cutoff $\frac{9}{20}$ (dot), since an offer below the market price attracts no seller; the optimal cutoff is $\frac{11}{20}$.

Intuitively, the government trades off the surplus of the extra projects against the rents it leaves to all financed firms, including the inframarginal ones that would have invested anyway. Because the market breaks even on what it buys, letting the market operate does not change this trade-off; it only changes who buys the better financed assets.

23.6 Fourth Step: Extension to Market Freezes

Let public funds be costly, $\lambda = 2$, so the government stays out of a working market. If $\alpha < \frac{1}{5}$, the market freezes, and the interior cutoff in (23.3) would give a price below k, so any useful intervention must reach the minimum scale $p_g = k$, that is, $\theta^* = \frac{2}{5}$. By (23.1), this intervention raises welfare by
$$F\left(\frac{2}{5}\right)\left(\alpha - \frac{1}{10}\right).$$

Result 4 (market freeze, Proposition 2(iii) of the paper). With $\lambda = 2$:

- if $\alpha \geq \frac{1}{5}$, the government does not intervene;
- if $\frac{1}{10} < \alpha < \frac{1}{5}$, the market is frozen and the government buys the assets in $\left[\alpha, \frac{1}{5}\right]$ at $\frac{3}{10}$, which restarts the market;
- if $\alpha < \frac{1}{10}$, restarting the market costs too much and the government does not intervene.

Figure 23.3 depicts the share of firms that invest, with $\lambda = 2$. Without intervention, the share drops from $\frac{1}{4}$ to zero as α falls below $\frac{1}{5}$. With the optimal intervention, it rises from $\frac{1}{4}$ to $\frac{1}{3}$ as α falls to $\frac{1}{10}$, and drops to zero below. At $\alpha = \frac{3}{20}$, for instance, the government buys the assets in $\left[\frac{3}{20}, \frac{1}{5}\right]$ at a deficit of $\frac{1}{136}$, and $\frac{5}{17}$ of the firms invest.

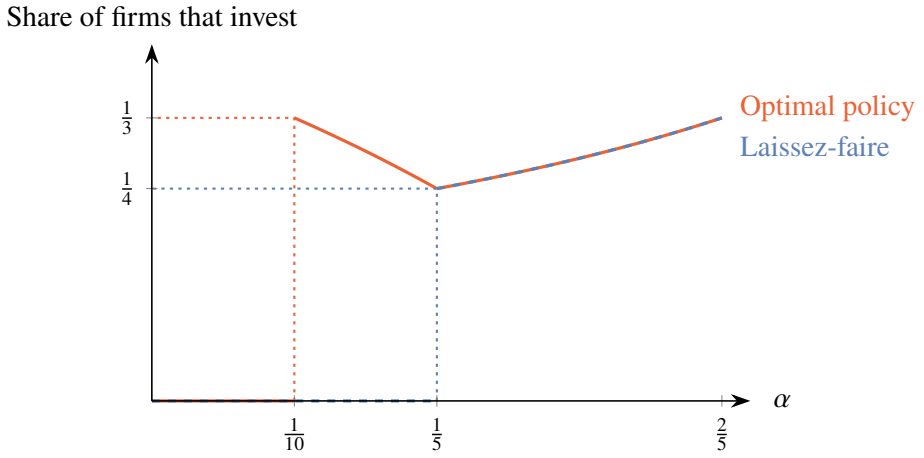


Figure 23.3: Share of firms that invest under laissez-faire (blue, dashed) and under the optimal policy (red), with $\lambda = 2$.

Intuitively, near the freeze threshold, the government only needs to remove a thin slice of bad assets to restore the market: as α approaches $\frac{1}{5}$ the deficit goes to zero while a quarter of the firms invest. Small bad news can thus switch the optimal policy from laissez-faire to intervention, and a small intervention has a large effect.

23.7 Fifth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Market freezes from small bad news, and the market rebound that a buyback creates.
- Price equalization between the government and the market, overpayment, and the irrelevance of the market for welfare.
- The optimal scope of intervention and the switch from laissez-faire to intervention after a freeze.

What it drops.

- Contracts in which firms keep a stake in their legacy assets (Sections III–IV of the paper), which produce a pecking order: buybacks for the weakest assets, recapitalization for intermediate ones, and the market for the strongest.
- Ex ante moral hazard: the prospect of an intervention reduces firms' incentive to build high-quality assets.
- General log-concave distributions of asset quality.

23.8 Connections to Other Chapters

- Chapter 22 studies a regulator that leaves information rents to firms with private information about their costs.
- Chapter 29 studies another market where buyers cannot tell good products from bad ones: firms selling under a common eco-label share its reputation, and some of them pass off conventional products as green.

References

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.

Coase, R. H. (1972). Durability and monopoly. *Journal of Law and Economics*, 15(1), 143–149.

Tirole, J. (2012). Overcoming adverse selection: How public intervention can restore market functioning. *American Economic Review*, 102(1), 29–59.

Part VI

Signaling, Disclosure, and Certification

The six chapters in this part study markets where consumers cannot observe an attribute of a product, such as its quality or how clean it is, and ask how information about it reaches them. In the first three chapters the firm conveys the information itself; in the last three a certifier does. Chapter 24 compares two channels: disclosing quality at a cost, or signaling it with a price that a firm of lower quality would not want to copy. Chapters 25 and 26 apply price signaling to green products. In Chapter 25, a clean firm signals in a single period, and, when cleaner products are more costly to produce, it charges a price above its full-information level. In Chapter 26, consumers learn the truth only after a verification lag, and a responsible entrant signals with a low introductory price when it faces conventional incumbents and with a high one in a market that no firm serves yet.

Chapters 27 through 29 turn to eco-labels, and each studies a different reason why a label is an imperfect signal. In Chapter 27, consumers do not know how demanding the label's standard is, so a label on a product with a poor reputation suggests that the standard is easy to meet. In Chapter 28, consumers know the standard, but the certification test makes errors. In Chapter 29, firms that sell under a common label share its reputation, and some of them pass off conventional products as green, so fraud rather than test error places brown products under the label. Across the part, the question is what consumers can infer from a price or a label when they cannot verify the claim behind it.

Chapter 24

Communicating Quality: Disclosure or Price Signaling

Companion note to Daughety, A. F., and J. F. Reinganum (2008), "Communicating Quality: A Unified Model of Disclosure and Signalling," RAND Journal of Economics 39(4): 973–989.

24.1 Introduction

Consumers often cannot see a product's quality before they buy it. Consider a winery that knows its new wine is excellent. It can pay an independent expert to certify the quality, which is *disclosure*, or it can charge a high price that only a high-quality producer would find worthwhile, which is *signaling*. Which channel does the firm choose, and does it choose the one that society would prefer?

Earlier work studied the two channels separately. Daughety and Reinganum (2008) combine them in a model where higher quality is costlier to produce, so a firm that does not disclose can still reveal its type through its price. A high price cuts the quantity sold, and this loss hurts a low-quality firm more, because its lower cost gives it a larger margin on each unit. To deter mimicking, a high-quality firm must raise its price above its full-information level and sell less. Disclosure, in contrast, has a fixed cost, so the firm picks the cheaper channel.

The paper finds that the lowest type signals without distortion, that higher types distort their price more, and that types above a cutoff disclose. Profit is therefore U-shaped in quality. Voluntary disclosure is socially *insufficient*: disclosure also lowers prices and raises output, a gain for consumers that the firm ignores. This reverses the classical result, in which quality does not affect cost and voluntary disclosure is excessive. A targeted subsidy for disclosure can restore efficiency.

The toy model uses linear demand and a uniform prior on quality. The first step solves the full-information benchmark, which is the outcome under disclosure, and the second step solves the separating signaling equilibrium. The third and fourth steps measure the cost of signaling and

solve the firm's choice of channel, the fifth step compares it with the social optimum, the sixth step studies the classical model where quality does not affect cost, and the seventh step lists what the toy model keeps and drops.

24.2 The Toy Model

Quality. A monopolist produces a good of quality $\theta \in [\underline{\theta}, \overline{\theta}] = [0, 1]$, which it observes and consumers do not. Quality is the probability that a unit is completely satisfactory, and consumers' prior is uniform on the unit interval. The sixth step uses the same type space and prior.

Consumers. A consumer's utility from q units is

$$\big(\alpha - (1 - \widetilde{\theta})\delta\big)q - \frac{\beta}{2}\,q^2 + I - pq,$$

where $\widetilde{\theta}$ is the perceived quality, p is the price, and I is income. A satisfactory unit yields a linear coefficient α and a disappointing unit yields $\alpha - \delta$, so a higher perceived quality shifts demand out, and demand is

$$q(p, \widetilde{\theta}) = \frac{A(\widetilde{\theta}) - p}{\beta}, \qquad A(\widetilde{\theta}) \equiv \alpha - (1 - \widetilde{\theta})\,\delta. \tag{24.1}$$

Perceived quality equals θ if the firm discloses. Otherwise consumers form it from the price, by a belief function $B(p)$ that is part of a perfect Bayesian equilibrium.

The firm. A firm of type θ produces at constant marginal cost $k\theta$, so higher quality is costlier, and earns

$$\pi(p, \theta, \widetilde{\theta}) = \frac{(p - k\theta)\big(A(\widetilde{\theta}) - p\big)}{\beta}. \tag{24.2}$$

Once it knows θ, the firm chooses its price and whether to disclose, which costs $D > 0$ (for example, an audit). Disclosure is truthful, so the perceived quality is θ.

Parameter values. $\alpha = \frac{5}{2}$, $\beta = 1$, $\delta = 1$, and $k = \frac{1}{2}$, so $A(\widetilde{\theta}) = \frac{3}{2} + \widetilde{\theta}$ and $q = \frac{3}{2} + \widetilde{\theta} - p$. The choices $\beta = \delta = 1$ are normalizations. The paper assumes $0 < \underline{\theta} < \overline{\theta} < 1$, so the unit type space is a limiting case, which the paper itself uses in its Section 4. The value $\alpha = \frac{5}{2}$ keeps the intercept $A(\widetilde{\theta}) = \frac{3}{2} + \widetilde{\theta}$ simple while satisfying the paper's assumptions: quality is socially valuable, $\delta > k$ (Assumption 1); every type can earn a positive margin at some price (Assumption 2),

$$A(\underline{\theta}) = \frac{3}{2} > k\overline{\theta} = \frac{1}{2};$$

and a slightly stronger version used in the proofs (Assumption A3),

$$A(\underline{\theta}) = \frac{3}{2} > k\overline{\theta} + k(\overline{\theta} - \underline{\theta}) = 1.$$

the sixth step reuses α, β, and δ with $k = 0$. Figure 24.1 depicts demand (24.1) for three perceived qualities: each unit of perceived quality raises the intercept $A(\widetilde{\theta})$ by $\delta = 1$, so the same price sells more when consumers believe quality is higher, which is why the firm cares about perceptions.

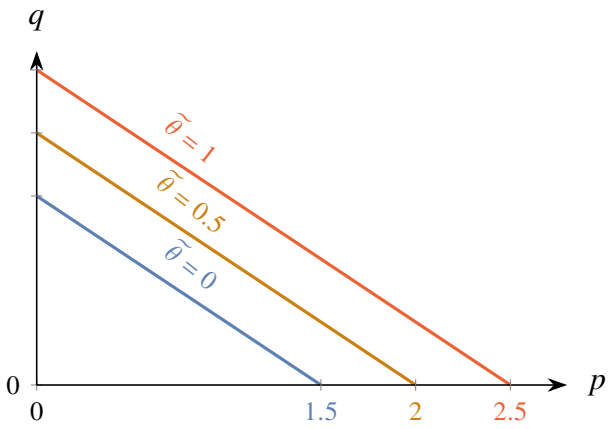


Figure 24.1: Demand $q = \frac{3}{2} + \widetilde{\theta} - p$ for three perceived qualities; the intercepts on the price axis are $A(\widetilde{\theta}) = \frac{3}{2} + \widetilde{\theta}$.

24.3 First Step: Full Information

Consider that consumers observe θ. The firm maximizes (24.2) with $\widetilde{\theta} = \theta$, so

$$P^f(\theta) = \frac{A(\theta) + k\theta}{2}, \qquad \pi^f(\theta) = \frac{\big(A(\theta) - k\theta\big)^2}{4\beta}, \qquad CS^f(\theta) = \frac{\big(A(\theta) - P^f(\theta)\big)^2}{2\beta}, \tag{24.3}$$

the price, profit, and consumer surplus. In the toy model,

$$P^f(\theta) = \frac{3}{4}(1+\theta), \qquad \pi^f(\theta) = \frac{(3+\theta)^2}{16}, \qquad \text{and} \qquad CS^f(\theta) = \frac{(3+\theta)^2}{32}.$$

All three rise with quality: a unit of quality raises the demand intercept by δ and marginal cost by $k < \delta$. A firm that discloses pays D and charges $P^f(\theta)$, earning $\pi^f(\theta) - D$. At the two ends of the interval, $P^f = 0.75$ and $\pi^f = 0.5625$ at $\theta = 0$, and $P^f = 1.5$ and $\pi^f = 1$ at $\theta = 1$.

24.4 Second Step: The Separating Signaling Equilibrium

Now consider that disclosure is too costly to be used, and that consumers infer quality from the price. A *separating* equilibrium is a strictly increasing price function $P^s(\theta)$ and beliefs $B^*(p)$ such that each type prefers its price to any other,

$$\pi\big(P^s(\theta), \theta, \theta\big) \geq \max_p \pi\big(p, \theta, B^*(p)\big),$$

and beliefs are correct, $B^*\big(P^s(\theta)\big) = \theta$.

The differential equation. A type that charges p is perceived as $B(p)$ and earns $\pi\big(p, \theta, B(p)\big)$. The first-order condition in p is

$$\big(A(B(p)) - p\big) + (p - k\theta)\big(\delta B'(p) - 1\big) = 0.$$

In equilibrium $B(P^s(\theta)) = \theta$ and $B'(p) = \frac{1}{P^{s\prime}(\theta)}$, so P^s solves

$$\frac{dp}{d\theta} = \frac{\delta\,(p - k\theta)}{2p - A(\theta) - k\theta} = \frac{p - \frac{1}{2}\theta}{2p - \frac{3}{2} - \frac{3}{2}\theta} \quad \text{in the toy model.} \tag{24.4}$$

The numerator is positive. The denominator is positive if and only if $p > P^f(\theta)$, since $2P^f = A + k\theta$. So the equilibrium price rises with quality, has a slope above δ (because $p - k\theta > 2p - A - k\theta$ when $p < A$), and lies above the full-information price wherever it is defined. The lowest type has no one below it to deter, so it has no reason to distort: the initial condition is $P^s(\underline{\theta}) = P^f(\underline{\theta}) = 0.75$, where the denominator is zero and the slope is infinite.

The solution. The paper solves (24.4) implicitly:

$$\big(A(\theta) - p\big)^{\delta - k}\left(2p - k\theta + \frac{k(\alpha - \delta)}{\delta - k}\right)^{\delta} = K, \qquad \text{that is,} \qquad \sqrt{\tfrac{3}{2} + \theta - p}\,\left(2p - \tfrac{1}{2}\theta + \tfrac{3}{2}\right) = K \tag{24.5}$$

in the toy model. Differentiating the logarithm of the left side along (24.4) gives zero, so the left side is constant along the solution, and K is set by the boundary condition at $\theta = 0$:

$$K = \sqrt{0.75} \cdot 3 = \frac{3}{2}\sqrt{3} = 2.598.$$

Solving for p at each θ gives the schedule of Table 24.1.

Result 1 (Proposition 1 of the paper). There is a unique separating perfect Bayesian equilibrium:

- the lowest type charges its full-information price, $P^s(\underline{\theta}) = P^f(\underline{\theta})$;
- every higher type charges more than its full-information price, and P^s is strictly increasing and concave, given by (24.5);
- beliefs are $B^*(p) = (P^s)^{-1}(p)$ for prices on the schedule, $\underline{\theta}$ below it, and $\overline{\theta}$ above it.

Figure 24.2 depicts the two price functions against quality, with $\delta = 1$ and $k = \frac{1}{2}$. They coincide at $\theta = 0$, the signaling price rises steeply at first and then flattens (concave), and the gap, which is the price distortion, widens with quality.

Intuitively, a higher price is a costly signal because it cuts the quantity sold, and a lower-quality type, which has the lower cost $k\theta$, would gain from being taken for a higher type. To deter that, the price schedule must climb faster than the value of quality to consumers, δ, and the distortion accumulates with quality. In the toy model, $P^s(1) = 2.282$ against $P^f(1) = 1.5$, and quantity at the top falls from 1 to $A - P^s = 0.218$.

The equilibrium price is the best reply of each type to the beliefs. Figure 24.3 depicts, for the type $\theta = 0.5$, its profit against the price it charges. The blue curve is the profit if consumers perceived the true quality at every price, which peaks at the full-information price $P^f(0.5) = 1.1250$ with height $\pi^f = 0.7656$. The red curve is the profit under the equilibrium beliefs $B^*(p)$: consumers infer a higher quality from a higher price, which raises demand along the curve. The red curve lies below the blue one at lower prices, where consumers underestimate quality, and above it at higher prices, where they overestimate it; the two cross at $P^s(0.5) = 1.6833$, where beliefs are correct, and the red curve peaks there with height $\pi^s = 0.4539$. Under these beliefs the type does not gain from any other price.

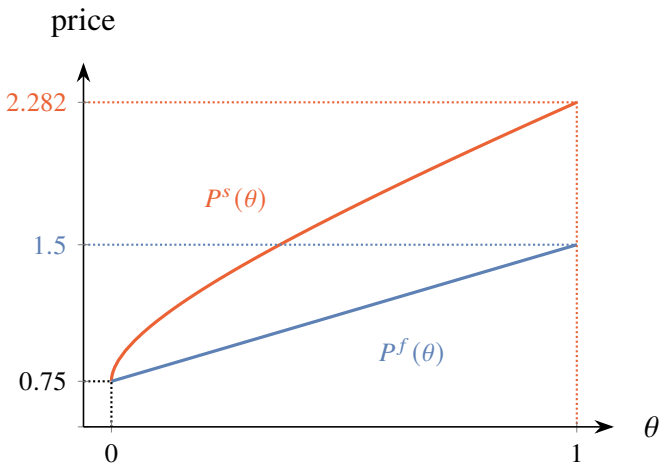


Figure 24.2: Signaling price $P^s(\theta)$ and full-information price $P^f(\theta) = \frac{3}{4}(1+\theta)$ against quality, with $\delta = 1$ and $k = \frac{1}{2}$.

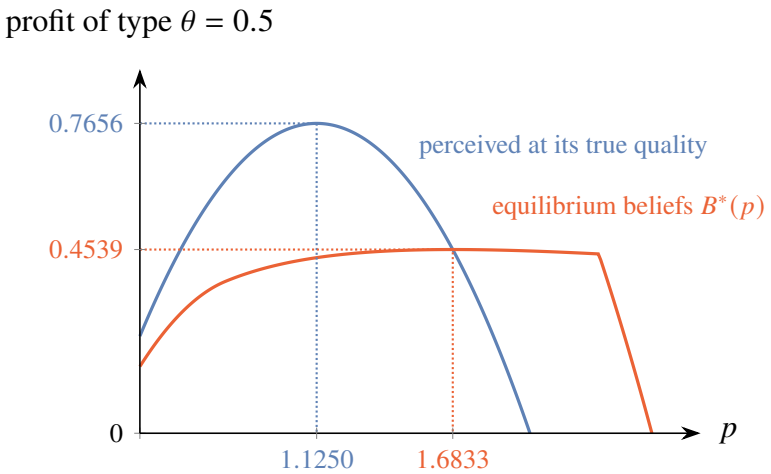


Figure 24.3: Profit of type $\theta = 0.5$ against its price, if perceived at its true quality at every price (blue) and under the equilibrium beliefs (red).

24.5 Third Step: The Cost of Signaling

Substituting the equilibrium price, the profit of type θ under signaling is

$$\pi^s(\theta) = \frac{(P^s - k\theta)(A - P^s)}{\beta}.$$

Beliefs are correct in equilibrium and the first-order condition holds, so by the envelope theorem only the direct effect of quality on cost remains,

$$\frac{d\pi^s}{d\theta} = -\frac{k\left(A(\theta) - P^s(\theta)\right)}{\beta} = -k\, q^s(\theta) < 0, \tag{24.6}$$

while the full-information profit rises,

$$\frac{d\pi^f}{d\theta} = \frac{(\delta - k)(A - k\theta)}{2\beta} > 0.$$

Signaling profit *falls* with quality because of the price distortion, and the two coincide at the lowest type, $\pi^s(\underline{\theta}) = \pi^f(\underline{\theta})$. The profit loss from having to signal, $\Delta\pi \equiv \pi^f - \pi^s$, is therefore strictly increasing in quality,

$$\frac{d\Delta\pi}{d\theta} = \frac{(\delta - k)(A - k\theta)}{2\beta} + k\, q^s(\theta) > 0. \tag{24.7}$$

Result 2 (Proposition 2 of the paper). The difference $\pi^f(\theta) - \pi^s(\theta)$ is strictly increasing in θ, from zero at the lowest type.

Table 24.1 reports the two prices and the two profits. The loss rises from 0 at $\theta = 0$ to 0.6115 at $\theta = 1$, and signaling profit falls from 0.5625 to 0.3885 while full-information profit rises from 0.5625 to 1.

θ	P^f	P^s	π^f	π^s	$\Delta\pi = \pi^f - \pi^s$
0.0	0.7500	0.7500	0.5625	0.5625	0.0000
0.2	0.9000	1.2613	0.6400	0.5094	0.1306
0.4	1.0500	1.5520	0.7225	0.4705	0.2520
0.6	1.2000	1.8093	0.8100	0.4388	0.3712
0.8	1.3500	2.0505	0.9025	0.4119	0.4906
1.0	1.5000	2.2820	1.0000	0.3885	0.6115

Table 24.1: Prices and profits under full information (f) and under signaling (s), with $\alpha = \frac{5}{2}$, $\beta = 1$, $\delta = 1$, and $k = \frac{1}{2}$.

24.6 Fourth Step: The Choice between Signaling and Disclosure

A type strictly prefers to disclose if and only if $\pi^f(\theta) - D > \pi^s(\theta)$, that is, if and only if $\Delta\pi(\theta) > D$. By Result 2, there are three cases:

- if $D \geq \Delta\pi(\overline{\theta}) = 0.6115$, every type signals;
- in the limit $D = 0$, every type discloses;
- if $0 < D < \Delta\pi(\overline{\theta})$, there is a marginal type θ^V with $\Delta\pi(\theta^V) = D$: types below it signal and types above it disclose. The marginal type is indifferent and, as in the paper, is assumed to disclose.

A nondisclosing price above $P^s(\theta^V)$ is off the equilibrium path, and consumers treat it as coming from the type $B^*(p)$; the incentive constraints then guarantee that no disclosing type gains by signaling instead. The firm's overall profit is

$$\Pi(\theta) = \begin{cases} \pi^s(\theta) & \text{if } \theta < \theta^V, \\ \pi^f(\theta) - D & \text{if } \theta \geq \theta^V, \end{cases} \tag{24.8}$$

which is continuous. It falls on $[\underline{\theta}, \theta^V]$, because π^s falls, and rises on $[\theta^V, \overline{\theta}]$, because π^f rises: profit is *U-shaped* in quality. The highest type earns more than the lowest if

$$\pi^f(\overline{\theta}) - D > \pi^s(\underline{\theta}) = \pi^f(\underline{\theta}),$$

that is, if

$$D < D^{\text{mod}} \equiv \pi^f(\overline{\theta}) - \pi^f(\underline{\theta}),$$

a *moderate* disclosure cost.

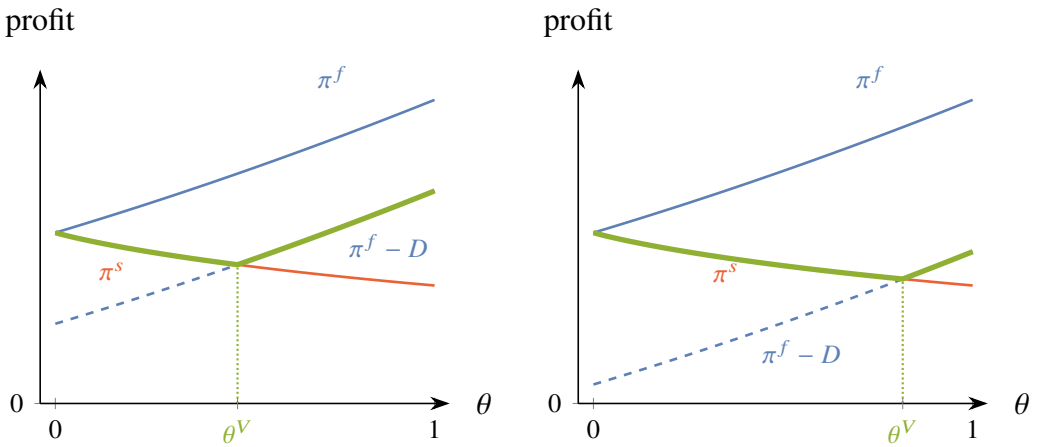
Result 3 (voluntary disclosure and U-shaped profit). For $0 < D < \Delta\pi(\overline{\theta})$:

- types below θ^V signal and types above it disclose, where $\Delta\pi(\theta^V) = D$;
- overall profit $\Pi(\theta)$ is U-shaped in θ, with its minimum at θ^V;
- if $D < D^{\text{mod}}$, the highest type earns the most; if $D^{\text{mod}} < D < \Delta\pi(\overline{\theta})$, the lowest type does.

Figure 24.4 depicts the profits for $D = 0.3$ in panel (a) and $D = 0.5$ in panel (b). The red curve is signaling profit π^s and the blue curves are full-information profit π^f (solid) and net of the disclosure cost (dashed). The firm earns the higher of the red curve and the dashed blue curve, shown in green, and the two cross at θ^V. In panel (a) the right end of the green curve is above its left end, and in panel (b) it is below.

In the toy model, $D^{\text{mod}} = 1 - 0.5625 = 0.4375$. For $D = 0.3$ (moderate), $\theta^V = 0.480$, the minimum profit is $\Pi(\theta^V) = 0.4571$, and

$$\Pi(\overline{\theta}) = 0.7000 > \Pi(\underline{\theta}) = 0.5625.$$



(a) $D = 0.3$ (b) $D = 0.5$

Figure 24.4: Full-information profit π^f, its value net of the disclosure cost, signaling profit π^s, and the firm's overall profit Π (green) against quality.

For $D = 0.5$ (between D^{mod} and 0.6115), $\theta^V = 0.816$, the minimum is 0.4099, and $\Pi(\overline{\theta}) = 0.5000 < \Pi(\underline{\theta})$. Disclosure at moderate cost partly resolves a problem of many price-signaling models: there, higher-quality firms earn less, so firms would want to produce low quality, whereas here the highest type earns the most.

24.7 Fifth Step: Social Efficiency, Mandatory Disclosure, and a Subsidy

The firm weighs only its own profit when it chooses its channel. Disclosure also lowers the price from the signaling schedule to the full-information line and raises output, which benefits consumers. Define consumer surplus under signaling and the surplus gain from disclosure as

$$CS^s(\theta) = \frac{(A - P^s)^2}{2\beta} \qquad \text{and} \qquad \Delta CS(\theta) \equiv CS^f(\theta) - CS^s(\theta),$$

where ΔCS is strictly increasing in θ (Proposition 3 of the paper). Table 24.2 reports the numbers. Let $W(\hat{\theta})$ be total surplus when types below $\hat{\theta}$ signal and types above it disclose,

$$W(\hat{\theta}) = \int_{\underline{\theta}}^{\hat{\theta}} \left(CS^s + \pi^s\right) dG(\theta) + \int_{\hat{\theta}}^{\overline{\theta}} \left(CS^f + \pi^f - D\right) dG(\theta), \tag{24.9}$$

so that

$$W'(\hat{\theta}) = \left[D - \Delta\pi(\hat{\theta}) - \Delta CS(\hat{\theta})\right] g(\hat{\theta}). \tag{24.10}$$

The bracket equals $D > 0$ at $\underline{\theta}$ and decreases in $\hat{\theta}$, by Result 2 and Proposition 3. When $0 < D < \Delta\pi(\overline{\theta})$, the case the paper studies, it equals $-\Delta CS(\theta^V) < 0$ at θ^V, so surplus is maximized at the unique θ^W with $\Delta\pi(\theta^W) + \Delta CS(\theta^W) = D$. The firm's own cutoff solves $\Delta\pi(\theta^V) = D$. Since $\Delta CS > 0$, $\theta^W < \theta^V$.

Result 4 (Section 3 of the paper).

- Voluntary disclosure is socially insufficient: types in $[\theta^W, \theta^V)$ signal although surplus would be higher if they disclosed.
- Mandatory disclosure lowers prices and raises output, but it makes types below θ^W pay D although surplus is higher if they signaled.
- A subsidy
$$s(\theta; D) = \pi^s(\theta) - \left(\pi^f(\theta) - D\right)$$
for $\theta \in [\theta^W, \theta^V]$ and zero otherwise, paid to a type that obtains an audit, induces the efficient channel with the smallest payments that achieve it and leaves the firm's overall profit $\Pi(\theta)$ unchanged.

Figure 24.5 depicts the two conditions, with $D = 0.3$: θ^V is where the red curve $\Delta\pi$ reaches D, and θ^W is where the green curve $\Delta\pi + \Delta CS$ does. In the toy model, $\theta^V = 0.480$ and $\theta^W = 0.152$. The subsidy equals $D - \Delta\pi(\theta)$: at θ^W it equals $\Delta CS(\theta^W) = 0.1995$, the consumer gain that the firm does not internalize, and it falls to zero at θ^V. Table 24.3 reports average surplus per type, with the uniform prior. For $D = 0.3$, voluntary disclosure gives 0.8170, mandatory disclosure gives 0.8563, and the optimum gives 0.8728, against 0.5271 if every type must signal. For $D = 0.5$ the numbers are 0.6204, 0.6563, and 0.7217, so the gap between the voluntary outcome and the optimum is larger when disclosure is costlier.

θ	CS^f	CS^s	ΔCS
0.0	0.2813	0.2813	0.0000
0.2	0.3200	0.0962	0.2238
0.4	0.3613	0.0606	0.3007
0.6	0.4050	0.0423	0.3627
0.8	0.4513	0.0311	0.4201
1.0	0.5000	0.0238	0.4762

Table 24.2: Consumer surplus under full information and under signaling.

	$D = 0.3$	$D = 0.5$
All types signal	0.5271	0.5271
Voluntary disclosure (types above θ^V)	0.8170	0.6204
Mandatory disclosure (all types)	0.8563	0.6563
Socially optimal (types above θ^W)	0.8728	0.7217

Table 24.3: Average total surplus per type, with a uniform prior on $[0, 1]$.

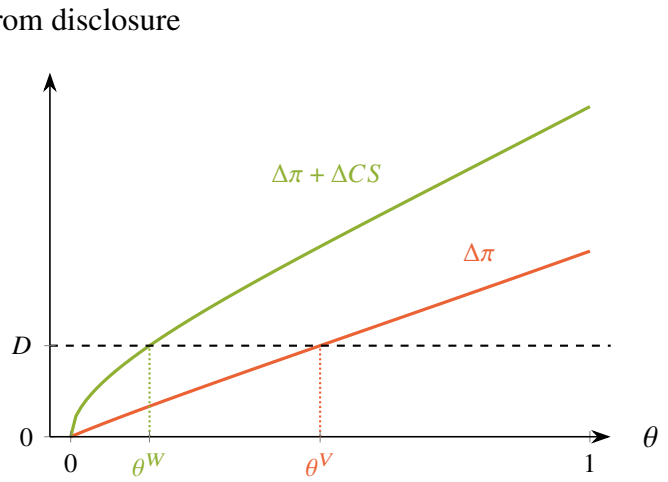


Figure 24.5: Profit gain $\Delta\pi$ and total gain $\Delta\pi + \Delta CS$ from disclosing, against quality, with the disclosure cost $D = 0.3$.

24.8 Sixth Step: The Classical Disclosure Model, Where Quality Does Not Affect Cost

The classical disclosure literature finds that a monopolist discloses too *much*. The paper attributes the difference to two assumptions: that cost does not depend on quality, so that nondisclosure can only pool, and that demand is inelastic. Set $k = 0$ and keep everything else unchanged: θ uniform on $[0, 1]$ and $\alpha = \frac{5}{2}$, $\beta = \delta = 1$, so that $A(\theta) = \frac{3}{2} + \theta$ as before. All types below a cutoff θ^P pool and charge a common price, and consumers' perceived quality is the average of the pool, $\frac{\theta^P}{2}$. The pooled price and profit are $P^f(\frac{\theta^P}{2})$ and $\pi^f(\frac{\theta^P}{2})$, because cost does not vary with type.

The firm's cutoff. The marginal type θ^{PV} is indifferent between pooling and disclosing, so

$$D = \pi^f(\theta) - \pi^f\left(\frac{\theta}{2}\right).$$

With $A(\theta) = \frac{3}{2} + \theta$ and $\pi^f = \frac{A^2}{4}$,

$$D = \frac{A(\theta)^2 - A\left(\frac{\theta}{2}\right)^2}{4} = \frac{3\theta}{8} + \frac{3\theta^2}{16}. \tag{24.11}$$

The planner's cutoff. With $k = 0$, a pooled type consumes $q^P = \frac{1}{2}A\left(\frac{\theta^P}{2}\right)$ at the pooled price and gains the utility of its true quality, so average surplus over the pool equals the surplus at the pool's average quality: the term that is linear in $\theta - \frac{\theta^P}{2}$ integrates to zero over the pool. Total surplus is

$$W^P(\theta^P) = \frac{3}{8}\left[\theta^P A\left(\frac{\theta^P}{2}\right)^2 + \int_{\theta^P}^{1} A(\theta)^2\, d\theta\right] - D(1 - \theta^P),$$

and $\frac{dW^P}{d\theta^P} = 0$ gives

$$D = \frac{3}{8}\left[A(\theta)^2 - A(\tfrac{\theta}{2})^2 - \theta\, A(\tfrac{\theta}{2})\right] = \frac{3\theta^2}{32}. \tag{24.12}$$

Figure 24.6 depicts both cutoffs against D. For the same D, (24.11) exceeds (24.12) at every $\theta > 0$, so $\theta^{PV} < \theta^{PW}$: too many types disclose, as in the classical result. Solving the two equations gives

$$\theta^{PV} = \sqrt{1 + \frac{16D}{3}} - 1 \qquad \text{and} \qquad \theta^{PW} = \sqrt{\frac{32D}{3}}.$$

For example, at $D = 0.06$ the firm's cutoff is $\theta^{PV} = 0.149$ and the planner's is $\theta^{PW} = 0.8$. The planner's cutoff reaches one at $D = \frac{3}{32} = 0.094$, beyond which no disclosure is socially optimal at all.

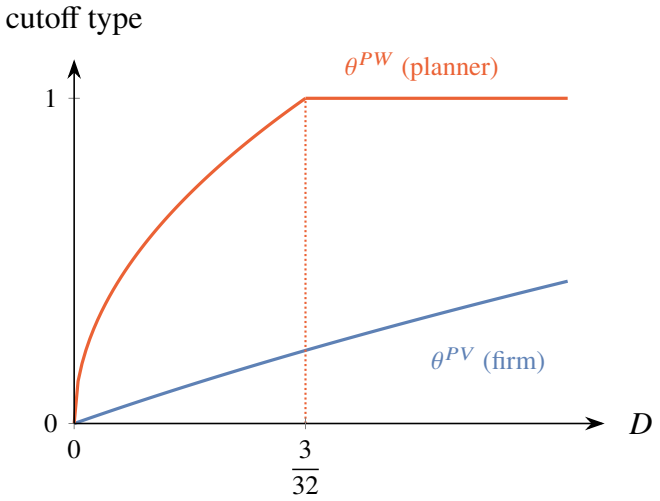


Figure 24.6: Cutoff type above which types disclose, chosen by the firm (θ^{PV}) and by the planner (θ^{PW}), against the disclosure cost D, with $k = 0$ and θ uniform on $[0, 1]$.

Two comments. First, this reproduces the paper's contrast: with a cost that does not depend on quality, voluntary disclosure is excessive, whereas with $k > 0$ in the fifth step it is insufficient. In the signaling model disclosure lowers a distorted price and raises output, a social gain that the firm does not capture; in the pooling model disclosure by a high type mainly redistributes, and it lowers the pool's average quality, which hurts the types that remain in it. Second, Proposition 4(ii) of the paper adds that the ranking reverses, with insufficient disclosure, when D exceeds $\frac{2(\alpha-\delta)^2}{3\beta}$, which is $\frac{3}{2}$ here. The reversal could not be reproduced: for the planner's problem in the paper's equation (6), with each pooled type treated as the pool's mean, as the paper states, the toy model gives $\theta^{PV} < \theta^{PW}$ at every D below $\frac{9}{16}$, the cost at which the firm's cutoff reaches one, and both cutoffs equal one above it. The same holds for any parameters that satisfy Assumption 2, which with $k = 0$ and $\underline{\theta} = 0$ requires $\alpha > \delta$: the firm's and the planner's conditions are

$$D = \frac{4\delta\theta(\alpha - \delta) + 3\delta^2\theta^2}{16\beta} \qquad \text{and} \qquad D = \frac{3\delta^2\theta^2}{32\beta},$$

and the first exceeds the second by

$$\frac{\delta\theta[8(\alpha - \delta) + 3\delta\theta]}{32\beta} > 0.$$

The threshold in Proposition 4(ii) is the value of D at which the two conditions cross: they are equal at

$$\theta = \frac{8(\delta - \alpha)}{3\delta}, \qquad \text{where} \qquad D = \frac{2(\alpha - \delta)^2}{3\beta},$$

but this crossing is positive only if $\alpha < \delta$, a case that Assumption 2 rules out.

24.9 Seventh Step: What the Toy Model Keeps and Drops

The toy model keeps the main conclusions of the paper:

1. When higher quality is costlier, the unique separating equilibrium has the lowest type undistorted and every higher type charging more than its full-information price (the second step, Proposition 1).
2. Signaling profit falls with quality, full-information profit rises, and the gap grows, so only high types disclose (the third and fourth steps, Proposition 2).
3. Overall profit is U-shaped in quality, and for a moderate disclosure cost the highest type earns the most (the fourth step).
4. Voluntary disclosure is socially insufficient, mandatory disclosure is excessive, and a decentralized subsidy achieves the optimum without changing the firm's profit (the fifth step; Section 3 of the paper).
5. When cost does not depend on quality, nondisclosure means pooling and voluntary disclosure is excessive (the sixth step).

The toy model fixes the parameters of Section 24.2, a linear-quadratic specification, and a uniform prior. Relative to Daughety and Reinganum (2008), the firm is a monopolist and price is the only signal; the paper notes that other signals used alongside price, such as advertising, could raise signaling profit and so raise the cutoff θ^V. Disclosure is a perfect, verifiable report bought from a competitive auditor, and a noisy report that keeps the same support would be worth nothing, because the separating price schedule depends only on the support of beliefs. The type space is exogenous: the paper speculates that a firm that chooses its quality by R&D would sample until it draws a low or a high quality, and that its type space could become the union of two intervals. An extension to oligopoly is hard, because each firm's incentive constraint is also a best reply to its rival's price schedule. Each of these features can be relaxed one at a time.

24.10 Connections to Other Chapters

- In Chapter 9, the receiver of the price signal is a potential entrant rather than consumers, and several incumbents send it at once.
- Chapter 25 studies price signaling when the hidden attribute is how clean a product is, and consumers care about total pollution.

- Chapter 26 studies price signaling over two periods, where a responsible firm signals with a low or a high introductory price, depending on competition.
- Chapter 30 studies a sender who commits in advance to the information she discloses, rather than deciding after learning the state.

References

Daughety, A. F., and J. F. Reinganum (2008). "Communicating Quality: A Unified Model of Disclosure and Signalling." *RAND Journal of Economics*, 39(4), pp. 973–989.

Chapter 25

Signaling a Clean Product with Its Price

Companion note to Mahenc, P. (2008), "Signaling the Environmental Performance of Polluting Products to Green Consumers," International Journal of Industrial Organization 26(1): 59–68.

25.1 Introduction

Many products claim to be green. Consider a detergent sold as "eco-friendly": shoppers cannot check how much pollution it generates, and the claim costs nothing to print on the box. Can the price tell consumers what the label cannot?

Mahenc (2008) extends the classic result that a high price can signal high quality (Milgrom and Roberts, 1986) to green products. Consumers dislike the pollution that a product generates, so every firm wants to be perceived as clean. A price reveals cleanness only if a dirty firm finds it too costly to copy. This requires *single crossing*: the clean firm must be more willing than the dirty firm to move its price in exchange for a better perception. When cleaner products are costlier to produce, the clean firm has a thinner margin, so a price increase raises its margin by a larger proportion, and a high price becomes the signal.

The paper finds that, when cleaner products are costlier, the clean firm charges a price above its full-information level. Because it sells less, it also pollutes less than under full information, so signaling helps the environment. When cleaner products are less costly to produce, as in the Porter hypothesis, the distortions go the other way: the clean firm signals with a low price, sells more, and pollutes more. Therefore, a label that reveals cleanness removes one distortion but can worsen another, and it raises welfare only for some levels of pollution damage.

The toy model uses uniform tastes and emissions that are linear in output. The first step solves the full-information benchmark, and the second step describes the firm's incentives when consumers cannot observe cleanness. The third and fourth steps find the separating equilibrium with two types when cleaner products are costlier and when they are less costly, the fifth step solves the model

with a continuum of types, the sixth step asks whether a label raises welfare, and the seventh step lists what the toy model keeps and drops.

25.2 The Toy Model

Emissions. A monopolist produces q units of a good. Production generates pollution

$$e = \frac{q}{\theta}, \tag{25.1}$$

where $\theta > 0$ is the environmental performance of the product: a higher θ means a cleaner product. This is the emission function $e = \frac{t(q)}{\theta}$ of the paper with $t(q) = q$.

Consumers. A unit mass of consumers have basic willingness to pay r, uniformly distributed on $[0, 1]$, and each buys at most one unit. A buyer bears a private damage $v(e) = e$ from the pollution of the product (a health cost, or the guilt from buying a dirty good), so a buyer of type r obtains $r - e - p$. The global damage of pollution falls on buyers and non-buyers alike and therefore does not affect who buys; it enters only the welfare analysis of the sixth step.

Demand. The marginal buyer $\widetilde{r}$ satisfies $\widetilde{r} - e - p = 0$, and the buyers are those with $r \geq \widetilde{r}$, so $q = 1 - \widetilde{r}$. Substituting (25.1),

$$\widetilde{r} - \frac{1 - \widetilde{r}}{\theta} - p = 0,$$

and solving,

$$D(p, \theta) = \alpha(\theta)\,(1 - p), \qquad \alpha(\theta) \equiv \frac{\theta}{1 + \theta}. \tag{25.2}$$

Demand is the standard linear demand $1 - p$ scaled down by the factor $\alpha(\theta) < 1$: every extra buyer adds pollution that lowers the willingness to pay of all buyers, and this effect is weaker for a cleaner product. Pollution is then

$$e = \frac{D(p, \theta)}{\theta} = \frac{1 - p}{1 + \theta}, \tag{25.3}$$

which falls with the price and with cleanness. The inverse demand is $p = 1 - q - \frac{q}{\theta}$, and consumer surplus is $CS = \frac{q^2}{2}$.

Costs. The cost of production is $c(q, \theta) = m(\theta)\, q$, with a constant marginal cost $m(\theta)$. The paper's main case is that a cleaner product is more costly to produce, $m'(\theta) > 0$, so that $c_{q\theta} > 0$.

Numerical example. A dirty type $\theta_L = 1$ and a clean type $\theta_H = 2$, so $\alpha_L = \frac{1}{2}$ and $\alpha_H = \frac{2}{3}$. In the first to third steps the marginal costs are $m_L = 0$ and $m_H = \frac{1}{4}$.

Figure 25.1 depicts demand (25.2) and pollution (25.3) for the two types, with the price on the horizontal axis. Both demands vanish at $p = 1$, but the clean type's demand is higher at every price below 1, because it scales the standard demand by $\alpha_H = \frac{2}{3}$ instead of $\alpha_L = \frac{1}{2}$. Pollution is lower for the clean type at every price, because each unit sold pollutes less, and it falls with the price for both types because fewer units are sold.

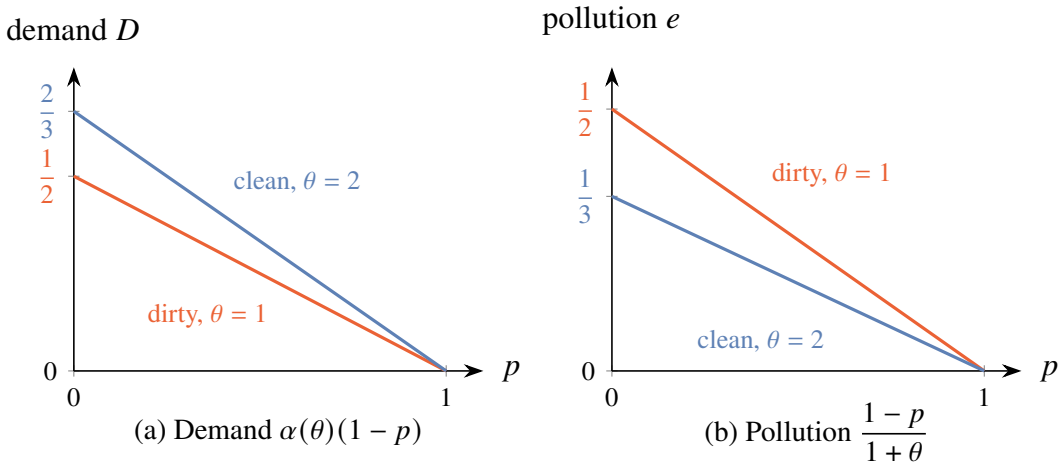


Figure 25.1: Demand and pollution against the price for the dirty type ($\theta_L = 1$) and the clean type ($\theta_H = 2$).

25.3 First Step: Full Information

If consumers observe θ, the firm maximizes $(p-m)\alpha(\theta)(1-p)$, so

$$p^*(\theta) = \frac{1+m(\theta)}{2}, \qquad q^*(\theta) = \frac{\alpha(\theta)\big(1-m(\theta)\big)}{2}, \qquad \pi^*(\theta) = \frac{\alpha(\theta)\big(1-m(\theta)\big)^2}{4}. \tag{25.4}$$

The price elasticity of demand is $\eta = \frac{p}{1-p}$, and at p^* the Lerner index satisfies

$$\frac{p^*-m}{p^*} = \frac{1-m}{1+m} = \frac{1}{\eta(p^*)}, \tag{25.5}$$

the standard result that Lemma 1 of the paper records. Cleanness scales demand by the factor $\alpha(\theta)$, which rises with θ, without changing the elasticity at a given price, so it raises profit through the size of the market. In the example, the dirty type charges $p_L^* = \frac{1}{2}$, sells $\frac{1}{4}$, earns $\pi_L^* = \frac{1}{8}$, and generates $e_L^* = \frac{1}{4}$; the clean type charges $p_H^* = \frac{5}{8}$, sells $\frac{1}{4}$, earns $\pi_H^* = \frac{3}{32}$, and generates $e_H^* = \frac{1}{8}$.

Figure 25.2 depicts the profit of each type against its price. The dirty type's profit $\frac{1}{2}\,p\,(1-p)$ is zero at $p = 0$ and at $p = 1$ and peaks at $p_L^* = \frac{1}{2}$ with height $\frac{1}{8}$. The clean type's profit $\frac{2}{3}\big(p-\frac{1}{4}\big)(1-p)$ is zero at its marginal cost $p = \frac{1}{4}$ and at $p = 1$, and peaks at $p_H^* = \frac{5}{8}$ with height $\frac{3}{32}$. The clean type charges more because its marginal cost is higher, and it earns less because its higher cost outweighs its larger market. These two curves, and their crossings with the profit levels in the third step, are the ingredients of Figure 25.4.

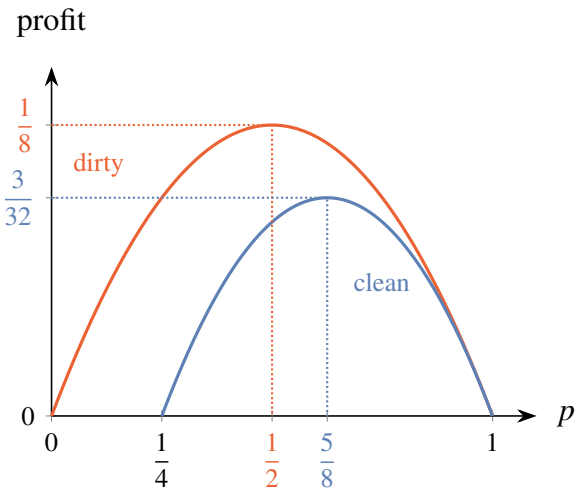


Figure 25.2: Profit of each type against its price under full information, with $(m_L, m_H) = (0, \frac{1}{4})$.

25.4 Second Step: Hidden Cleanness and the Firm's Incentives

Now consider that only the firm observes θ. Consumers see the price and form an inference $\widehat{\theta}$, and they buy according to $D(p, \widehat{\theta})$. A firm of type θ that is perceived as $\widehat{\theta}$ earns

$$\pi(p, \theta, \widehat{\theta}) = \alpha(\widehat{\theta})\,\big(p - m(\theta)\big)(1 - p). \tag{25.6}$$

We look for separating perfect Bayesian equilibria, in which different types charge different prices and consumers infer the type from the price. Two properties of (25.6) drive the analysis.

Every type wants to be perceived as clean. Since α is increasing, $\pi_{\widehat{\theta}} > 0$ whenever $p > m(\theta)$ (Condition 2 of the paper). The worst inference for any type is therefore the dirtiest one.

Single crossing. The perceived type multiplies profit, so a type's ranking of prices for a given perception depends only on $\ln\big[(p-m)(1-p)\big]$, whose slope in p is

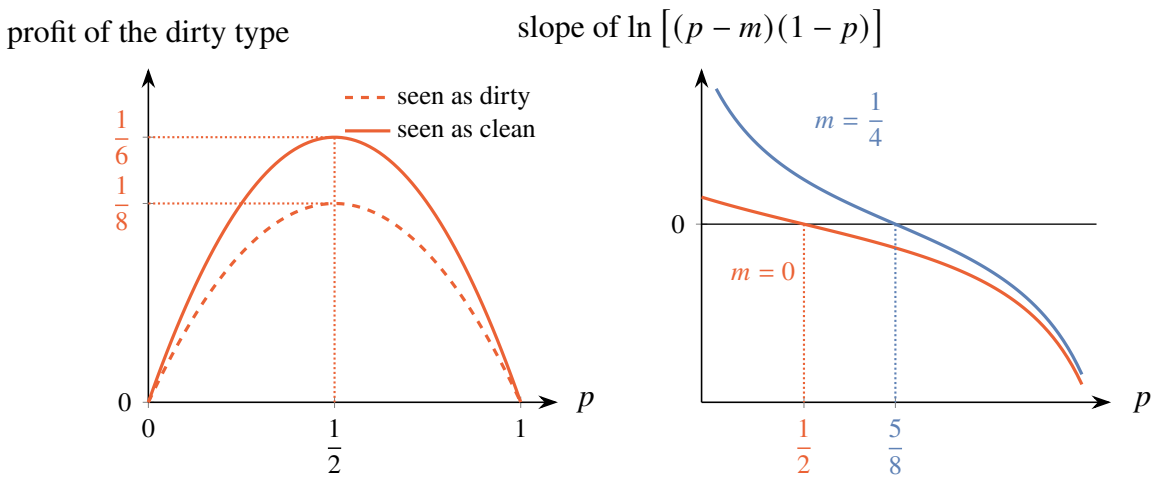
$$\frac{\partial}{\partial p}\ln\big[(p-m)(1-p)\big] = \frac{1}{p-m} - \frac{1}{1-p}, \qquad \frac{\partial}{\partial m}\left[\frac{1}{p-m} - \frac{1}{1-p}\right] = \frac{1}{(p-m)^2} > 0. \tag{25.7}$$

A firm with a higher marginal cost gains more, in proportional terms, from raising its price: a price increase raises its thinner margin by a larger proportion, while the proportional loss of sales, $\frac{1}{1-p}$ per unit of price, is the same for every type. When cleaner products are more costly ($m' > 0$), a clean firm is more willing than a dirty firm to trade a higher price for a better perception. In terms of the paper's Condition 4, the marginal rate of substitution between price and perceived type satisfies

$$\frac{\partial}{\partial \theta}\left[-\frac{\pi_{\widehat{\theta}}}{\pi_p}\right] = \frac{c_{q\theta}\, D_{\widehat{\theta}}\, D}{\pi_p^2}, \tag{25.8}$$

and since $D_{\widehat{\theta}} = \frac{1-p}{(1+\widehat{\theta})^2} > 0$, the sign of the single-crossing condition is the sign of $c_{q\theta}$.

Figure 25.3 depicts both properties. Panel (a) shows the dirty type's profit when it is perceived as dirty and when it is perceived as clean: the second curve is the first one multiplied by $\frac{\alpha_H}{\alpha_L} = \frac{4}{3}$, so it lies above it at every price and the two peak at the same price. Panel (b) shows the slope (25.7) of $\ln\left[(p-m)(1-p)\right]$ for $m = 0$ and $m = \frac{1}{4}$. The slope is higher for the costlier type at every price, and it is zero at the price that maximizes $(p-m)(1-p)$, which is $\frac{1}{2}$ for $m = 0$ and $\frac{5}{8}$ for $m = \frac{1}{4}$. At a price between $\frac{1}{2}$ and $\frac{5}{8}$, for example, the dirty type's profit falls in proportional terms while the clean type's profit is still rising.



(a) Perception raises profit (b) Single crossing

Figure 25.3: Profit of the dirty type ($m = 0$) when perceived as dirty (dashed, α_L) and as clean (solid, α_H), and the slope of $\ln\left[(p-m)(1-p)\right]$ for $m = 0$ and $m = \frac{1}{4}$.

25.5 Third Step: Two Types when Cleaner Products Are Costlier

Let consumers believe that the firm is clean if it charges p_H and dirty otherwise. In a separating equilibrium the dirty type is recognized, so it charges its full-information price $p_L^* = \frac{1}{2}$ and earns $\frac{1}{8}$: there is no distortion at the bottom (Lemma 2 of the paper). The clean type's price is determined by two incentive constraints. Each one compares a profit that depends on the price, which is a curve in p_H, with a payoff that does not depend on it, which is a horizontal line, and Figure 25.4 is built from these four objects.

The dirty type must not mimic. Charging p_H would make the dirty type look clean, and it would earn

$$\alpha_H\,(p_H - m_L)(1-p_H) = \frac{2}{3}\,p_H(1-p_H),$$

a curve, instead of its equilibrium profit $\frac{1}{8}$, a line. Separation requires

$$\frac{2}{3}\,p_H(1-p_H) \le \frac{1}{8}. \tag{25.9}$$

The clean type must prefer to separate. If the clean type charges any other price, consumers believe it is dirty, and its best such deviation is its full-information price against the dirty type's demand, which earns

$$\alpha_L \frac{(1-m_H)^2}{4} = \frac{9}{128},$$

a line. Charging p_H earns

$$\alpha_H (p_H - m_H)(1-p_H) = \frac{2}{3}\Big(p_H - \frac{1}{4}\Big)(1-p_H),$$

a curve. Separation requires

$$\frac{2}{3}\Big(p_H - \frac{1}{4}\Big)(1-p_H) \geq \frac{9}{128}. \tag{25.10}$$

Rearranging the conditions. Each condition is a quadratic inequality in p_H. Moving everything to one side of (25.9), and multiplying (25.10) by $\frac{3}{2}$ to get $\big(p_H - \frac{1}{4}\big)(1-p_H) \geq \frac{27}{256}$ before doing the same,

$$\begin{aligned} (25.9) \quad &\Longleftrightarrow \quad p_H^2 - p_H + \frac{3}{16} \geq 0 \quad \Longleftrightarrow \quad \Big(p_H - \frac{1}{4}\Big)\Big(p_H - \frac{3}{4}\Big) \geq 0 \\ &\Longleftrightarrow \quad p_H \leq \frac{1}{4} \text{ or } p_H \geq \frac{3}{4}, \end{aligned} \tag{25.11}$$

$$\begin{aligned} (25.10) \quad &\Longleftrightarrow \quad p_H^2 - \frac{5}{4}p_H + \frac{91}{256} \leq 0 \quad \Longleftrightarrow \quad \Big(p_H - \frac{7}{16}\Big)\Big(p_H - \frac{13}{16}\Big) \leq 0 \\ &\Longleftrightarrow \quad \frac{7}{16} \leq p_H \leq \frac{13}{16}. \end{aligned} \tag{25.12}$$

The four roots are the prices at which a curve crosses its line, and they are the numbers marked on the horizontal axis of Figure 25.4(a).

Constructing Figure 25.4(a). Students can reproduce the panel in six steps, using the numbers of Table 25.1.

1. Draw the price p on the horizontal axis, from 0 to 1, and profit on the vertical axis, from 0 to about 0.2.
2. Plot the profit that the dirty type would earn if it were seen as clean, $\frac{2}{3}\,p(1-p)$. It is a downward parabola that is zero at $p = 0$ and $p = 1$ and peaks at $p = \frac{1}{2}$ with height $\frac{1}{6}$.
3. Draw the horizontal line at the dirty type's equilibrium profit, $\frac{1}{8}$. The parabola crosses it where $p(1-p) = \frac{3}{16}$, that is, at $p = \frac{1}{4}$ and $p = \frac{3}{4}$, and lies above it in between, where the dirty type would mimic.
4. Plot the profit of the clean type if it is seen as clean, $\frac{2}{3}\big(p - \frac{1}{4}\big)(1-p)$, for $p \geq \frac{1}{4}$. It is zero at $p = \frac{1}{4}$ and $p = 1$ and peaks at $p = \frac{5}{8}$ with height $\frac{3}{32}$, its full-information profit.
5. Draw the horizontal line at the clean type's profit if it is seen as dirty, $\frac{9}{128}$. The curve crosses it where $\big(p - \frac{1}{4}\big)(1-p) = \frac{27}{256}$, at $p = \frac{7}{16}$ and $p = \frac{13}{16}$, and lies above it in between, where the clean type prefers to separate.

6. Shade the prices at which the dirty type's curve is below its line and the clean type's curve is above its line, which gives $\left[\frac{3}{4}, \frac{13}{16}\right]$. The prices $p \leq \frac{1}{4}$ also satisfy (25.11), but there the clean type's curve is below its line.

Object in Figure 25.4(a)	Formula	Zeros or crossings	Peak
Dirty type, seen as clean (red curve)	$\frac{2}{3}\,p(1-p)$	0 and 1	$\frac{1}{6}$ at $p = \frac{1}{2}$
Dirty type's equilibrium profit (red line)	$\frac{1}{8}$	$\frac{1}{4}$ and $\frac{3}{4}$	–
Clean type, seen as clean (blue curve)	$\frac{2}{3}\left(p - \frac{1}{4}\right)(1-p)$	$\frac{1}{4}$ and 1	$\frac{3}{32}$ at $p = \frac{5}{8}$
Clean type, seen as dirty (blue line)	$\frac{9}{128}$	$\frac{7}{16}$ and $\frac{13}{16}$	–

Table 25.1: Ingredients of Figure 25.4(a). The crossings are those of each curve with the line of the same color.

Combining (25.9) and (25.10), the separating prices for the clean type are $p_H \in \left[\frac{3}{4}, \frac{13}{16}\right]$. The low prices $p_H \leq \frac{1}{4}$ that would also deter the dirty type are useless to the clean type, since they leave it no margin over its marginal cost $\frac{1}{4}$. Among the separating equilibria, the Intuitive Criterion of Cho and Kreps (1987) selects the least costly one, $p_H = \frac{3}{4}$: if the clean type charged $p_H > \frac{3}{4}$, a deviation to a price in $\left(\frac{3}{4}, p_H\right)$ would not be profitable for the dirty type even if consumers believed it was clean, but it would be profitable for the clean type, so consumers should attribute the deviation to the clean type.

Result 1 (Proposition 1 of the paper, two types). When cleaner products are more costly to produce:

- the clean type signals with a price above its full-information level, $\frac{3}{4} > \frac{5}{8}$;
- the dirty type charges its full-information price, $\frac{1}{2}$;
- the clean type sells less, $\frac{1}{6}$ instead of $\frac{1}{4}$, and pollutes less, $\frac{1}{12}$ instead of $\frac{1}{8}$.

Figure 25.4(a) depicts the two incentive constraints, with $(m_L, m_H) = (0, \frac{1}{4})$: the dirty type mimics at the prices where the red curve lies above the dotted red line, the clean type separates at the prices where the blue curve lies above the dotted blue line, and the separating prices, shaded, satisfy both conditions.

Signaling is costly for the clean type: its profit falls from $\frac{3}{32}$ to $\frac{1}{12}$, a loss of $\frac{1}{96}$. Pollution falls by one third, because the signal works by restricting output and output is what pollutes.

25.6 Fourth Step: When Cleaner Products Are Less Costly

Porter and van der Linde (1995) argue that cleaner products can be less costly to produce. Swap the marginal costs: $m_L = \frac{1}{4}$ and $m_H = 0$, so that $c_{q\theta} < 0$. Under full information the dirty type charges $\frac{5}{8}$ and earns $\frac{9}{128}$, and the clean type charges $\frac{1}{2}$, sells $\frac{1}{3}$, earns $\frac{1}{6}$, and generates $e^*_H = \frac{1}{6}$.

The constraint that the dirty type must not mimic compares its curve $\frac{2}{3}\big(p-\frac{1}{4}\big)(1-p)$ with its full-information profit $\frac{9}{128}$, and the constraint that the clean type must prefer to separate compares its curve $\frac{2}{3}\,p(1-p)$ with the profit that it earns if it is seen as dirty,

$$\frac{1}{2}\cdot\frac{1}{4}=\frac{1}{8}.$$

They become

$$\frac{2}{3}\Big(p_H-\frac{1}{4}\Big)(1-p_H)\le\frac{9}{128}\quad\Longleftrightarrow\quad p_H\le\frac{7}{16}\ \text{ or }\ p_H\ge\frac{13}{16}, \tag{25.13}$$

$$\frac{2}{3}\,p_H(1-p_H)\ge\frac{1}{8}\quad\Longleftrightarrow\quad\frac{1}{4}\le p_H\le\frac{3}{4}, \tag{25.14}$$

so the separating prices are $p_H\in\big[\frac{1}{4},\frac{7}{16}\big]$. Panel (b) of Figure 25.4 is built with the same six steps, with the roles of the two types exchanged: the red curve is the same parabola as the blue curve of panel (a) and its line is $\frac{9}{128}$, and the blue curve is the same parabola as the red curve of panel (a), with the line $\frac{1}{8}$, so the crossings $\frac{1}{4}$, $\frac{3}{4}$ and $\frac{7}{16}$, $\frac{13}{16}$ appear again, now for the other type. The least costly separating price for the clean type is now the highest one in the interval, $p_H=\frac{7}{16}$, because its profit $\frac{2}{3}p(1-p)$ rises up to $p=\frac{1}{2}$.

Result 2 (cleaner is less costly). When cleaner products are less costly to produce:

- the clean type signals with a price below its full-information level, $\frac{7}{16}<\frac{1}{2}$;
- the dirty type charges its full-information price, $\frac{5}{8}$;
- the clean type sells more, $\frac{3}{8}$ instead of $\frac{1}{3}$, and pollutes more, $\frac{3}{16}$ instead of $\frac{1}{6}$.

Figure 25.4(b) depicts the two incentive constraints, with $(m_L,m_H)=(\frac{1}{4},0)$: the shaded separating prices are now $\big[\frac{1}{4},\frac{7}{16}\big]$, below the clean type's full-information price $\frac{1}{2}$.

Intuitively, the clean firm now has the wider margin, so a price cut lowers its margin by a smaller proportion than it lowers the dirty firm's, while both gain the same proportion of sales. A price cut is therefore the signal that the dirty firm finds too costly to copy, and the direction of the signal follows the sign of $c_{q\theta}$, as (25.8) shows.

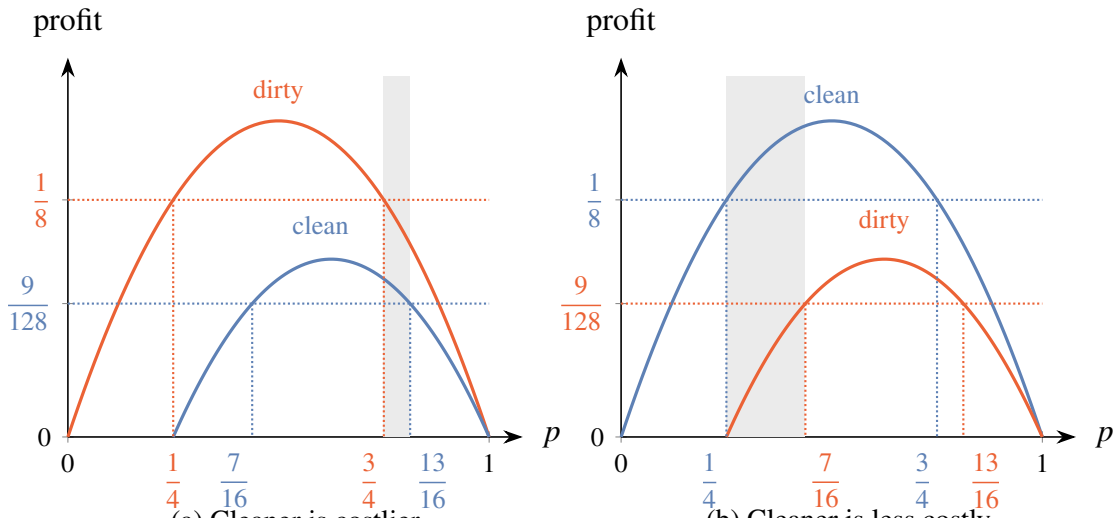
Table 25.2 reports the full-information and signaling outcomes in both cases.

25.7 Fifth Step: A Continuum of Types

The paper works with a continuum of types $\theta\in[\underline{\theta},\overline{\theta}]$ and the method of Mailath (1987). Let $\theta\in[1,2]$ with $m(\theta)=\frac{\theta-1}{4}$, which reproduces the two types of the third step at the end points. A separating equilibrium is a strictly monotone price function $\varphi(\theta)$, and a consumer who sees p infers $\widehat{\theta}=\varphi^{-1}(p)$.

Local incentive compatibility. Type θ chooses which type $\widehat{\theta}$ to imitate, earning

$$\Pi(\theta,\widehat{\theta})\equiv\pi\big(\varphi(\widehat{\theta}),\theta,\widehat{\theta}\big).$$



(a) Cleaner is costlier (b) Cleaner is less costly

Figure 25.4: Profit from charging p and being perceived as clean, for each type, and each type's payoff outside the clean type's price (dotted). Shaded: separating prices for the clean type. Types $\theta_L = 1$ and $\theta_H = 2$; marginal costs $(m_L, m_H) = (0, \frac{1}{4})$ in panel (a) and $(\frac{1}{4}, 0)$ in panel (b).

Case	Type	p^*	p^s	q^*	q^s	e^*	e^s	π^*	π^s
Cleaner is costlier	Dirty	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{8}$
$(m_L, m_H) = (0, \frac{1}{4})$	Clean	$\frac{5}{8}$	$\frac{3}{4}$	$\frac{1}{4}$	$\frac{1}{6}$	$\frac{1}{8}$	$\frac{1}{12}$	$\frac{3}{32}$	$\frac{1}{12}$
Cleaner is less costly	Dirty	$\frac{5}{8}$	$\frac{5}{8}$	$\frac{3}{16}$	$\frac{3}{16}$	$\frac{3}{16}$	$\frac{3}{16}$	$\frac{9}{128}$	$\frac{9}{128}$
$(m_L, m_H) = (\frac{1}{4}, 0)$	Clean	$\frac{1}{2}$	$\frac{7}{16}$	$\frac{1}{3}$	$\frac{3}{8}$	$\frac{1}{6}$	$\frac{3}{16}$	$\frac{1}{6}$	$\frac{21}{128}$

Table 25.2: Price, output, pollution, and profit under full information ($*$) and in the least costly separating equilibrium (s), with $\theta_L = 1$ and $\theta_H = 2$.

Truth-telling requires $\Pi_{\widehat{\theta}}(\theta,\theta)=0$, that is, $\pi_p\,\varphi'(\theta)+\pi_{\widehat{\theta}}=0$, which gives the differential equation

$$\varphi'(\theta)=-\frac{\pi_{\widehat{\theta}}}{\pi_p}=\frac{\big(\varphi-m(\theta)\big)\big(1-\varphi\big)}{\theta(1+\theta)\big(2\varphi-1-m(\theta)\big)}, \tag{25.15}$$

where we used

$$\frac{\alpha'(\theta)}{\alpha(\theta)}=\frac{1}{\theta(1+\theta)} \qquad \text{and} \qquad \pi_p=\alpha(\theta)\big(1+m(\theta)-2\varphi\big).$$

The initial condition is $\varphi(1)=p^*(1)=\frac{1}{2}$: the dirtiest type is recognized as the dirtiest, so it has no reason to distort (Lemma 2 of the paper).

The price exceeds the full-information price. Along a separating equilibrium $\varphi'>0$ and $\pi_{\widehat{\theta}}>0$, so (25.15) requires $\pi_p<0$, that is,

$$\varphi(\theta)>\frac{1+m(\theta)}{2}=p^*(\theta)$$

for every $\theta>1$. Each type prices above the level at which its profit stops rising, which is what makes the price costly enough to deter the types just below it.

Result 3 (Proposition 1 of the paper). With a continuum of types and $c_{q\theta}>0$:

- the separating price $\varphi(\theta)$ is strictly increasing;
- it exceeds $p^*(\theta)$ for every type except the dirtiest;
- every type except the dirtiest pollutes less than under full information.

Figure 25.5 depicts the solution of (25.15), with $\theta\in[1,2]$ and $m(\theta)=\frac{\theta-1}{4}$. The equation has no closed-form solution, but it is easy to integrate numerically (as a function $\theta(\varphi)$, whose slope is zero at the initial point). In panel (a), the signaling price rises from 0.500 at $\theta=1$ to 0.714 at $\theta=1.5$ and 0.765 at $\theta=2$, while the full-information price $p^*(\theta)=\frac{\theta+3}{8}$ rises only from 0.500 to 0.625. In panel (b), pollution under signaling falls from 0.250 to 0.078, against 0.125 under full information for the cleanest type. The price function has an infinite slope at $\theta=1$: the types just above the dirtiest one must move their prices away from $\frac{1}{2}$ quickly, because near p^* a small price change costs almost nothing.

The cleanest type charges 0.765 in the continuum, more than the $\frac{3}{4}$ it charges with two types. With two types, the clean firm only needs to deter the dirtiest firm, whose profit is undistorted. With a continuum, each type must deter the types just below it, which already charge distorted prices, and the distortions accumulate along the chain. A discrete chain of types confirms this: the top price is 0.750 with two types, 0.762 with 11 equally spaced types, and approaches 0.765 as the grid becomes finer.

25.8 Sixth Step: Should a Label Reveal Cleanness?

The paper notes that a labeling program that resolves consumers' uncertainty removes one distortion but can worsen another: by ending the need to signal, it can raise pollution. The toy model makes

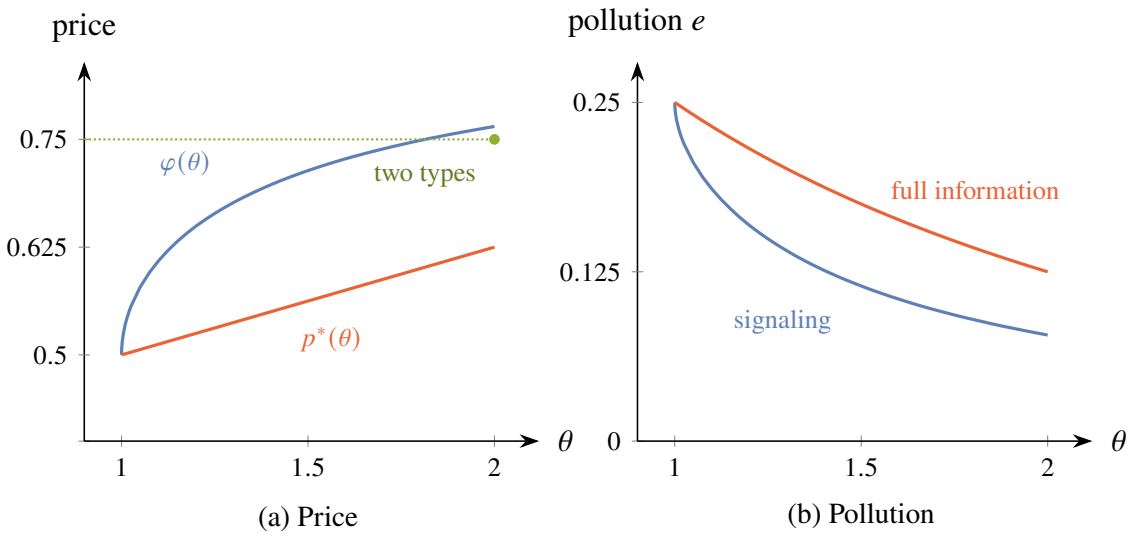


Figure 25.5: Separating price and pollution with a continuum of types $\theta \in [1,2]$ and $m(\theta) = \frac{\theta-1}{4}$, against full information.

the trade-off explicit. Let the global damage of pollution be δe, borne by everyone, and measure welfare as $CS + \pi - \delta e$, with $CS = \frac{q^2}{2}$. A label that reveals θ moves the clean type from its signaling outcome to its full-information outcome and leaves the dirty type unchanged.

Cleaner is costlier. With the label the clean type sells $\frac{1}{4}$ instead of $\frac{1}{6}$. Consumer surplus rises from $\frac{1}{72}$ to $\frac{1}{32}$ and profit from $\frac{1}{12}$ to $\frac{3}{32}$, a gain of $\frac{1}{36}$ in $CS + \pi$, while pollution rises by

$$\frac{1}{8} - \frac{1}{12} = \frac{1}{24}.$$

The label raises welfare if and only if

$$\frac{1}{36} > \frac{\delta}{24} \quad \Longleftrightarrow \quad \delta < \frac{2}{3}. \tag{25.16}$$

Cleaner is less costly. With the label the clean type raises its price from $\frac{7}{16}$ to $\frac{1}{2}$ and sells $\frac{1}{3}$ instead of $\frac{3}{8}$. $CS + \pi$ falls by $\frac{7}{576}$, because the signaling price cut had moved the monopoly output closer to the efficient level, while pollution falls by

$$\frac{3}{16} - \frac{1}{6} = \frac{1}{48}.$$

The label raises welfare if and only if

$$\frac{\delta}{48} > \frac{7}{576} \quad \Longleftrightarrow \quad \delta > \frac{7}{12}. \tag{25.17}$$

Result 4 (labels). A label that reveals cleanness raises welfare:

- when cleaner products are costlier, if and only if the pollution damage is low, $\delta < \frac{2}{3}$;

- when cleaner products are less costly, if and only if the pollution damage is high, $\delta > \frac{7}{12}$.

Signaling distorts the clean firm's output, and the direction of the distortion determines whether removing it helps the environment or the market. This illustrates the paper's point that a single instrument can mitigate one distortion and worsen another; with three distortions (the externality, market power, and asymmetric information), the paper calls for three instruments: a Pigouvian tax on emissions, a tax or subsidy that offsets the price distortion, and a lump-sum transfer.

25.9 Seventh Step: What the Toy Model Keeps and Drops

The toy model keeps the paper's main conclusions:

1. Every type wants to look clean (Condition 2), and the dirtiest type charges its full-information price (Lemma 2, the third step).
2. The direction of the signal follows the sign of $c_{q\theta}$: a price above the full-information level when cleaner products are costlier, below it when they are less costly (Proposition 1 and its discussion, the third and fourth steps).
3. Signaling lowers pollution when cleaner products are costlier, from $\frac{1}{8}$ to $\frac{1}{12}$ for the clean type with two types, and for every type above the dirtiest with a continuum (the third and fifth steps).
4. Full disclosure can raise pollution, so a label alone need not raise welfare (the sixth step).

The linear-uniform specification adds two features to the general analysis. First, cleanness enters demand as the multiplicative factor $\alpha(\theta) = \frac{\theta}{1+\theta}$, which makes single crossing a one-line calculation (25.7). Second, the comparison between two types and a continuum shows that the distortion of the cleanest type grows with the number of types it must separate from, from 0.750 to 0.765.

Relative to Mahenc (2008), the toy model fixes $v(e) = e$, $t(q) = q$, basic willingness to pay uniform on $[0, 1]$ (the paper's $[0, R]$ with $R = 1$), and a constant marginal cost; the paper allows a convex private damage, a convex emission function, and a general cost $c(q, \theta)$. As in the paper, the firm is a monopolist, cleanness is exogenous, the price is the only signal, and the analysis focuses on separating equilibria. The paper's conclusion suggests advertising and green investment as additional signals; each of these features, and competition between a clean and a dirty firm, can be added one at a time.

25.10 Connections to Other Chapters

- Chapter 24 lets the firm choose between signaling through its price and disclosing its quality at a cost.
- Chapter 26 studies price signaling of a credence attribute over two periods, before consumers learn the truth.

- In Chapter 29, consumers cannot verify green claims either, and firms under a common label pass off conventional products as green.

References

Cho, I.-K., and D. M. Kreps (1987). "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics*, 102(2), pp. 179–221.

Mahenc, P. (2008). "Signaling the Environmental Performance of Polluting Products to Green Consumers." *International Journal of Industrial Organization*, 26(1), pp. 59–68.

Mailath, G. J. (1987). "Incentive Compatibility in Signaling Games with a Continuum of Types." *Econometrica*, 55(6), pp. 1349–1365.

Milgrom, P., and J. Roberts (1986). "Price and Advertising Signals of Product Quality." *Journal of Political Economy*, 94(4), pp. 796–821.

Porter, M. E., and C. van der Linde (1995). "Toward a New Conception of the Environment-Competitiveness Relationship." *Journal of Economic Perspectives*, 9(4), pp. 97–118.

Chapter 26

Penetration or Skimming Pricing for Credence Products?

Companion note to Mahenc, P. (2025), "Penetration or Skimming Pricing for Credence Products?," Journal of Economic Dynamics and Control 180: 105196.

26.1 Introduction

Some product qualities cannot be checked even after purchase. Consider a new coffee brand that claims its beans are sustainably sourced, a *credence attribute* that no shopper can taste. The truth comes out only after a *verification phase*, in which certifiers, NGOs, or regulators investigate the claim. Until then, the price is the only signal the entrant can send. Which price path convinces consumers that the claim is true?

Mahenc (2025) shows that the answer depends on competition. A conventional firm that pretends to be responsible gains only during the verification phase, since it is exposed afterwards. The responsible entrant must choose a first-period price at which mimicking brings this firm no gain. In an *occupied* market, where conventional incumbents already sell at cost, the entrant charges the competitive price. A conventional firm earns nothing at that price, but the responsible entrant, with its higher cost, takes a loss. In an *untapped* market, where the entrant is a monopolist, it charges a high price instead and gives up sales. In both cases the firm wastes money in a visible way to prove its type, an instance of Veblen's "conspicuous waste."

Therefore, a responsible entrant uses *penetration pricing* (a low introductory price that rises later) when it faces competitors, and *skimming pricing* (a high introductory price that falls later) when it does not. Penetration requires a patient entrant, since the loss comes first and the reward later. Signaling is easier when the product is greener and its cost premium is smaller, and harder when verification takes longer, in which case greenwashing (pooling) fills the early periods.

The toy model collapses the verification lag to a single period. The first step solves the full-information benchmark and the second step sets up the signaling problem. The third and fourth

steps study penetration and skimming, the fifth step derives the comparative statics and the role of the verification lag, the sixth step works through a numerical example, and the seventh step lists what the toy model keeps and drops.

26.2 The Toy Model

Consumers. A unit mass of consumers each buy at most one unit. A unit with *social performance* $e \in [0, 1]$ generates $1 - e$ units of a harmful by-product; the conventional product has $e = 0$. Consumers differ in their *aversion* x to the harm, with x distributed uniformly on $[0, 1]$. A consumer of aversion x who buys a unit of performance e at price p gets

$$u = v - (1 - e)\, x - p, \tag{26.1}$$

where v is the gross value of the product category. A higher e lowers the harm term $(1 - e)x$, so consumers pay more for a greener product.

The entrant and costs. An entrant introduces a product of performance e at a constant unit cost c_e, against a conventional alternative of unit cost c_0. Write the *cost premium* of responsibility as

$$\varphi \;\equiv\; c_e - c_0 \;>\; 0. \tag{26.2}$$

A fixed setup cost $F > 0$ is sunk once the firm enters; because it is paid whether or not the firm signals, it cancels from every comparison below and we drop it.

Two scenarios. The entrant faces either

- an *untapped market*: no competitor, so the entrant is a monopolist; or
- an *occupied market*: perfectly competitive incumbents already sell the conventional product at price c_0.

The two settings differ only in the outside option consumers have, and that difference is what flips the sign of the price signal.

Normalization. To keep the algebra light we set

$$v = 1, \qquad c_0 = 0, \tag{26.3}$$

so the conventional unit cost (and competitive price) is 0 and the responsible unit cost is $c_e = \varphi$. Throughout we assume $\varphi < e$ (the responsible product is socially desirable) and $v - c_0 \geq 1$ (the occupied market is fully covered); both hold under (26.3) for the parameters we use.

Timing and the verification lag. Social responsibility is the entrant's private information. In Mahenc's full model it takes T trading periods for consumers to learn the truth; here we take the simplest case $T = 1$. The game then has two stages:

Stage 1. Verification period ($t = 1$). The entrant posts a price; consumers buy and try to infer the type from the price.

Stage 2. Full information ($t \geq 2$). If the entrant sold in period 1, the type becomes public (as in Mahenc, verification requires that the product be sold), and the entrant charges the full-information price forever after.

Players discount the future with discount factor $\delta \in (0, 1)$, so the value of the full-information continuation, seen from $t = 1$, is $\frac{\delta}{1-\delta}$ times the per-period full-information profit. Section 26.7 restores general T in one line.

26.3 First Step: The Full-Information Benchmark

Consider that the type is known. We record the responsible entrant's optimal price and profit in each scenario; these are the post-verification payoffs the signal is trying to protect.

Untapped market. The indifferent consumer solves $0 = v - (1 - e)x - p$, giving demand

$$D_u(p) = \min\left\{1, \frac{v-p}{1-e}\right\} = \min\left\{1, \frac{1-p}{1-e}\right\} :$$

as in the paper, demand cannot exceed the unit mass of consumers, so it equals one when $p \leq e$. We also assume, as the paper does, $\frac{1+\varphi}{2} \geq e$, so the full-information price below lies where demand is $\frac{1-p}{1-e}$. The responsible entrant maximizes $(p - \varphi)D_u(p)$, yielding

$$\hat{p}^u_{\mathrm{r}} = \frac{v + c_e}{2} = \frac{1+\varphi}{2}, \qquad \hat{\pi}^u_{\mathrm{r}} = \frac{(v-c_e)^2}{4(1-e)} = \frac{(1-\varphi)^2}{4(1-e)}. \tag{26.4}$$

A conventional monopolist would instead charge and earn

$$\hat{p}^u_{\mathrm{c}} = \frac{v + c_0}{2} = \frac{1}{2} \qquad \text{and} \qquad \hat{\pi}^u_{\mathrm{c}} = \frac{(v-c_0)^2}{4} = \frac{1}{4}.$$

Occupied market. A consumer is indifferent between the conventional product at c_0 and the new product at p when

$$v - x - c_0 = v - (1 - e)x - p, \qquad \text{that is, at} \qquad x = \frac{p - c_0}{e} = \frac{p}{e}.$$

Consumers above this cutoff buy the responsible product, giving demand

$$D_o(p) = 1 - \frac{p - c_0}{e} = 1 - \frac{p}{e}.$$

Maximizing $(p - \varphi)D_o(p)$,

$$\hat{p}^o_{\mathrm{r}} = \frac{e + c_0 + c_e}{2} = \frac{e+\varphi}{2}, \qquad \hat{\pi}^o_{\mathrm{r}} = \frac{(e-\varphi)^2}{4e}. \tag{26.5}$$

A conventional entrant here is forced by competition to price at $c_0 = 0$ and earns a *zero* margin: $\hat{\pi}^o_{\mathrm{c}} = 0$. This zero-profit outside option drives the occupied-market result.

26.4 Second Step: The Signaling Problem

Now let the type be private in period 1. The entrant is responsible (r) or conventional (c); consumers hold a prior μ_0 that it is responsible and update to a posterior μ after seeing the price. Consumers who believe the product is responsible with probability μ treat its perceived performance as μe, so demand under belief μ is

$$D(p,\mu) = \begin{cases} 1 - \dfrac{p - c_0}{\mu e} & \text{occupied market,} \\ \dfrac{v - p}{1 - \mu e} & \text{untapped market.} \end{cases} \tag{26.6}$$

For any price, demand rises with μ: being *perceived* as responsible is always worth more, regardless of the truth. That is why a conventional entrant is tempted to mimic, and why the responsible entrant needs a signal a mimic would not copy.

Separating logic. We look for a *separating* equilibrium in which the responsible type charges a signaling price $\underline{p}$ in period 1, is believed ($\mu = 1$), and then earns the full-information profit forever; the conventional type charges its own full-information price and is believed conventional ($\mu = 0$). Two conditions must hold.

- Incentive compatibility (IC): the conventional type must not want to mimic $\underline{p}$. Its gain from being taken as responsible for one period must not exceed what it earns honestly. IC determines the *direction* of the price distortion.
- Individual rationality (IR): the responsible type must prefer to signal rather than accept being mistaken for conventional. IR determines *when* signaling is worthwhile.

Among all separating prices, the least distorted one (hence the most profitable for the responsible type) is the *Riley outcome*, the unique signal that survives the Intuitive Criterion of Cho and Kreps (1987). We take that price in each scenario.

26.5 Third Step: Occupied Market, Penetration Pricing

IC forces a low price. Consider that the responsible type is expected to charge $\underline{p}$. If the conventional type mimics it and is believed responsible, it earns the flow profit

$$\pi_{\mathrm{c}}(\underline{p},1) = (\underline{p} - c_0)\, D_o(\underline{p},1) = \underline{p}\left(1 - \frac{\underline{p}}{e}\right),$$

which is strictly positive for any $\underline{p} \in (0,e)$ and peaks at $\underline{p} = \frac{e}{2}$. But honest behavior earns the conventional type only $\hat{\pi}^o_{\mathrm{c}} = 0$. So it would mimic *any* interior price: high prices cannot separate. To deter mimicry the responsible type must drive the price down to where the conventional type gains nothing,

$$\underline{p}\left(1 - \frac{\underline{p}}{e}\right) \le 0 \quad \Longleftrightarrow \quad \underline{p} \le c_0 = 0 \ \text{ or } \ \underline{p} \ge e. \tag{26.7}$$

(A distortion the other way, $\underline{p} \geq e$, would choke demand to zero, so no one could learn.) The least costly such price (the highest one that still separates) is $\underline{p} = c_0$: the responsible entrant matches the competitive price.

Result 1 (penetration). In the occupied market the responsible entrant signals by charging the competitive price $\underline{p} = c_0$, then raises the price to $\hat{p}_{\mathrm{r}}^{o}$ once its type is public. The introductory price is below the full-information price: a penetration path.

The signal is a conspicuous loss. Selling at $\underline{p} = c_0 = 0$ while bearing unit cost $c_e = \varphi$, the responsible type serves the whole market ($D_o(0, \overline{1}) = 1$) at a per-period loss

$$\pi_{\mathrm{r}}(c_0, 1) = (c_0 - c_e) \cdot 1 = -\varphi. \tag{26.8}$$

This is Veblen's "conspicuous waste": the responsible type gives up φ per period to prove itself. A conventional firm, already earning zero, would never add a loss on top, so the loss is what makes the signal credible.

IR: when is the loss worth bearing? The responsible type accepts the period-1 loss φ only if the discounted full-information profit it unlocks covers it:

$$\frac{\delta}{1-\delta}\,\hat{\pi}_{\mathrm{r}}^{o} \;\geq\; \varphi \quad\Longleftrightarrow\quad \frac{\delta}{1-\delta}\cdot\frac{(e-\varphi)^2}{4e} \;\geq\; \varphi. \tag{26.9}$$

Solving for δ gives a critical discount factor

$$\tilde{\delta} \;=\; \frac{4e\varphi}{(e+\varphi)^2} \;=\; \frac{4r}{(1+r)^2}, \qquad r \equiv \frac{\varphi}{e} \in (0,1), \tag{26.10}$$

where the second form uses the premium-to-performance ratio $r = \frac{\varphi}{e}$: dividing numerator and denominator of the first form by e^2 leaves $\tilde{\delta}$ a function of r alone. The map $r \mapsto \frac{4r}{(1+r)^2}$ rises on $(0,1)$ and reaches 1 at $r = 1$, that is, at $\varphi = e$; there $\tilde{\delta} = 1$, so no $\delta < 1$ can separate. Desirability ($\varphi < e$, that is, $r < 1$) keeps us strictly below that boundary.

Result 2 (occupied market). Penetration pricing separates the types if and only if $\delta \geq \tilde{\delta}$.

Intuitively, signaling is worthwhile only for a patient entrant, because it must survive a real, out-of-pocket loss during the verification phase before collecting the reward.

26.6 Fourth Step: Untapped Market, Skimming Pricing

IC forces a high price. Now there is no competitor. A conventional type that mimics $\underline{p}$ and is believed responsible earns

$$\pi_{\mathrm{c}}(\underline{p}, 1) = (\underline{p} - c_0)\, D_u(\underline{p}, 1) = \underline{p}\,\frac{1-\underline{p}}{1-e},$$

while honestly it earns the conventional monopoly profit $\hat{\pi}^u_{\mathrm{c}} = \frac{1}{4}$. Mimicry is deterred when

$$\underline{p}\,\frac{1-\underline{p}}{1-e} \;\le\; \frac{1}{4} \quad\Longleftrightarrow\quad \underline{p} \le \frac{1-\sqrt{e}}{2} \;\text{ or }\; \underline{p} \ge \frac{1+\sqrt{e}}{2}. \tag{26.11}$$

The upper edge always lies where demand is $\frac{1-p}{1-e}$, since $\frac{1+\sqrt{e}}{2} \ge \sqrt{e} \ge e$.[1] The responsible type's profit is single-peaked at its full-information price $\hat{p}^u_{\mathrm{r}} = \frac{1+\varphi}{2}$, which lies inside the forbidden interval

$$\left(\frac{1-\sqrt{e}}{2}, \frac{1+\sqrt{e}}{2}\right)$$

precisely when $\varphi < \sqrt{e}$. That inequality is not a separate assumption: it follows from desirability. Because $e \in (0,1)$ we have $\sqrt{e} > e$, so

$$\varphi \;<\; e \;<\; \sqrt{e}, \tag{26.12}$$

and $\varphi < \sqrt{e}$ holds automatically. The optimum therefore always sits inside the interval, and the least costly separating price is its *upper* edge (closer than the lower edge to $\hat{p}^u_{\mathrm{r}}$, since $\varphi > 0$),

$$\underline{p} = \frac{v + c_0}{2} + \frac{(v - c_0)\sqrt{e}}{2} = \frac{1+\sqrt{e}}{2} \;>\; \hat{p}^u_{\mathrm{r}}. \tag{26.13}$$

Result 3 (skimming). In the untapped market the responsible entrant signals by charging the high price $\underline{p} = \frac{1+\sqrt{e}}{2}$, then lowers it to $\hat{p}^u_{\mathrm{r}}$ once its type is public. The introductory price is above the full-information price: a skimming path.

Figure 26.1 depicts the period-1 profit from a price p that consumers take as a sign of responsibility, for the conventional mimic, $\pi_{\mathrm{c}}(p, 1)$, and for the responsible type, $\pi_{\mathrm{r}}(p, 1)$, with $e = \frac{1}{4}$ and $\varphi = \frac{1}{8}$, the values of the sixth step. (For $p < e = \frac{1}{4}$ every consumer buys, so demand is capped at 1 and both profits are linear in p.)

Intuitively, the mimic's profit exceeds its honest profit $\hat{\pi}^u_{\mathrm{c}} = \frac{1}{4}$ on the whole shaded interval, so every price in it invites mimicry, including the responsible type's full-information price $\frac{9}{16}$. The responsible type moves to the nearest price outside the interval, its upper edge $\frac{3}{4}$, where its own profit is still positive.

The signal is forgone sales, not a loss. At the skimming price the responsible type still earns a *positive* margin; it simply serves fewer consumers than at $\hat{p}^u_{\mathrm{r}}$, skimming the cream of the market most willing to pay for responsibility, as in the high introductory prices of Bagwell and Riordan (1991). The high price is one a conventional type does not want to match: its unit cost is lower, so it earns a larger margin on every unit and each sale lost to the high price costs it more, while its honest monopoly profit is already high. Again the signal works by wasting something the mimic values more: here, sales volume.

[1] The lower edge in (26.11) uses the same demand, which is valid when $e \le \frac{1}{4}$, as in the numerical example of the sixth step. When $e > \frac{1}{4}$, demand equals one at the lower edge, the mimic earns $\underline{p}$, and the lower edge becomes $\frac{1}{4}$. The upper edge, which is the one that matters below, is unaffected.

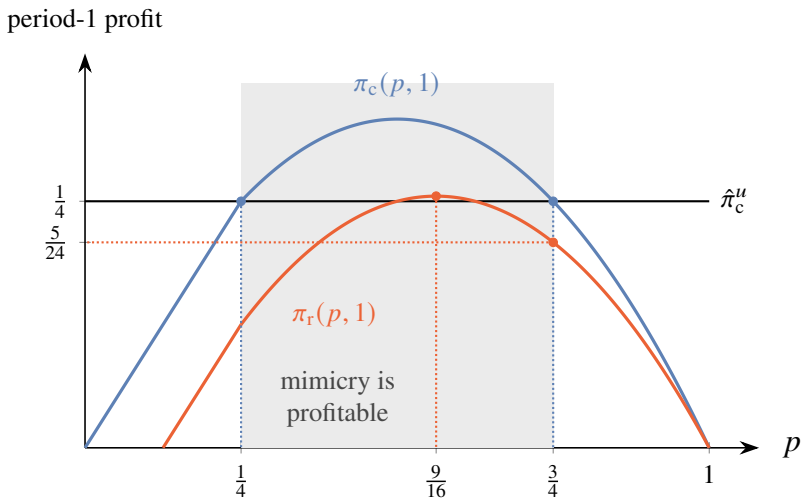


Figure 26.1: Untapped market, period-1 profits when believed responsible, at $e = \frac{1}{4}$ and $\varphi = \frac{1}{8}$. The conventional type mimics any price in the shaded interval, so the responsible type skims at its upper edge, $\underline{p} = \frac{3}{4}$.

IR holds for all δ. Because signaling here means a *reduced* profit rather than an outright loss, the responsible type is willing to signal regardless of how patient it is. Formally, the firm enters and is verified with or without the signal, so the future profit is the same in both cases. In period 1 the skimming price earns

$$\frac{1+\sqrt{e}-2\varphi}{4(1+\sqrt{e})},$$

which exceeds the best profit of a firm believed conventional, $\frac{(1-\varphi)^2}{4}$, if and only if

$$\varphi < \frac{2\sqrt{e}}{1+\sqrt{e}};$$

this holds because

$$\varphi < e \le \frac{2\sqrt{e}}{1+\sqrt{e}}.$$

Result 4 (untapped market). Skimming pricing separates the types for every δ.

Unlike penetration, skimming needs no patience threshold: both types have a profitable future, so the only question is the direction of the period-1 distortion, and IC sends it upward.

26.7 Fifth Step: Comparative Statics and the Verification Lag

The occupied-market threshold (26.10) carries the paper's comparative statics. Differentiating,

$$\frac{\partial\tilde{\delta}}{\partial e} = \frac{4\varphi(\varphi-e)}{(e+\varphi)^3} < 0, \qquad \frac{\partial\tilde{\delta}}{\partial \varphi} = \frac{4e(e-\varphi)}{(e+\varphi)^3} > 0 \qquad (\text{since } \varphi < e). \tag{26.14}$$

Result 5 (comparative statics). Signaling is easier (a lower $\tilde{\delta}$):

- when the product is more differentiated (higher e);
- when the cost premium of responsibility is smaller (lower φ).

A greener product commands a larger future profit worth protecting, and a lower premium shrinks the period-1 loss; both lower the patience requirement. The ratio form (26.10) shows the same thing directly: $\tilde{\delta}$ rises in $r = \frac{\varphi}{e}$, and higher e or lower φ both lower r.

The verification lag. With a general verification length T, the responsible type repeats the penetration price for T periods, so the separation condition (26.9) becomes

$$\frac{\delta^T}{1-\delta}\,\hat{\pi}_{\mathrm{r}}^{o} \geq \varphi\,\frac{1-\delta^T}{1-\delta},$$

which reduces to

$$\delta \;\geq\; \tilde{\delta}^{\,1/T}. \tag{26.15}$$

The threshold $\tilde{\delta}^{1/T}$ rises toward 1 as $T \to \infty$: a longer verification lag makes penetration harder to sustain and demands a more patient entrant.

Result 6 (greenwashing).

- In the occupied market, the entrant cannot separate when $\delta < \tilde{\delta}^{1/T}$, and pooling prices (greenwashing) take over the early periods.
- In the untapped market, the entrant separates for any T and any δ.

Anything that speeds verification (faster certification, active NGOs) lowers the bar for honest signaling; anything that slows it (fake certification, information lobbying) protects greenwashing.

Figure 26.2 depicts the threshold $\tilde{\delta}^{1/T}$ as a function of e, with $\varphi = \frac{1}{8}$ and $T \in \{1, 2, 5\}$.

Table 26.1 summarizes the two scenarios, and Figure 26.3 depicts the two price paths.

26.8 Sixth Step: A Numerical Example

Take

$$v = 1, \qquad c_0 = 0, \qquad e = \frac{1}{4}, \qquad \varphi = \frac{1}{8},$$

so the responsible unit cost is $c_e = \varphi = \frac{1}{8}$. Both scenarios are socially desirable, $\varphi < e$, since $\frac{1}{8} < \frac{1}{4}$; and by (26.12) this already delivers the upward skimming distortion, as $\varphi < e < \sqrt{e}$ reads $\frac{1}{8} < \frac{1}{4} < \frac{1}{2}$. The premium-to-performance ratio is $r = \frac{\varphi}{e} = \frac{1}{2}$.

Occupied market (penetration). From (26.5),

$$\hat{p}_{\mathrm{r}}^{o} = \frac{e+\varphi}{2} = \frac{3}{16}, \qquad \hat{\pi}_{\mathrm{r}}^{o} = \frac{(e-\varphi)^2}{4e} = \frac{1}{64}.$$

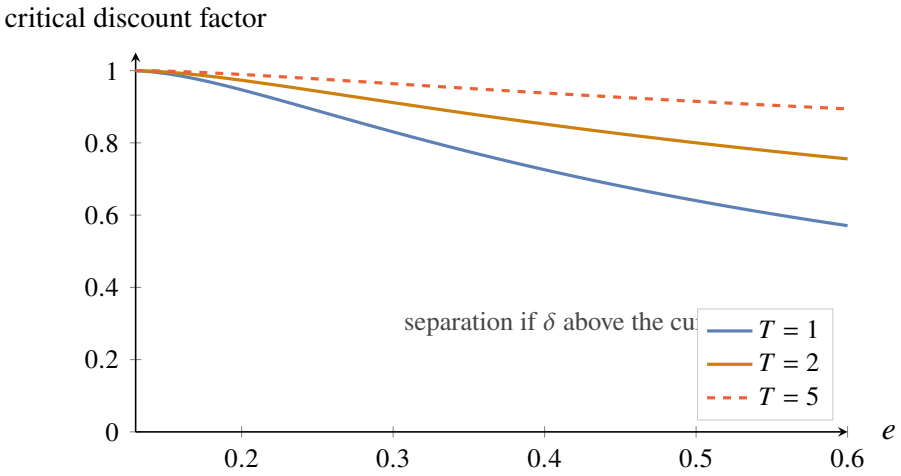


Figure 26.2: The occupied-market threshold $\tilde{\delta}^{1/T} = \left[\frac{4e\varphi}{(e+\varphi)^2}\right]^{1/T}$ against differentiation e, at $\varphi = \frac{1}{8}$, for three verification lengths. It falls with e (more differentiation makes signaling easier) and rises with T (a longer verification lag makes it harder), approaching 1 as $T \to \infty$. Penetration separates only for discount factors above the relevant curve; the untapped market needs no such condition.

	Occupied market (competition)	Untapped market (monopoly)
Signaling strategy	Penetration (low price)	Skimming (high price)
Signaling price $\underline{p}$	c_0	$\frac{v + c_0}{2} + \frac{(v - c_0)\sqrt{e}}{2}$
Bias relative to full information	Downward	Upward
Period-1 sacrifice	A loss, $-\varphi$	Forgone sales (profit still > 0)
Separation requires	$\delta \geq \tilde{\delta} = \frac{4e\varphi}{(e+\varphi)^2}$	every δ
Real-world examples	Costco/Kroger organic; BYD EVs; free samples	Tesla Model S; luxury eco-brands

Table 26.1: The two scenarios of Mahenc (2025). Competition flips the sign of the price signal: a responsible entrant penetrates a market held by conventional incumbents but skims an untapped one.

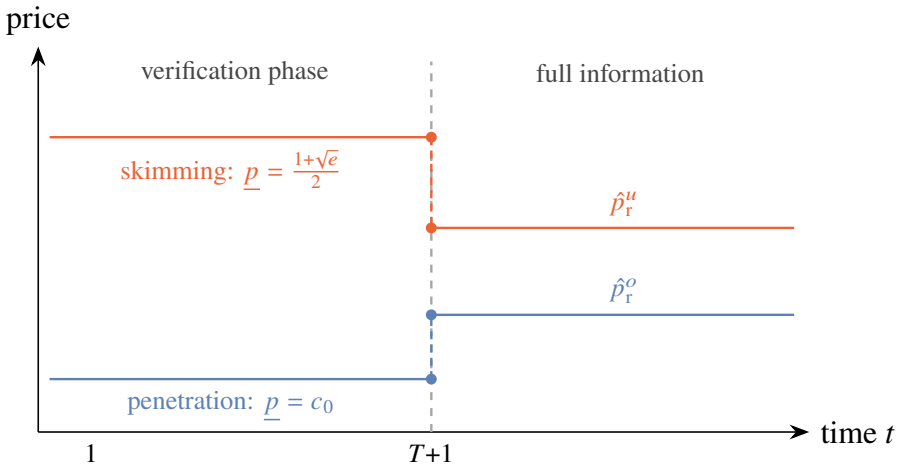


Figure 26.3: The two signaling price paths. In the occupied market the responsible entrant *penetrates*: it holds the low competitive price c_0 through the verification phase, then raises the price to $\hat{p}^o_{\mathrm{r}}$. In the untapped market it *skims*: it holds a high price $\underline{p} = \frac{1+\sqrt{e}}{2}$, then lowers it to $\hat{p}^u_{\mathrm{r}}$. Both distortions shrink to zero once the type is public.

The entrant signals by charging the competitive price $\underline{p} = c_0 = 0$ during verification, taking a loss of $\varphi = \frac{1}{8}$ per period, then raises the price to $\frac{3}{16}$. The critical discount factor, in ratio form with $r = \frac{1}{2}$, is

$$\tilde{\delta} = \frac{4r}{(1+r)^2} = \frac{4 \cdot \frac{1}{2}}{\left(\frac{3}{2}\right)^2} = \frac{2}{\frac{9}{4}} = \frac{8}{9}.$$

So penetration separates the types only if $\delta \geq \frac{8}{9}$: the entrant must be quite patient.

Untapped market (skimming). From (26.4) and (26.13),

$$\hat{p}^u_{\mathrm{r}} = \frac{1+\varphi}{2} = \frac{9}{16}, \qquad \underline{p} = \frac{1+\sqrt{e}}{2} = \frac{1+\frac{1}{2}}{2} = \frac{3}{4}.$$

The entrant signals by charging the high price $\frac{3}{4}$ during verification, then lowers it to $\frac{9}{16}$. Since $\frac{3}{4} > \frac{9}{16}$, the introductory price is biased *upward*. A conventional mimic would earn

$$\underline{p}\,\frac{1-\underline{p}}{1-e} = \frac{3}{4} \cdot \frac{1-\frac{3}{4}}{1-\frac{1}{4}} = \frac{1}{4},$$

equal to its honest monopoly profit $\hat{\pi}^u_{\mathrm{c}} = \frac{1}{4}$: it is just deterred. Skimming separates for every δ.

The same performance level $e = \frac{1}{4}$ and premium $\varphi = \frac{1}{8}$ thus produce *opposite* price signals (a rock-bottom 0 under competition, a premium $\frac{3}{4}$ under monopoly), purely because of the competitive environment.

26.9 Seventh Step: What the Toy Model Keeps and Drops

Mahenc's main messages all survive the simplification:

1. Competition flips the signal. Penetration (a low, rising price) in the occupied market; skimming (a high, falling price) in the untapped market (Results 1 to 4).
2. Signaling is conspicuous waste. The responsible type sacrifices profit during verification (an out-of-pocket loss under competition, forgone sales under monopoly), and the mimic cannot afford to copy it.
3. Patience matters only under competition. Penetration needs $\delta \geq \tilde{\delta}$; skimming works for all δ.
4. Comparative statics. Signaling is easier with more differentiation and a smaller cost premium, and harder the longer verification takes, with greenwashing filling the gap when the entrant is too impatient (Results 5 and 6).

What the toy model drops, relative to Mahenc (2025): (i) the general verification length T is collapsed to $T = 1$, so the trigger-strategy and never-dissuaded-once-convinced machinery that sustains repeated separation is replaced by a single period (Section 26.7 restores the $\tilde{\delta}^{1/T}$ scaling in one line); (ii) the multi-period pooling/greenwashing phase and its timing are only sketched, not derived; and (iii) the equilibrium-selection argument (Intuitive Criterion, Riley outcome) is stated rather than proved. Each can be reinstated at the cost of the dynamic bookkeeping that makes the full paper dense. For teaching, the two-period version already delivers the penetration-versus-skimming dichotomy and the Veblen logic behind it.

26.10 Connections to Other Chapters

- Chapter 25 studies price signaling of a clean product in a single period.
- Chapter 24 lets the firm choose between signaling through its price and disclosing its quality at a cost.
- Chapter 28 studies certification instead of the price: a noisy test lets sellers of green products obtain a label.

References

Bagwell, K., and M. H. Riordan (1991). "High and Declining Prices Signal Product Quality." *American Economic Review*, 81(1), pp. 224–239.

Cho, I.-K., and D. M. Kreps (1987). "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics*, 102(2), pp. 179–221.

Mahenc, P. (2025). "Penetration or Skimming Pricing for Credence Products?" *Journal of Economic Dynamics and Control*, 180, 105196. https://doi.org/10.1016/j.jedc.2025.105196

Chapter 27

Label Confusion and the Groucho Effect

Companion note to Harbaugh, R., J. W. Maxwell, and B. Roussillon (2011), "Label Confusion: The Groucho Effect of Uncertain Standards," Management Science 57(9): 1512–1527. Full reference at the end.

27.1 Introduction

Supermarket shelves are full of eco-labels. Consider a shopper who sees a "green" label on a bottle of detergent. The label certifies that the product meets a pass-fail standard, but the shopper does not know how demanding that standard is. What does the label tell her?

The label carries news about two unknowns at once. It says that the product is at least as good as the standard, but also that the standard is no harder than the product. Harbaugh et al. (2011) call the second inference the *Groucho effect*, after Groucho Marx's quip that he would not join any club willing to accept him as a member. When a product with a poor reputation displays a label, consumers conclude that the label must be easy to get. In the same way, when a product with a good reputation has no label, consumers conclude that the standard must be hard, a *reverse Groucho effect*. Both effects make a label weaker news about quality.

The paper shows that uncertain standards make labels less informative on average, so firms have a weaker reason to certify. Labeling is then harder to sustain as an equilibrium, and an equilibrium in which no firm labels is easier to sustain. Adding more labels makes things worse: as labels multiply, an equilibrium without labeling exists for any certification cost, and the equilibrium with labeling stops conveying information. A single focal label that all firms use restores information.

The toy model draws the standard uniformly on the unit interval, the paper's leading example, so a firm of quality q passes with probability q. The first step sets up the benchmark with a known standard, the second step derives the Groucho effect, and the third step measures how much a label reveals. The fourth step finds the labeling and non-labeling equilibria, the fifth step adds firm reputations, the sixth and seventh steps study label proliferation and several firms, and the eighth step lists what the toy model keeps and drops.

27.2 The Toy Model

Quality. A firm sells a product of environmental quality $Q \in [0,1]$, drawn from a cumulative distribution F with full support, mean $\mu \equiv E[Q]$, and variance $\sigma^2 \equiv \mathrm{Var}(Q)$. The firm observes the realization q; consumers know only F. Quality is fixed, so the firm makes no investment in it.

The label. A single eco-label has a standard $S \in [0,1]$, drawn uniformly and independently of Q. The firm observes the realization s. If $q \geq s$ the firm may certify at a fixed cost $c > 0$ (certification fees, testing, documentation); if $q < s$ it cannot certify. The standard is *certain* if consumers also observe s, and *uncertain* if they know only that S is uniform.

Payoffs. The firm earns consumers' expected quality given what they observe, minus c if it certifies. (The paper notes that its results hold whenever the firm's payoff is a strictly increasing function of expected quality.)

The key normalization. Because S is uniform, a firm of quality q meets the standard with probability

$$\Pr(S \leq q) = q. \tag{27.1}$$

Figure 27.1 depicts (27.1). Panel (a) plots the density of the standard, $g(s) = 1$ on $[0,1]$, and shades its part to the left of a quality level q: the shaded area is the probability that the standard is no harder than q. Panel (b) plots this probability against q, which gives the 45-degree line: a firm of higher quality is more likely to meet the standard, and a firm of quality 1 meets it for sure. A label therefore reweights each quality level q by its pass probability q.

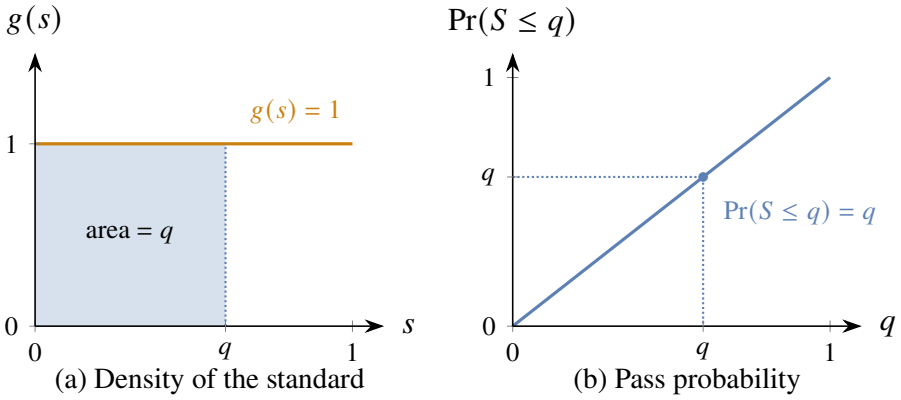


Figure 27.1: Probability that a uniform standard is no harder than quality q (shown for $q = 0.6$).

The numerical examples below also take F uniform, so that $\mu = \frac{1}{2}$, $E[Q^2] = \frac{1}{3}$, and $\sigma^2 = \frac{1}{12}$.

27.3 First Step: The Benchmark with a Known Standard

Consider first a certain standard s, and assume that every firm able to certify does so. Consumers who see a label learn $Q \geq s$, and consumers who see no label learn $Q < s$. Write $\overline{q}(s) \equiv E[Q \mid Q \geq s]$ and $\underline{q}(s) \equiv E[Q \mid Q < s]$. Figure 27.2 depicts the two events on the quality interval $[0,1]$

for a given s. With uniform F, quality is spread evenly over each event, so the expected quality in an event is the midpoint of its interval: $\underline{q}(s)$ is the midpoint of $[0, s]$ and $\overline{q}(s)$ is the midpoint of $[s, 1]$.

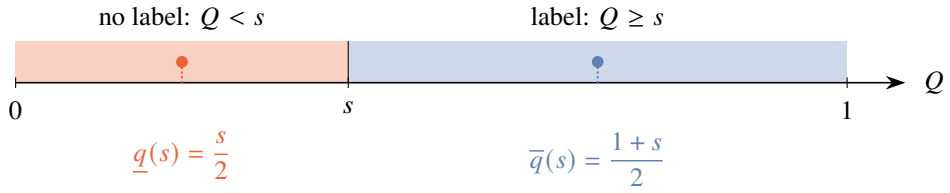


Figure 27.2: Expected quality given $Q < s$ and given $Q \geq s$, with uniform quality (shown for $s = 0.4$).

In formulas,

$$\overline{q}(s) = \frac{1+s}{2}, \qquad \underline{q}(s) = \frac{s}{2}. \tag{27.2}$$

Averaging over the possible standards,

$$E\left[\overline{q}(S)\right] = \int_0^1 \frac{1+s}{2}\, ds = \frac{3}{4}, \qquad E\left[\underline{q}(S)\right] = \int_0^1 \frac{s}{2}\, ds = \frac{1}{4}. \tag{27.3}$$

Figure 27.3 depicts the two integrals in (27.3): the shaded area under $\overline{q}(s)$ in panel (a) equals $\frac{3}{4}$, and the shaded area under $\underline{q}(s)$ in panel (b) equals $\frac{1}{4}$. Because the standard has density 1 on $[0, 1]$, each area is the average of the line over the standards.

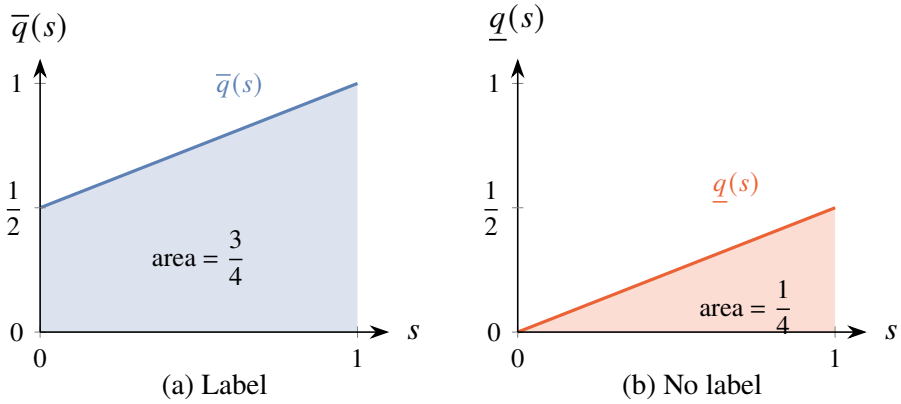


Figure 27.3: Expected quality given a known standard s, with uniform quality; shaded areas are $E[\overline{q}(S)]$ and $E[\underline{q}(S)]$.

These are the posterior means consumers would form if the conditional distribution of the standard did not change with the firm's labeling decision. The second step shows that it does change.

27.4 Second Step: Joint Updating and the Groucho Effect

When the standard is uncertain, a label tells consumers only that $Q \geq S$ for some unobserved S. By (27.1), the expected quality of a labeled product is

$$E[Q \mid Q \geq S] = \frac{E\big[Q \Pr(S \leq Q \mid Q)\big]}{\Pr(S \leq Q)} = \frac{E[Q \cdot Q]}{E[Q]} = \frac{E[Q^2]}{\mu}, \tag{27.4}$$

and, by the same steps, the expected quality of an unlabeled product is

$$E[Q \mid Q < S] = \frac{E\big[Q(1-Q)\big]}{1-\mu} = \frac{\mu - E[Q^2]}{1-\mu}. \tag{27.5}$$

With uniform F, $E[Q^2] = \frac{1}{3}$ and $\mu = \frac{1}{2}$, so these equal $\frac{2}{3}$ and $\frac{1}{3}$. Figure 27.4 depicts the reweighting behind the two means. A label multiplies the prior density $f(q)$ by the pass probability q and divides by μ, the probability of a label, which gives the density $2q$; no label multiplies it by $1-q$ and divides by $1-\mu$, which gives $2(1-q)$. The first density puts more weight on high qualities, so its mean lies above the prior mean $\frac{1}{2}$, and the second puts more weight on low qualities, so its mean lies below it.

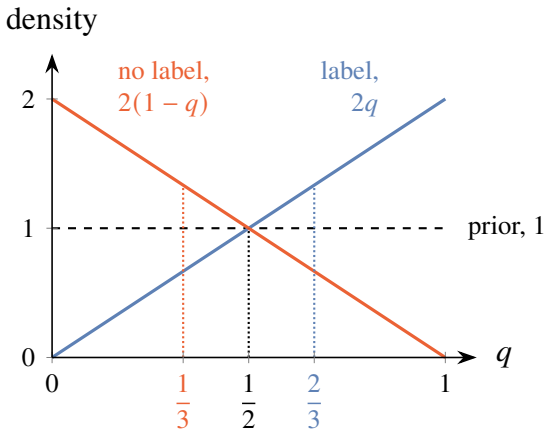


Figure 27.4: Density of quality before the label decision is observed (prior) and after a label or no label, with uniform quality and a uniform standard.

Consumers also revise their estimate of the standard. Conditional on a label, the density of S is

$$g(s \mid Q \geq S) = \frac{1 - F(s)}{\mu}, \tag{27.6}$$

which is decreasing in s: easy standards are over-represented among labeled products, because more firms clear them. Integrating,

$$E[S \mid Q \geq S] = \frac{E[Q^2]}{2\mu} = \frac{1}{3}, \qquad E[S \mid Q < S] = \frac{1 - E[Q^2]}{2(1-\mu)} = \frac{2}{3}, \tag{27.7}$$

where the second equality in each expression uses uniform F. Meeting the standard is good news about quality and bad news about the standard: the estimate of S falls from $\frac{1}{2}$ to $\frac{1}{3}$. Harbaugh et

al. (2011) call this downward revision the *Groucho effect*. Failing to meet the standard works the other way: the estimate of S rises to $\frac{2}{3}$, the *reverse Groucho effect*.

Figure 27.5 depicts the revision of the standard. Conditional on a label, the density of S is $2(1-s)$, from (27.6); conditional on no label, it is $\frac{F(s)}{1-\mu} = 2s$, which is increasing in s because hard standards are over-represented among unlabeled products. The means of the two densities are the estimates $\frac{1}{3}$ and $\frac{2}{3}$ in (27.7).

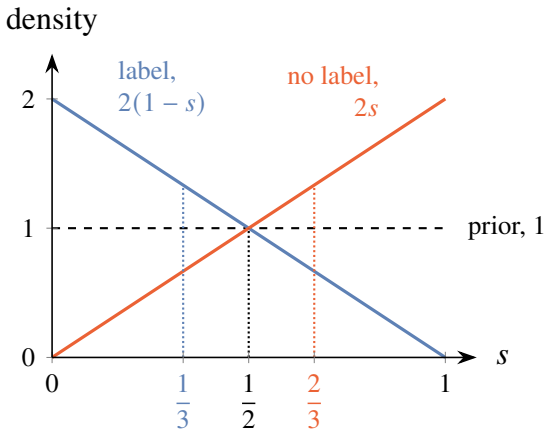


Figure 27.5: Density of the standard before the label decision is observed (prior) and after a label or no label, with uniform quality and a uniform standard.

Figure 27.6 depicts the two cases in the unit square, with the standard on the horizontal axis and quality on the vertical axis, for uniform F. In panel (a), with an uncertain standard, the prior point $\left(\frac{1}{2}, \frac{1}{2}\right)$ moves to the centroid of the upper triangle $\{Q \geq S\}$ (shaded blue) after a label and to the centroid of the lower triangle $\{Q < S\}$ (shaded red) without one; both moves are diagonal, because the label shifts the estimates of Q and S in opposite directions. In panel (b), with a known standard s, the estimate of s cannot move, so the full update falls on quality: the posterior mean jumps to the upper line $\overline{q}(s)$ or the lower line $\underline{q}(s)$ of (27.2).

Comparing the two panels of Figure 27.6 with the averages in (27.3),

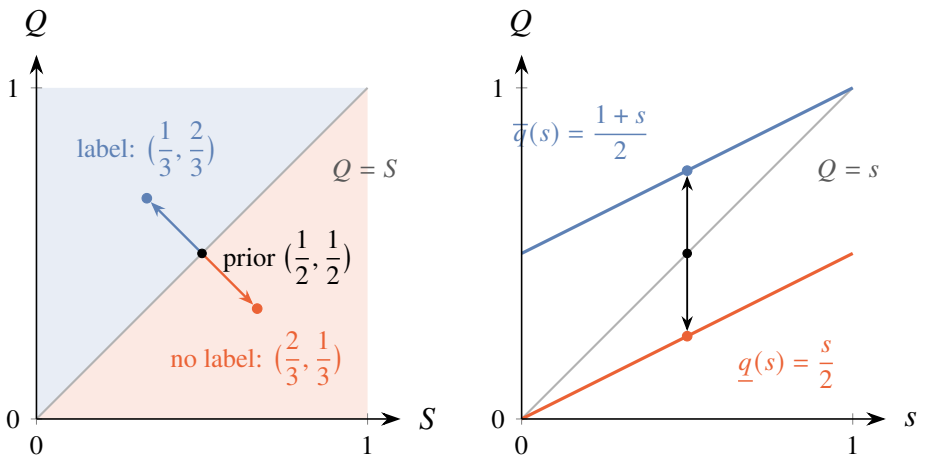
$$\underbrace{E\left[\underline{q}(S)\right]}_{\frac{1}{4}} < \underbrace{E[Q \mid Q < S]}_{\frac{1}{3}} < \underbrace{\mu}_{\frac{1}{2}} < \underbrace{E[Q \mid Q \geq S]}_{\frac{2}{3}} < \underbrace{E\left[\overline{q}(S)\right]}_{\frac{3}{4}}. \tag{27.8}$$

Result 1 (moderation). Compared with a known standard, with an uncertain standard:

- a label is weaker good news;
- the absence of a label is weaker bad news.

The paper proves the ranking (27.8) for any F and G (inequality (6) of the paper); the uniform numbers make the size of the gap visible.

Figure 27.7 places the five numbers of (27.8) on the unit interval.



(a) Uncertain standard

(b) Known standard $S = s$

Figure 27.6: Updated estimates of quality and of the standard after a label and after no label, with uniform quality and a uniform standard.

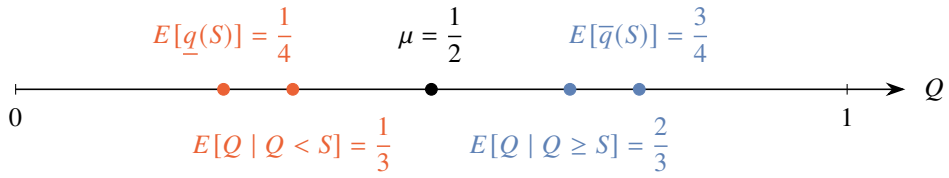


Figure 27.7: Expected quality after a label or no label, with a known and with an uncertain standard, for uniform quality and a uniform standard.

27.5 Third Step: How Much Does a Label Reveal?

Measure the information in a label by the share of the prior variance of quality that it removes,

$$R^2 \equiv 1 - \frac{E\left[(Q - \widehat{Q})^2\right]}{\sigma^2}, \tag{27.9}$$

where $\widehat{Q}$ is consumers' posterior mean. The numerator is the expected mean-squared error (MSE) that the paper uses, so a higher R^2 is a lower MSE. Without a label the MSE is $\sigma^2 = \frac{1}{12}$. With a known standard s the MSE equals

$$\frac{s^3 + (1-s)^3}{12},$$

whose average over s is $\frac{1}{24}$, so $R^2 = \frac{1}{2}$. With an uncertain standard the two posterior means are $\frac{2}{3}$ and $\frac{1}{3}$, each reached with probability $\frac{1}{2}$, so the MSE is

$$\frac{1}{3} - \left(\frac{1}{2} \cdot \frac{4}{9} + \frac{1}{2} \cdot \frac{1}{9}\right) = \frac{1}{18},$$

and $R^2 = \frac{1}{3}$.

Result 2 (Proposition 1 of the paper). A label is less informative on average when its standard is uncertain.

A sum-of-squares identity shows why: the expected MSE with an uncertain standard minus the expected MSE with a known standard equals

$$\int_0^1 \left[F(s)\big(E[Q \mid Q < S] - \underline{q}(s)\big)^2 + \big(1 - F(s)\big)\big(E[Q \mid Q \geq S] - \overline{q}(s)\big)^2\right] dG(s) \;>\; 0, \tag{27.10}$$

because consumers who do not know s must use a single posterior mean in each event, while consumers who know s use the best mean for each s. With uniform F and G the integral equals $\frac{1}{18} - \frac{1}{24} = \frac{1}{72}$. Figure 27.8 depicts the expected MSE against the standard s for consumers who know s (blue curve) and for consumers who use the posterior means $\frac{2}{3}$ and $\frac{1}{3}$ regardless of s (red curve). The shaded area between the curves is the integral in (27.10), and the dotted lines mark the two averages $\frac{1}{24}$ and $\frac{1}{18}$. Consumers who know s place the posterior means at the midpoints of the two events, which minimizes the error at every s, so the blue curve lies below the red one everywhere. The red curve rises above the dashed line $\frac{1}{12}$, the MSE without a label, for standards below about 0.09 and above about 0.91, where consumers who do not know s misread the label; on average over s, the label still lowers the MSE from $\frac{1}{12}$ to $\frac{1}{18}$.

Table 27.1 reports the uniform values.

27.6 Fourth Step: Labeling and Non-Labeling Equilibria

The equilibrium concept is perfect Bayesian equilibrium with one refinement, described below. In a *labeling equilibrium* every firm that meets the standard certifies, so the lack of a label reveals

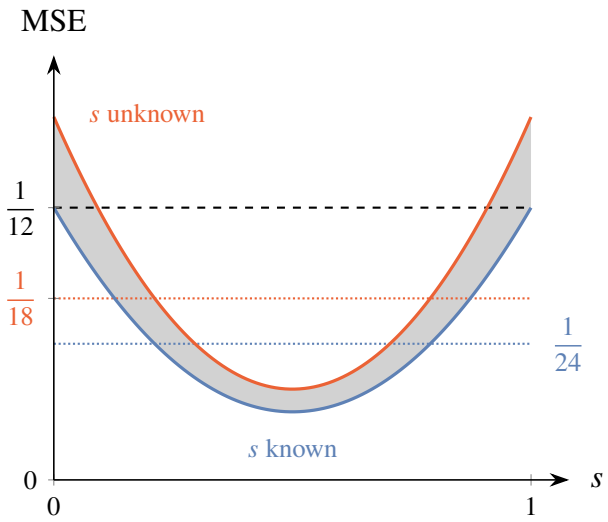


Figure 27.8: Expected MSE of consumers' estimate of quality against the standard s, with uniform quality; the shaded area is the integral in (27.10), and the dashed line is the MSE without a label.

Standard	$E[Q \mid \text{label}]$	$E[Q \mid \text{no label}]$	$E[S \mid \text{label}]$	$E[S \mid \text{no label}]$	R^2
Known (average over s)	$\frac{3}{4}$	$\frac{1}{4}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
Uncertain	$\frac{2}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{1}{3}$

Table 27.1: Posterior means and informativeness when every firm able to certify does so, with uniform quality and a uniform standard.

failure. The firm certifies if the gain is at least the cost,

$$E[Q \mid Q \geq S] - E[Q \mid Q < S] \;\geq\; c. \tag{27.11}$$

In a *non-labeling equilibrium* no firm certifies, so the absence of a label is no news and the firm earns μ. A label is then off the equilibrium path, and the paper assumes that consumers attribute an unexpected label to any type that meets the standard with equal likelihood, so a firm that deviates earns $E[Q \mid Q \geq S]$. No firm deviates if

$$E[Q \mid Q \geq S] - \mu \;\leq\; c. \tag{27.12}$$

Closed forms. Substituting (27.4) and (27.5) into (27.11),

$$\begin{aligned} \overline{c} \;\equiv\; E[Q \mid Q \geq S] - E[Q \mid Q < S] &= \frac{E[Q^2]}{\mu} - \frac{\mu - E[Q^2]}{1-\mu} \\ &= \frac{E[Q^2](1-\mu) - \mu\big(\mu - E[Q^2]\big)}{\mu(1-\mu)} \\ &= \frac{E[Q^2] - \mu^2}{\mu(1-\mu)} = \frac{\sigma^2}{\mu(1-\mu)}, \end{aligned} \tag{27.13}$$

and, from (27.12),

$$\underline{c} \;\equiv\; E[Q \mid Q \geq S] - \mu = \frac{E[Q^2] - \mu^2}{\mu} = \frac{\sigma^2}{\mu} = (1-\mu)\,\overline{c}. \tag{27.14}$$

With a uniform standard, the labeling cutoff is the variance of quality relative to its largest possible value given the mean, $\mu(1-\mu)$. With uniform F, $\underline{c} = \frac{1}{6}$ and $\overline{c} = \frac{1}{3}$.

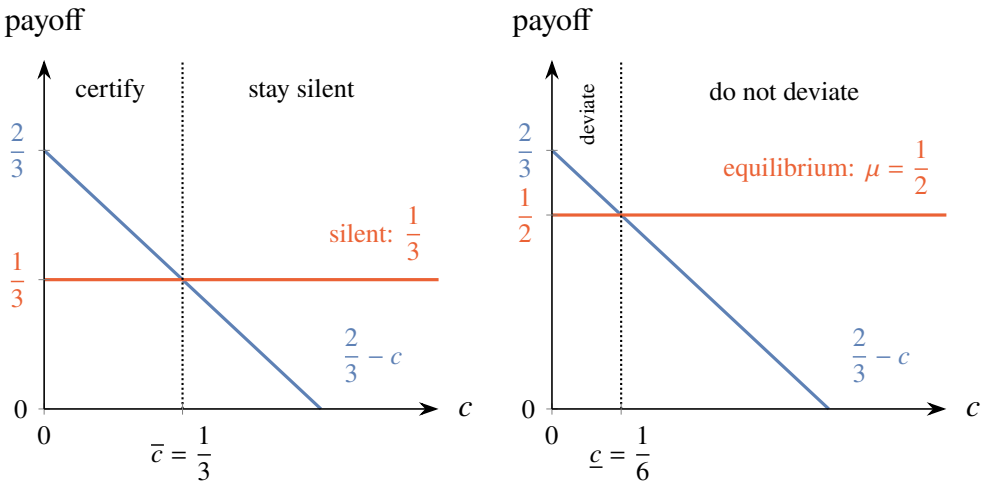
Figure 27.9 depicts the two incentive conditions for uniform F, with the certification cost on the horizontal axis. In panel (a), the labeling equilibrium, a firm that meets the standard compares the payoff from certifying, $E[Q \mid Q \geq S] - c = \frac{2}{3} - c$ (blue line), with the payoff from staying silent, $E[Q \mid Q < S] = \frac{1}{3}$ (red line): certifying is profitable up to $c = \frac{1}{3} = \overline{c}$. In panel (b), the non-labeling equilibrium, a firm compares the payoff from deviating to a label, $\frac{2}{3} - c$, with the equilibrium payoff $\mu = \frac{1}{2}$: deviating is unprofitable from $c = \frac{1}{6} = \underline{c}$ on.

Result 3 (Proposition 2 of the paper).

- A labeling equilibrium exists if and only if $c \leq \overline{c}$.
- A non-labeling equilibrium exists if and only if $c \geq \underline{c}$.
- Both exist if and only if $\underline{c} \leq c \leq \overline{c}$.

Since $\underline{c} = (1-\mu)\overline{c} < \overline{c}$, the interval of multiple equilibria is never empty. The two cutoffs split the cost axis into three regions, which the figures below name by the equilibria they support:

- L: $c < \underline{c}$, only the labeling equilibrium exists;
- LN: $\underline{c} \leq c \leq \overline{c}$, both equilibria exist;



(a) Labeling equilibrium (b) Non-labeling equilibrium

Figure 27.9: A firm's payoff from each action against the certification cost c, with uniform quality and a uniform standard.

- N: $c > \overline{c}$, only the non-labeling equilibrium exists.

The known-standard benchmark. With a known standard s and uniform F, (27.2) gives

$$\overline{c}(s) = \overline{q}(s) - \underline{q}(s) = \frac{1}{2}, \qquad \underline{c}(s) = \overline{q}(s) - \mu = \frac{s}{2}, \tag{27.15}$$

so a labeling equilibrium exists for $c \leq \frac{1}{2}$ regardless of the standard, and a non-labeling equilibrium exists for $c \geq \frac{s}{2}$, a cutoff whose average over s is $\frac{1}{4}$. Figure 27.10 depicts the three regions for a known standard, with the standard s on the horizontal axis and the cost on the vertical axis. The labeling cutoff $\overline{c}(s) = \frac{1}{2}$ is flat, and the non-labeling cutoff $\underline{c}(s) = \frac{s}{2}$ rises with s, so the region of multiple equilibria LN is the wedge between them; the dotted line marks the average $\frac{1}{4}$ of the rising cutoff.

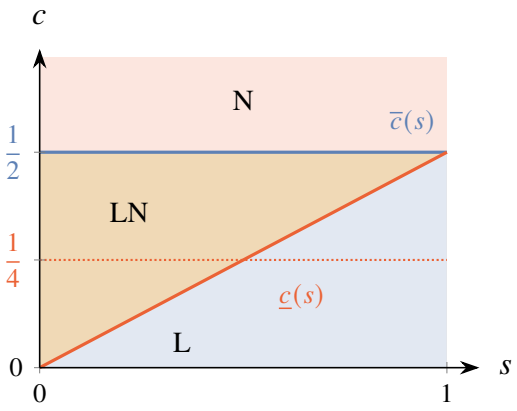


Figure 27.10: Labeling (L), multiple-equilibria (LN), and non-labeling (N) regions with a known standard s, for uniform quality.

Result 4 (Proposition 3 of the paper). Uncertainty about the standard:

- shrinks the range of costs that supports labeling;
- enlarges the range of costs that supports non-labeling.

Figure 27.11 depicts the two ranges on a cost axis for uniform quality: labeling survives up to $c = \frac{1}{3}$ instead of $\frac{1}{2}$, and non-labeling starts at $c = \frac{1}{6}$ instead of $\frac{1}{4}$.

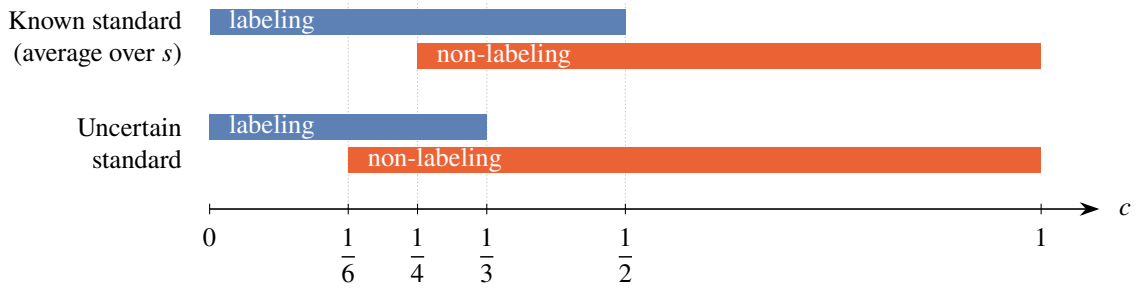


Figure 27.11: Costs that support a labeling equilibrium and a non-labeling equilibrium, with uniform quality and a uniform standard.

Intuitively, the Groucho effect makes a label less valuable, which hurts the labeling equilibrium, and the reverse Groucho effect makes the lack of a label less damaging, which helps the non-labeling equilibrium.

27.7 Fifth Step: Reputations

Let quality follow a Beta distribution with parameters (α, β), so that

$$\mu = \frac{\alpha}{\alpha + \beta}, \qquad \sigma^2 = \frac{\mu(1-\mu)}{\alpha + \beta + 1}.$$

Substituting into (27.13) and (27.14),

$$\overline{c} = \frac{1}{\alpha + \beta + 1}, \qquad \underline{c} = \frac{\beta}{(\alpha + \beta)(\alpha + \beta + 1)}. \tag{27.16}$$

The paper models a firm with a *bad reputation* as $Q \sim \text{Beta}(1, b)$ with $b \geq 1$, so that $\mu = \frac{1}{1+b} \leq \frac{1}{2}$, and a firm with a *good reputation* as $Q \sim \text{Beta}(a, 1)$ with $a \geq 1$, so that $\mu \geq \frac{1}{2}$. Writing each family in terms of its mean, (27.16) becomes

$$\overline{c} = \begin{cases} \dfrac{\mu}{1+\mu} & \text{if } \mu \leq \dfrac{1}{2} \text{ (bad reputation)}, \\[2ex] \dfrac{1-\mu}{2-\mu} & \text{if } \mu \geq \dfrac{1}{2} \text{ (good reputation)}, \end{cases} \qquad \underline{c} = (1-\mu)\,\overline{c}. \tag{27.17}$$

Figure 27.12 depicts the two families for $b = a = 3$, together with the uniform density. A bad reputation piles the density of quality near $q = 0$ (mean $\frac{1}{4}$), and a good reputation piles it near $q = 1$ (mean $\frac{3}{4}$).

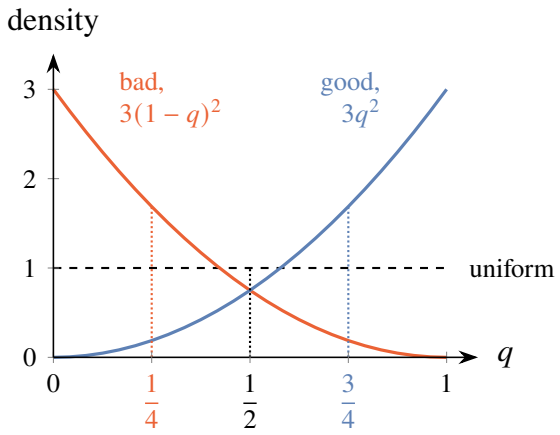


Figure 27.12: Density of quality for a bad reputation, Beta$(1, 3)$, a good reputation, Beta$(3, 1)$, and the uniform distribution.

Figure 27.13 depicts the regions L, LN, and N of the fourth step against the prior mean μ, with the certification cost on the vertical axis: panel (a) for an uncertain standard, from (27.17), and panel (b) for a known standard, where the cutoffs are averaged over s. In panel (a) the labeling cutoff $\overline{c}$ peaks at $\mu = \frac{1}{2}$, where $\overline{c} = \frac{1}{3}$, and falls to zero as the reputation becomes very bad or very good; the non-labeling cutoff $\underline{c}$, below which labeling is the only equilibrium, peaks at $\mu = \sqrt{2} - 1$, where $\underline{c} = 3 - 2\sqrt{2}$. Intuitively, consumers who see a label on a product with a bad reputation conclude that the standard must be easy: in the bad-reputation family,

$$E[S \mid Q \geq S] = \frac{E[Q^2]}{2\mu} = \frac{1}{b+2}, \tag{27.18}$$

which falls to zero as b grows, so the label is worth little. For a product with a good reputation the reverse Groucho effect dominates: consumers who see no label conclude that the standard must be hard, since in the good-reputation family

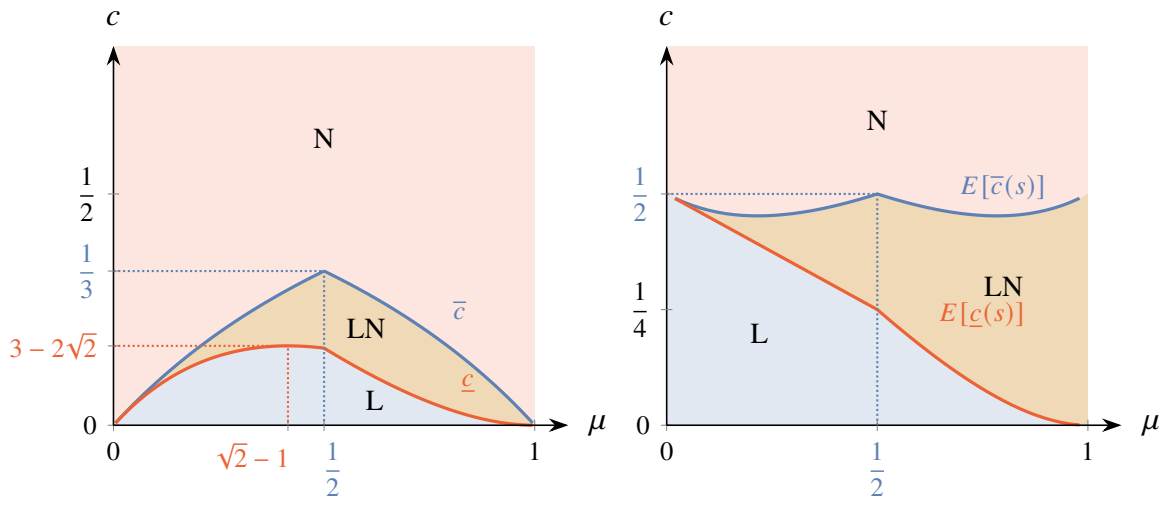
$$E[S \mid Q < S] = \frac{1 - E[Q^2]}{2(1-\mu)} = \frac{a+1}{a+2}, \tag{27.19}$$

which rises to one as a grows, so the firm loses little by not certifying. In panel (b), with a known standard, a firm with a bad reputation that is in fact of high quality can prove it, and the labeling cutoff stays above 0.45 for every reputation.

Figure 27.14 depicts the posterior estimates of the standard in (27.18) and (27.19) for every mean μ, from the same two families. In terms of μ,

$$E[S \mid Q \geq S] = \begin{cases} \dfrac{\mu}{1+\mu} & \text{if } \mu \leq \dfrac{1}{2}, \\[2ex] \dfrac{1}{2(2-\mu)} & \text{if } \mu \geq \dfrac{1}{2}, \end{cases} \qquad E[S \mid Q < S] = \begin{cases} \dfrac{1+2\mu}{2(1+\mu)} & \text{if } \mu \leq \dfrac{1}{2}, \\[2ex] \dfrac{1}{2-\mu} & \text{if } \mu \geq \dfrac{1}{2}. \end{cases} \tag{27.20}$$

A label lowers the estimate of the standard below the prior value $\frac{1}{2}$ (blue curve), most strongly when the reputation is bad; the lack of a label raises it above $\frac{1}{2}$ (red curve), most strongly when the reputation is good.



(a) Uncertain standard (b) Known standard

Figure 27.13: Labeling (L), non-labeling (N), and multiple-equilibria (LN) regions against the prior mean μ, with a uniform standard; $Q \sim \text{Beta}(1, b)$ for $\mu \leq \frac{1}{2}$ and $Q \sim \text{Beta}(a, 1)$ for $\mu \geq \frac{1}{2}$.

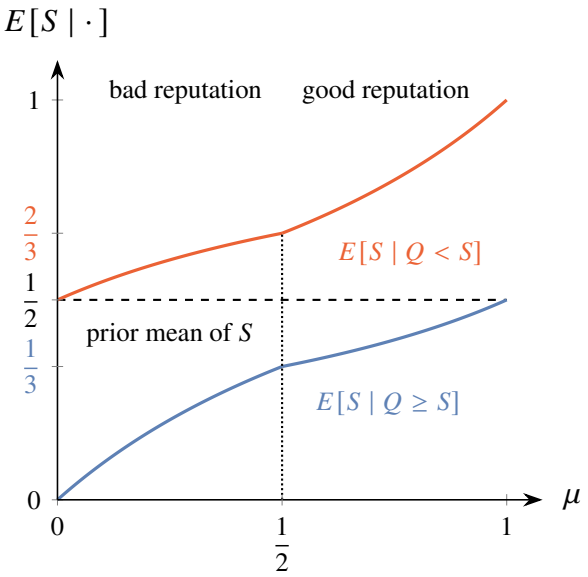


Figure 27.14: Expected standard after a label and after no label against the prior mean μ, with a uniform standard; $Q \sim \text{Beta}(1, b)$ for $\mu \leq \frac{1}{2}$ and $Q \sim \text{Beta}(a, 1)$ for $\mu \geq \frac{1}{2}$.

What drives the hump. Equation (27.16) shows that, with a uniform standard, the labeling cutoff $\overline{c} = \frac{1}{\alpha+\beta+1}$ depends only on how concentrated the prior is, measured by $\alpha + \beta$, and not on where it is centered. In the two families above, moving the mean toward 0 or 1 also raises $\alpha + \beta$ (to $1 + b$ or $a + 1$), and that is what shrinks the labeling region at both ends of Figure 27.13(a). The paper points to this in its footnote 6, where it notes that these families are most diffuse at $\mu = \frac{1}{2}$. A firm with mean $\frac{1}{4}$ and Beta$\left(\frac{1}{2}, \frac{3}{2}\right)$ quality, for instance, has the same $\overline{c} = \frac{1}{3}$ as a firm with uniform quality. The strength of the Groucho effect does depend on the mean: for any Beta prior,

$$E[S \mid Q \geq S] = \frac{\alpha + 1}{2(\alpha + \beta + 1)},$$

which, for a given concentration $\alpha+\beta$, is lower the worse the reputation. A label is less valuable for a product with a bad reputation because it signals an easy standard, and also because a concentrated prior leaves little to learn.

27.8 Sixth Step: Label Proliferation and Focal Labels

Now let there be $n \geq 1$ labels with standards drawn independently from the uniform distribution and the same cost c, and let the firm display at most one label. Consider the *symmetric* strategy: the firm adopts the toughest label it meets. Consumers cannot tell which label is toughest, and the labels are interchangeable to them, so a label tells consumers only that the firm meets at least the *easiest* standard, $S_{1:n} \equiv \min\{S_1, \ldots, S_n\}$. A firm of quality q meets it with probability

$$\Pr(S_{1:n} \leq q) = 1 - (1 - q)^n. \tag{27.21}$$

Using (27.21) in place of (27.1) in the steps behind (27.4) and (27.5), and taking uniform F,

$$E[Q \mid Q \geq S_{1:n}] = \frac{n + 3}{2(n + 2)}, \qquad E[Q \mid Q < S_{1:n}] = \frac{1}{n + 2}. \tag{27.22}$$

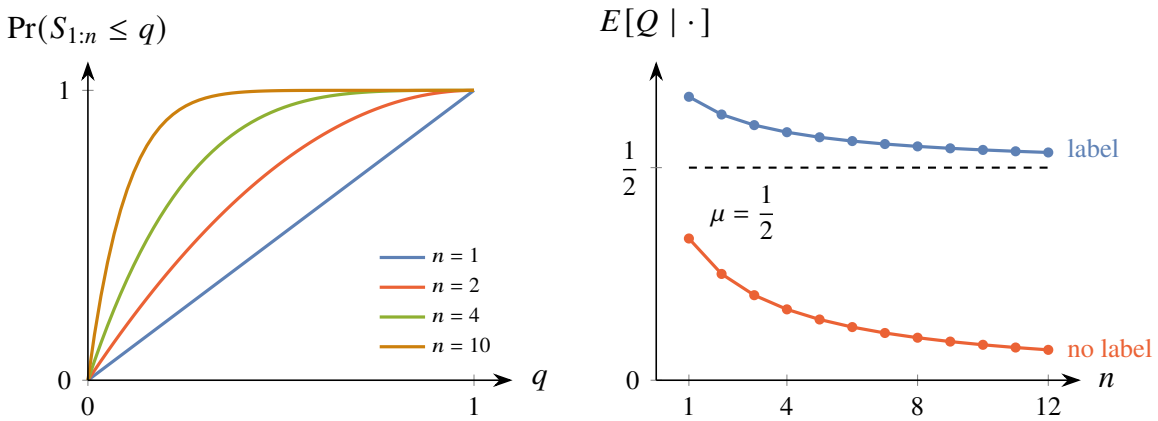
Figure 27.15 depicts (27.21) in panel (a) and (27.22) in panel (b). Panel (a) shows that, for any given quality, the probability of meeting the easiest standard rises with n, so a label becomes easier to get; panel (b) shows that both posterior means fall with n, toward $\mu = \frac{1}{2}$ and toward 0, so failing every label becomes a sign of very low quality.

The two cutoffs are

$$\overline{c}_n = \frac{n + 1}{2(n + 2)}, \qquad \underline{c}_n = \frac{1}{2(n + 2)}. \tag{27.23}$$

Result 5 (Proposition 4 of the paper). With uncertain standards, adding labels enlarges the range of costs that supports the non-labeling equilibrium. As n grows:

- non-labeling becomes an equilibrium for every $c > 0$;
- labeling is an equilibrium for every $c < \mu$ once n is large enough;
- the labeling equilibrium becomes uninformative.



(a) Probability of a label (b) Posterior means

Figure 27.15: Probability of meeting the easiest of n uniform standards, and expected quality given a label and given no label, with uniform quality.

In the uniform case the cutoffs in (27.23) hold for every n, not only in the limit: for a given cost c, a non-labeling equilibrium exists if and only if

$$n \geq \frac{1}{2c} - 2,$$

and, for $c < \frac{1}{2}$, a labeling equilibrium exists if and only if

$$n \geq \frac{4c - 1}{1 - 2c}.$$

As n grows, $\underline{c}_n \to 0$ and $\overline{c}_n \to \frac{1}{2} = \mu$, so the interval of multiple equilibria widens toward $\left[0, \frac{1}{2}\right]$.

Informativeness. Under the symmetric strategy the label splits quality at the random cutoff $S_{1:n}$, and from (27.9) and (27.22),

$$R^2_{\text{sym}}(n) = \frac{3n}{(n+2)^2}. \tag{27.24}$$

With one label $R^2_{\text{sym}} = \frac{1}{3}$, with two labels it rises to $\frac{3}{8}$, it returns to $\frac{1}{3}$ at $n = 4$, and it falls toward zero after that. Intuitively, a second label makes the lack of any label a sharper signal of low quality, and at first this outweighs the fact that a label becomes easier to get; with many labels, almost every firm meets the easiest standard, so a label proves almost nothing. Part (iv) of the paper's Proposition 4 concerns the limit $n \to \infty$; the uniform example shows that the decline in informativeness need not start at $n = 1$.

Focal labels. Consider now that one label, say label X with standard S_X, is *focal*: consumers expect a firm to adopt it whenever it can, and to adopt another label only if it fails X. The market then sorts firms into three groups:

$$\{Q \geq S_X\}, \qquad \{S_{1:n} \leq Q < S_X\}, \qquad \{Q < S_{1:n}\}, \tag{27.25}$$

with, in the uniform case, posterior means

$$\frac{2}{3}, \qquad \frac{n+4}{3(n+2)}, \qquad \text{and} \qquad \frac{1}{n+2}.$$

Figure 27.16 depicts the partitions of the quality interval for two labels, X and Y, with realized standards $s_Y < s_X$, so that $s_{1:2} = s_Y$. Under the symmetric strategy consumers see only whether a firm has some label, and they split quality at $s_{1:2}$; under the focal strategy they also see whether the label is X, and the labeled group splits at s_X into two groups.

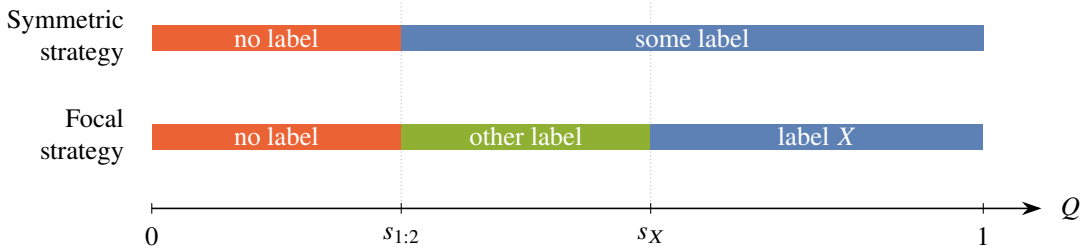


Figure 27.16: Groups of firms by quality under the symmetric and the focal strategies, for two labels with realized standards $s_{1:2} = 0.3$ and $s_X = 0.6$.

A firm in the middle group must prefer some label to none, and a firm in the top group must prefer label X to no label:

$$\frac{n+4}{3(n+2)} - \frac{1}{n+2} \geq c, \qquad \frac{2}{3} - \frac{1}{n+2} \geq c.$$

A firm in the top group also prefers X to any other label, since $\frac{2}{3}$ exceeds the middle-group mean. The first condition binds, so for $n \geq 2$ a focal labeling equilibrium exists if and only if

$$c \;\leq\; \frac{n+1}{3(n+2)}. \tag{27.26}$$

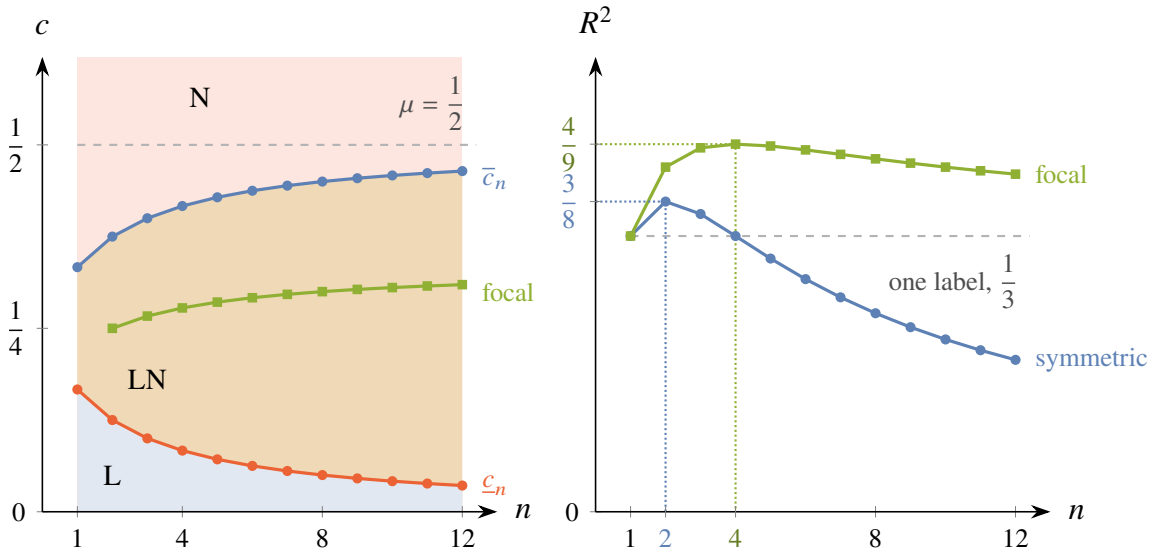
The cutoff in (27.26) equals $\frac{2}{3}\overline{c}_n$, so a focal equilibrium requires a lower cost than the symmetric labeling equilibrium. The partition (27.25) is finer than the symmetric one, and

$$R^2_{\text{focal}}(n) = \frac{n(n+8)}{3(n+2)^2}, \qquad R^2_{\text{focal}}(n) - R^2_{\text{sym}}(n) = \frac{n(n-1)}{3(n+2)^2} \;\geq\; 0. \tag{27.27}$$

Result 6 (Proposition 5 of the paper).

- With uncertain standards, a focal labeling equilibrium exists if the cost satisfies (27.26), and it is more informative than the symmetric one.
- With known standards, no focal labeling equilibrium exists.

Figure 27.17 depicts the cutoffs (27.23) and (27.26) in panel (a), with the regions L, LN, and N of the symmetric labeling strategy shaded, and the informativeness of the label in (27.24) and (27.27) in panel (b), with the number of labels on the horizontal axis. Panel (b) shows R^2_{sym} peaking at $\frac{3}{8}$ with two labels, and R^2_{focal} peaking at $\frac{4}{9}$ with four labels and converging to $\frac{1}{3}$, the value for a single uncertain label, because non-focal labels become uninformative as they multiply. With known standards a firm always prefers the toughest label it meets, and under the paper's refinement a deviation to a tougher label is rewarded, so focality cannot be sustained. This is the paper's case for "look for the label" campaigns and mandatory labels: they keep the information that a single label conveys, which proliferation destroys under the symmetric strategy.



(a) Equilibrium cutoffs

(b) Informativeness

Figure 27.17: Equilibrium cutoffs and informativeness of a label against the number of labels n, with uniform quality and uniform standards.

Labels n	$\underline{c}_n$	$\overline{c}_n$	Focal cutoff	R^2_{sym}	R^2_{focal}
1	$\frac{1}{6}$	$\frac{1}{3}$	–	$\frac{1}{3}$	$\frac{1}{3}$
2	$\frac{1}{8}$	$\frac{3}{8}$	$\frac{1}{4}$	$\frac{3}{8}$	$\frac{5}{12}$
3	$\frac{1}{10}$	$\frac{2}{5}$	$\frac{4}{15}$	$\frac{9}{25}$	$\frac{11}{25}$
4	$\frac{1}{12}$	$\frac{5}{12}$	$\frac{5}{18}$	$\frac{1}{3}$	$\frac{4}{9}$
6	$\frac{1}{16}$	$\frac{7}{16}$	$\frac{7}{24}$	$\frac{9}{32}$	$\frac{7}{16}$
10	$\frac{1}{24}$	$\frac{11}{24}$	$\frac{11}{36}$	$\frac{5}{24}$	$\frac{5}{12}$
$\to \infty$	0	$\frac{1}{2}$	$\frac{1}{3}$	0	$\frac{1}{3}$

Table 27.2: Cutoffs and informativeness under the symmetric and the focal labeling strategies, with uniform quality and uniform standards.

Table 27.2 reports the uniform values.

Known standards. With known standards, more labels let a firm reveal its quality more finely, and in the limit the problem becomes the one Lizzeri (1999) studies, in which every quality level can be certified. The paper shows that as n grows, non-labeling is an equilibrium almost surely if and only if $\mu \geq 1 - c$, and, for log-concave F, labeling is an equilibrium almost surely if and only if $\mu \leq 1 - c$. With uniform quality both conditions turn on whether c is above or below $\frac{1}{2}$, and a label becomes highly informative, instead of uninformative as with uncertain standards.

27.9 Seventh Step: Several Firms and the Strategic Choice of Labels

One label, many firms. Let m firms with independent qualities from F face one label. In a labeling equilibrium the share of firms with a label converges to $1 - F(s)$, so consumers learn the standard, and each firm effectively faces a known standard: with uniform quality, labeling is supported for $c \leq \overline{c}(s) = \frac{1}{2}$ regardless of the standard. In a non-labeling equilibrium no firm certifies, so consumers have nothing to learn from, and the condition stays $c \geq \underline{c} = \frac{1}{6}$ (Proposition 6 of the paper). As the number of firms grows, the interval of multiple equilibria therefore widens from $\left[\frac{1}{6}, \frac{1}{3}\right]$ to $\left[\frac{1}{6}, \frac{1}{2}\right]$. Figure 27.18 depicts the two ranges of costs, as in Figure 27.11: with many firms, labeling survives up to $c = \frac{1}{2}$ instead of $\frac{1}{3}$, and non-labeling still starts at $c = \frac{1}{6}$.

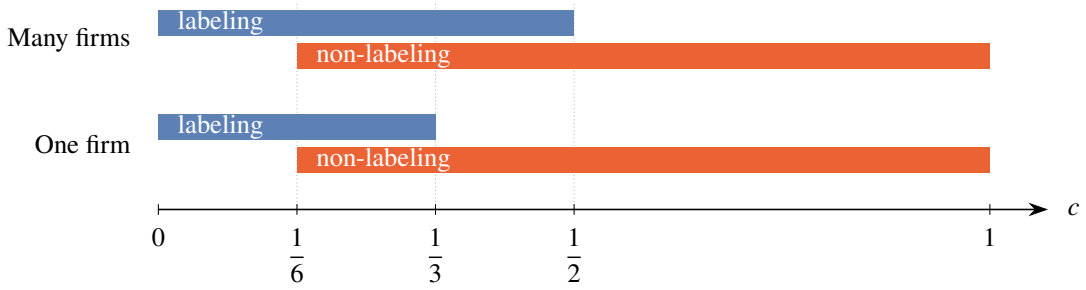


Figure 27.18: Costs that support a labeling equilibrium and a non-labeling equilibrium with one firm and with many firms, for uniform quality and a uniform standard.

Two firms, two labels. With two firms and two labels, the label one firm adopts carries news about the standards the other firm faces. Under the symmetric strategy, consumers see either both firms with the *same* label or the firms with *different* labels, and in the second case they cannot tell which firm met the tougher standard. Denoting by $S_{1:2} < S_{2:2}$ the easier and the tougher standard, the events are

$$\text{Same} = \left\{S_{1:2} \leq Q_1 < S_{2:2},\ S_{1:2} \leq Q_2 < S_{2:2}\right\} \cup \left\{Q_1 \geq S_{2:2},\ Q_2 \geq S_{2:2}\right\}, \tag{27.28}$$

$$\text{Different} = \left\{S_{1:2} \leq Q_1 < S_{2:2} \leq Q_2\right\} \cup \left\{S_{1:2} \leq Q_2 < S_{2:2} \leq Q_1\right\}. \tag{27.29}$$

Figure 27.19 depicts the two events for given standards, in the plane of the two firms' qualities. The events Same and Different are the shaded cells; in the remaining region at least one firm meets neither standard and has no label.

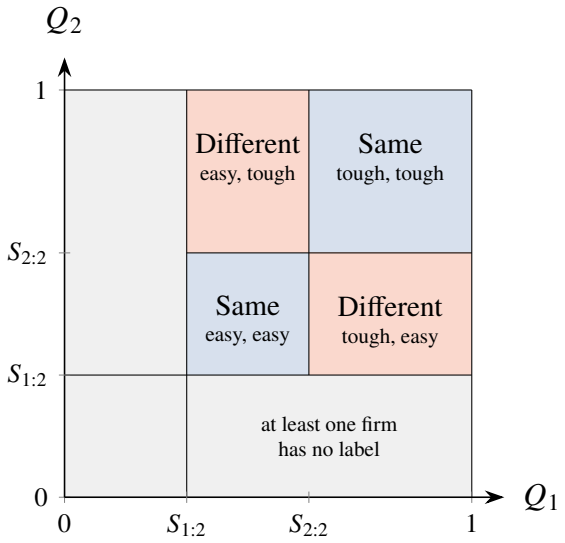


Figure 27.19: Events Same and Different in the plane of the two qualities, for two labels with given standards $S_{1:2} < S_{2:2}$; each cell lists the label of firm 1 and of firm 2.

A firm that meets both standards while its rival meets only the easier one can pool with the rival by adopting the easier label, and a firm that meets both standards while its rival does too can separate by adopting the easier label. The symmetric strategy survives only if neither deviation is profitable, which requires

$$E[Q_i \mid \text{Same}] = E[Q_i \mid \text{Different}].$$

With uniform quality the two events occur with probabilities $\frac{1}{3}$ and $\frac{1}{6}$, and

$$E[Q_i \mid \text{Same}] = E[Q_i \mid \text{Different}] = \frac{3}{5}, \tag{27.30}$$

so the uniform case is the knife-edge at which the symmetric strategy is an equilibrium.

Figure 27.20 depicts both conditional means against the prior mean μ, for the two reputation families of the fifth step, with uniform standards. The curves cross at $\mu = \frac{1}{2}$, at the value $\frac{3}{5}$ in (27.30). To the right of the crossing, the same label is better news than different labels: when reputations are good, the label a rival adopts is likely to be the tougher one, so each firm wants to adopt it too. The paper says that a firm with a good reputation *legitimizes* the label it adopts. To the left of the crossing, different labels are better news: the label a rival with a bad reputation adopts is likely to be the easier one, so each firm wants to avoid it, and a firm with a bad reputation *spoils* the label it adopts. Away from $\mu = \frac{1}{2}$ the symmetric labeling equilibrium does not exist, while a focal labeling equilibrium exists under the same cost condition as with a single firm (Proposition 7 of the paper). When firms choose in sequence, these incentives produce labeling cascades: the second firm herds onto the first firm's label if reputations are good and moves away from it if they are bad. The paper offers this spillover as one reason why certifiers recruit reputable firms as early adopters of a new label.

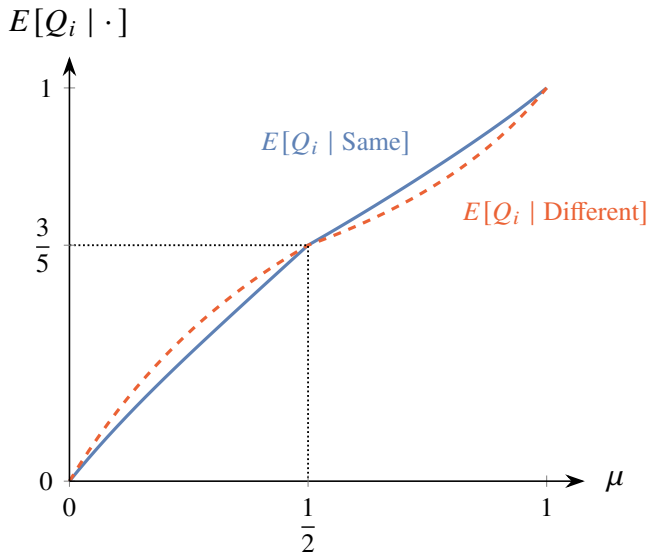


Figure 27.20: Expected quality of firm i when both firms display the same label and when they display different labels, against the prior mean μ; two firms, two uniform standards, $Q \sim \text{Beta}(1, b)$ for $\mu \leq \frac{1}{2}$ and $Q \sim \text{Beta}(a, 1)$ for $\mu \geq \frac{1}{2}$.

27.10 Eighth Step: What the Toy Model Keeps and Drops

Every one of the paper's main conclusions survives the uniform standard:

1. The Groucho and reverse Groucho effects: a label lowers the estimate of the standard from $\frac{1}{2}$ to $\frac{1}{3}$ and its absence raises it to $\frac{2}{3}$ (the second step).
2. Uncertain standards reveal less: R^2 falls from $\frac{1}{2}$ to $\frac{1}{3}$ (the third step, Proposition 1).
3. Labeling is harder to sustain and non-labeling easier: the cutoffs move from $\frac{1}{2}$ and $\frac{1}{4}$ to $\frac{1}{3}$ and $\frac{1}{6}$, and for any F, $\overline{c} = \frac{\sigma^2}{\mu(1-\mu)}$ and $\underline{c} = \frac{\sigma^2}{\mu}$ (the fourth step, Propositions 2 and 3).
4. Proliferation makes non-labeling an equilibrium for every $c > 0$ and the labeling equilibrium uninformative, while a focal label restores information (the sixth step, Propositions 4 and 5).
5. Several firms widen the interval of multiple equilibria, and with several labels firms legitimize or spoil labels (the seventh step, Propositions 6 and 7).

With a uniform standard, the labeling cutoff depends on the prior only through $\frac{\sigma^2}{\mu(1-\mu)}$, which separates the role of a firm's reputation (its mean) from the precision of consumers' prior (its dispersion); the general analysis leaves this split implicit. In addition, the informativeness of the symmetric labeling equilibrium is not monotone in the number of labels: in the uniform case it peaks at two labels and declines after that.

Relative to Harbaugh et al. (2011), the toy model takes the standard to be uniform, so the pass probability (27.1) is linear in quality; the paper allows any distribution G, and its propositions hold for general F and G (one known-standard limit in Proposition 4 also requires a log-concave F).

As in the paper, the firm's payoff is linear in expected quality, without prices or product-market competition between firms, and quality and standards are fixed: firms do not invest in quality, and certifiers do not choose their standards or fees. Each of these features can be relaxed one at a time.

27.11 Connections to Other Chapters

- Chapter 28 studies the opposite uncertainty: consumers know what the label certifies, but the test makes errors.
- In Chapter 29, the products under a common label share its reputation, and firms can pass off conventional products as green.

References

Harbaugh, R., J. W. Maxwell, and B. Roussillon (2011). "Label Confusion: The Groucho Effect of Uncertain Standards." *Management Science*, 57(9), pp. 1512–1527. https://doi.org/10.1287/mnsc.1110.1412

Lizzeri, A. (1999). "Information Revelation and Certification Intermediaries." *RAND Journal of Economics*, 30(2), pp. 214–231.

Chapter 28

Eco-labeling with Noisy Certification Tests

Companion note to Mason, C. F. (2011), "Eco-Labeling and Market Equilibria with Noisy Certification Tests," Environmental and Resource Economics 48(4): 537–560. Full reference at the end.

28.1 Introduction

Consumers cannot tell whether a product was made in a green or a polluting way. Consider a buyer of timber who cannot see whether the wood came from a sustainably managed forest. A certifier can test the product and award a label, but its tests make mistakes: some green products fail, and some brown (polluting) products pass. Can a noisy label still help consumers, and which sellers will pay to be tested?

Mason (2011) studies this question in a competitive market where consumers pay their expected value for a unit. A label works because green sellers pass the test with a higher probability than brown sellers. Labeled units are therefore more likely to be green than an average unit, and unlabeled units are less likely to be green. A green seller also gains more from applying than a brown seller, because its higher pass rate turns the same price gap into a larger expected gain. Brown sellers therefore never apply unless green sellers do.

The paper shows that the certification fee decides who applies. With a low fee, every seller applies. With an intermediate fee, all green sellers apply and only some brown sellers do. With a higher fee, only green sellers apply, and with a prohibitive fee no one does. A more precise test widens the price gap between labeled and unlabeled units. These results matter for certifiers and regulators, who choose the fee and the accuracy of the test.

The toy model gives each seller one unit and values a green unit at one and a brown unit at zero, so prices equal consumers' beliefs. The first to third steps set up the benchmark without a label, introduce the noisy test, and derive beliefs and prices. The fourth step compares the two types' gains from applying, the fifth step finds the equilibrium classes, the sixth step works through a numerical example, and the seventh step lists what the toy model keeps and drops.

28.2 The Toy Model

Consumers and demand. Consumers value a green unit more than a brown one but cannot tell them apart directly. As in Mason (2011), demand is perfectly elastic. We normalize the two willingness-to-pay levels to

$$v_G = 1, \qquad v_B = 0,$$

so the green "premium" is the whole willingness to pay. A consumer who believes a unit is green with probability μ pays the expected value

$$P(\mu) = \mu\, v_G + (1 - \mu)\, v_B = \mu,$$

so *prices are beliefs*.

Sellers. There is a unit mass of sellers, each offering one unit. A fraction $\lambda \in (0, 1)$ are green and the complementary fraction $1 - \lambda$ are brown; a seller privately knows its own type. Because each seller brings a single unit to market, a seller's payoff is simply the price it receives, net of any certification fee it pays. This single-unit structure is what keeps the equilibrium conditions linear in the beliefs.

Market structure. Within each observable segment (labeled or unlabeled) the market is competitive: the price equals consumers' expected willingness to pay given the belief that a random unit in that segment is green. This is Mason's rational-expectations pricing rule.

28.3 First Step: The No-Information Benchmark

Consider that no eco-label is available. The whole market pools, and consumers pay the expected willingness to pay for a random unit,

$$P_0 \;=\; \lambda\, v_G + (1 - \lambda)\, v_B \;=\; \lambda. \tag{28.1}$$

Every seller, green or brown, receives $P_0 = \lambda$. This is the reference price against which the label is measured, and the outcome the market reverts to whenever the fee is too high for anyone to seek certification.

28.4 Second Step: The Noisy Certification Test

A certifying agency offers a test. A seller that applies pays a fee $A \geq 0$ and receives a public verdict (pass or fail). We index the test by a single *precision* $\tau \in [0, 1]$: a green seller passes with probability ϕ_G and a brown seller with probability ϕ_B, where

$$\phi_G \;\equiv\; \Pr(\text{pass} \mid \text{G}) \;=\; \frac{1 + \tau}{2}, \qquad \phi_B \;\equiv\; \Pr(\text{pass} \mid \text{B}) \;=\; \frac{1 - \tau}{2}. \tag{28.2}$$

At $\tau = 0$ the test is uninformative ($\phi_G = \phi_B = \frac{1}{2}$); at $\tau = 1$ it is perfect ($\phi_G = 1$, $\phi_B = 0$). A green seller fails with probability $1 - \phi_G = \frac{1-\tau}{2}$ (a *type I* error, or false rejection), and a brown seller passes with probability $\phi_B = \frac{1-\tau}{2}$ (a *type II* error, or false acceptance). The two error rates are equal, so τ is the single dial that sets test quality, with informativeness $\phi_G - \phi_B = \tau$. Table 28.1 reports the four outcomes conditional on applying, and Table 28.2 shows its two extremes. Figure 28.1 plots the two pass probabilities against τ, and Figure 28.2 shows the same four flows as a tree.

Seller	Labeled (pass)	Unlabeled (fail)
Green seller	$\frac{1+\tau}{2}$	$\frac{1-\tau}{2}$ (type I error)
Brown seller	$\frac{1-\tau}{2}$ (type II error)	$\frac{1+\tau}{2}$

Table 28.1: Certification outcomes by seller type, conditional on applying. Under the symmetric precision test (28.2), the diagonal (correct classification) carries mass $\frac{1+\tau}{2}$ and each off-diagonal error carries mass $\frac{1-\tau}{2}$; a more precise test (higher τ) moves mass onto the diagonal.

$\tau = 0$	Labeled	Unlabeled
Green seller	$\frac{1}{2}$	$\frac{1}{2}$
Brown seller	$\frac{1}{2}$	$\frac{1}{2}$

(a) Uninformative test.

$\tau = 1$	Labeled	Unlabeled
Green seller	1	0
Brown seller	0	1

(b) Perfect test.

Table 28.2: The two extremes of the confusion matrix. At $\tau = 0$ (panel a) every cell equals $\frac{1}{2}$: the label is pure noise, so the two rows are identical and a pass says nothing about type. At $\tau = 1$ (panel b) the diagonal is 1 and both errors vanish: the label separates green from brown perfectly.

28.5 Third Step: Posterior Beliefs and Prices

Consider that (for now) every seller applies. Because each seller brings one unit, the mass of green units in a segment relative to the total is a plain Bayes ratio, and the two posteriors are

$$\mu \equiv \Pr(\text{G} \mid \text{labeled}) = \frac{\lambda\phi_G}{\lambda\phi_G + (1-\lambda)\phi_B}, \tag{28.3}$$

$$\nu \equiv \Pr(\text{G} \mid \text{unlabeled}) = \frac{\lambda(1-\phi_G)}{\lambda(1-\phi_G) + (1-\lambda)(1-\phi_B)}. \tag{28.4}$$

Substituting the precision test (28.2),

$$\mu = \frac{\lambda(1+\tau)}{\lambda(1+\tau) + (1-\lambda)(1-\tau)}, \qquad \nu = \frac{\lambda(1-\tau)}{\lambda(1-\tau) + (1-\lambda)(1+\tau)}. \tag{28.5}$$

At $\lambda = \frac{1}{2}$ these collapse to $\mu = \frac{1+\tau}{2}$ and $\nu = \frac{1-\tau}{2}$, the pass probabilities themselves.

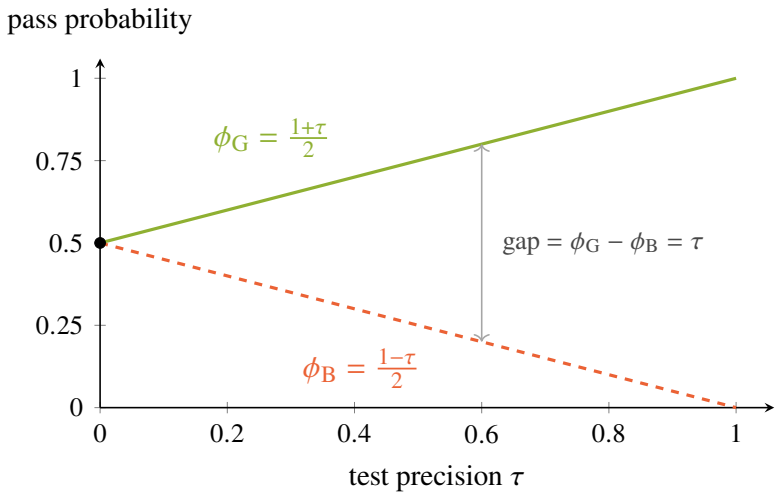


Figure 28.1: The two pass probabilities as functions of precision τ: a green seller passes with probability $\phi_G = \frac{1+\tau}{2}$ (rising, green) and a brown seller with probability $\phi_B = \frac{1-\tau}{2}$ (falling, red). They coincide at $\frac{1}{2}$ when $\tau = 0$ (uninformative) and separate to 1 and 0 when $\tau = 1$ (perfect). The vertical gap between them is the test's informativeness, $\phi_G - \phi_B = \tau$.

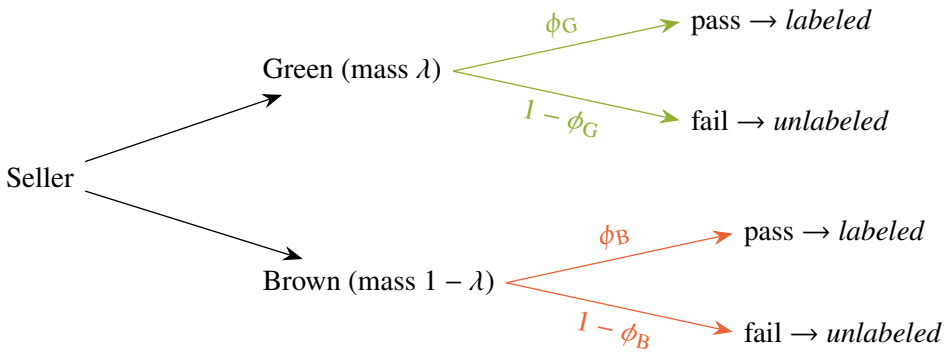


Figure 28.2: The four flows through the noisy test. Green sellers pass with the higher probability ϕ_G; brown sellers pass with the lower probability ϕ_B. Failures and non-applicants are pooled together in the *unlabeled* segment.

Prices equal posteriors. Within each segment competition sets the price to the expected willingness to pay, which under $v_G = 1$ and $v_B = 0$ equals the posterior:

$$P_c = \mu, \qquad P_u = \nu, \qquad P_0 = \lambda. \tag{28.6}$$

Sorting. Any precision $\tau > 0$ makes $\phi_G > \phi_B$, which translates into $\mu > \lambda > \nu$: an informative test raises the green share (and hence the price) in the labeled segment above the population share and lowers it in the unlabeled segment. Using $\phi_G - \phi_B = \tau$,

$$\mu - \lambda = \frac{\lambda\phi_G - \lambda\left[\lambda\phi_G + (1-\lambda)\phi_B\right]}{\lambda\phi_G + (1-\lambda)\phi_B} = \frac{\lambda(1-\lambda)\,\tau}{\lambda\phi_G + (1-\lambda)\phi_B} \;>\; 0,$$

and $\lambda - \nu > 0$ by an analogous computation.

Result 1 (sorting). The label puts the two segments on opposite sides of the prior, $\nu < \lambda < \mu$, with gaps

$$\mu - \lambda = \frac{2\lambda(1-\lambda)\,\tau}{1 + (2\lambda - 1)\,\tau} \qquad \text{and} \qquad \lambda - \nu = \frac{2\lambda(1-\lambda)\,\tau}{1 - (2\lambda - 1)\,\tau},$$

which are zero for an uninformative test and increase in the precision τ; they are proportional to τ when $\lambda = \frac{1}{2}$.

28.6 Fourth Step: Application Payoffs and Ordered Incentives

A seller of type $k = G, B$ that applies is labeled with probability ϕ_k, earning the certified price $P_c = \mu$, and unlabeled with probability $1 - \phi_k$, earning $P_u = \nu$; it pays the fee A either way. If it does not apply it is unlabeled for sure and earns ν. The gain from applying is therefore

$$W_k \;\equiv\; \left[\phi_k\,\mu + (1-\phi_k)\,\nu - A\right] - \nu \;=\; \phi_k\,(\mu - \nu) - A, \tag{28.7}$$

which is *linear* in the price gap $\mu - \nu$. Subtracting the two gains, and using $\phi_G - \phi_B = \tau$,

$$W_G - W_B = (\phi_G - \phi_B)(\mu - \nu) = \tau\,(\mu - \nu) \;>\; 0. \tag{28.8}$$

Result 2 (ordered incentives). Green sellers have a strictly stronger incentive to seek certification than brown sellers, $W_G > W_B$. Hence no equilibrium can have some brown sellers applying while no green sellers do.

Intuitively, a green seller's higher pass probability multiplies the same price gap into a larger gain, so if brown sellers apply, green sellers apply too.

28.7 Fifth Step: The Equilibrium Classes

Depending on the size of the fee A, the equilibrium falls into one of four regions, corresponding to three kinds of perfect Bayesian equilibrium (PBE): a *separating* PBE, a *partially separating* PBE,

and a *pooling* PBE that appears in two forms (both sellers apply, or no seller applies). The cutoffs come out in closed form.

Pooling PBE (both sellers apply). Both types apply, so the market cannot tell green from brown among applicants. This requires $W_{\mathrm{B}} \geq 0$, i.e.,

$$A \;\leq\; \phi_{\mathrm{B}}\,(\mu - \nu) \;\equiv\; \overline{A}_{\mathrm{pool}}, \tag{28.9}$$

with μ, ν the pooling posteriors (28.3)–(28.4).

Separating PBE (green apply, brown do not). Since only greens apply, every labeled unit is green, so $\hat{\mu} = 1$ and $\hat{P}_c = 1$. The unlabeled posterior contains greens that failed plus all browns:

$$\hat{\nu} = \frac{\lambda(1 - \phi_{\mathrm{G}})}{\lambda(1 - \phi_{\mathrm{G}}) + (1 - \lambda)}, \qquad \hat{P}_u = \hat{\nu}. \tag{28.10}$$

Because the certified price is now 1, the price gap is $1 - \hat{\nu}$, and the separating range on the fee is

$$\underbrace{\phi_{\mathrm{B}}\,(1 - \hat{\nu})}_{\text{brown IC: do not apply}} \;\leq\; A \;\leq\; \underbrace{\phi_{\mathrm{G}}\,(1 - \hat{\nu})}_{\text{green IC: do apply}}. \tag{28.11}$$

Since $\phi_{\mathrm{G}} > \phi_{\mathrm{B}}$, this interval has positive length.

Partially separating PBE (green apply, brown mix). For fees in the gap

$$\overline{A}_{\mathrm{pool}} \;<\; A \;<\; \phi_{\mathrm{B}}\,(1 - \hat{\nu}), \tag{28.12}$$

brown sellers mix so as to be indifferent between applying and not applying. The mixing probability $\lambda_{\mathrm{B}} \in (0, 1)$ solves the indifference condition

$$\phi_{\mathrm{B}}\left(\mu^* - \nu^*\right) = A, \tag{28.13}$$

where the labeled and unlabeled posteriors, when the green seller applies for sure and the brown seller applies with probability λ_{B}, are

$$\mu^* = \frac{\lambda\phi_{\mathrm{G}}}{\lambda\phi_{\mathrm{G}} + (1 - \lambda)\,\lambda_{\mathrm{B}}\phi_{\mathrm{B}}}, \tag{28.14}$$

$$\nu^* = \frac{\lambda(1 - \phi_{\mathrm{G}})}{\lambda(1 - \phi_{\mathrm{G}}) + (1 - \lambda)\,(1 - \lambda_{\mathrm{B}}\phi_{\mathrm{B}})}. \tag{28.15}$$

These reduce to the pooling posteriors (28.3)–(28.4) when $\lambda_{\mathrm{B}} = 1$ and to the separating posteriors ($\mu^* = 1$, $\nu^* = \hat{\nu}$) when $\lambda_{\mathrm{B}} = 0$. This regime deserves a closer look, since the mixing probability is what makes the label only partly trustworthy.

The mixing probability lies strictly between 0 *and* 1. Write the brown seller's gain from applying as a function of the mixing probability,

$$g(\lambda_{\mathrm{B}}) \;\equiv\; \phi_{\mathrm{B}}\big(\mu^* - \nu^*\big), \tag{28.16}$$

so the indifference condition (28.13) is $g(\lambda_B) = A$. As more brown sellers apply, the labeled pool is diluted and the unlabeled pool is cleansed: in (28.14) the numerator of μ^* is fixed while its denominator grows, so μ^* *falls* in λ_B; in (28.15) the term $(1 - \lambda_B\phi_B)$ shrinks, so ν^* *rises*. Both moves shrink the gap $\mu^* - \nu^*$, so g is strictly decreasing, from

$$g(0) = \phi_B(1 - \hat{\nu}) = \underline{A}_{\text{sep}} \qquad \text{down to} \qquad g(1) = \phi_B(\mu - \nu) = \overline{A}_{\text{pool}}.$$

A partially separating fee lies strictly between these two cutoffs, $\overline{A}_{\text{pool}} < A < \underline{A}_{\text{sep}}$, so by continuity and strict monotonicity the equation $g(\lambda_B) = A$ has a *unique* solution $\lambda_B \in (0, 1)$. Outside the gap there is no interior solution: for $A \le \overline{A}_{\text{pool}}$ the brown seller strictly prefers to apply ($\lambda_B = 1$, pooling), and for $A \ge \underline{A}_{\text{sep}}$ it strictly prefers not to ($\lambda_B = 0$, separation).

Comparative static in the fee. Differentiating $g(\lambda_B) = A$ and using $g' < 0$,

$$\frac{\partial \lambda_B}{\partial A} = \frac{1}{g'(\lambda_B)} \; < \; 0, \tag{28.17}$$

so a higher fee lowers the fraction of brown sellers that apply. The intuition is a dilution feedback: a higher fee makes the label less attractive to a brown seller, so fewer apply; but fewer brown applicants make the label *cleaner* (a higher μ^* and a wider price gap), which pushes the gain from applying back up until the marginal brown seller is again indifferent. At $\lambda = \frac{1}{2}$ and $\tau = \frac{4}{5}$, where (28.14)–(28.15) simplify to

$$\mu^* = \frac{9}{9 + \lambda_B}, \qquad \nu^* = \frac{1}{11 - \lambda_B},$$

the solution runs monotonically from $\lambda_B = 1$ at $A = \frac{2}{25}$ to $\lambda_B = 0$ at $A = \frac{1}{11}$, as Table 28.3 reports.

Fee A	Brown application probability λ_B
$\frac{2}{25} = 0.080$	1 (pooling boundary)
0.082	0.80
0.084	0.61
0.086	0.43
0.088	0.25
$\frac{1}{11} \approx 0.091$	0 (separating boundary)

Table 28.3: The brown seller's equilibrium application probability λ_B across the partially separating gap, at $\lambda = \frac{1}{2}$ and $\tau = \frac{4}{5}$. As the fee rises from $\frac{2}{25}$ to $\frac{1}{11}$, λ_B falls from 1 to 0.

Figure 28.3 depicts the brown seller's gain $g(\lambda_B)$ in (28.16), with λ_B on the horizontal axis, for $\lambda = \frac{1}{2}$ and $\tau = \frac{4}{5}$, together with the fee $A = 0.086$ of Table 28.3.

Intuitively, every extra brown applicant dilutes the labeled segment and cleans the unlabeled one, so the gain falls as λ_B rises. A higher fee shifts the horizontal line up, and it meets the curve further to the left, at a lower λ_B.

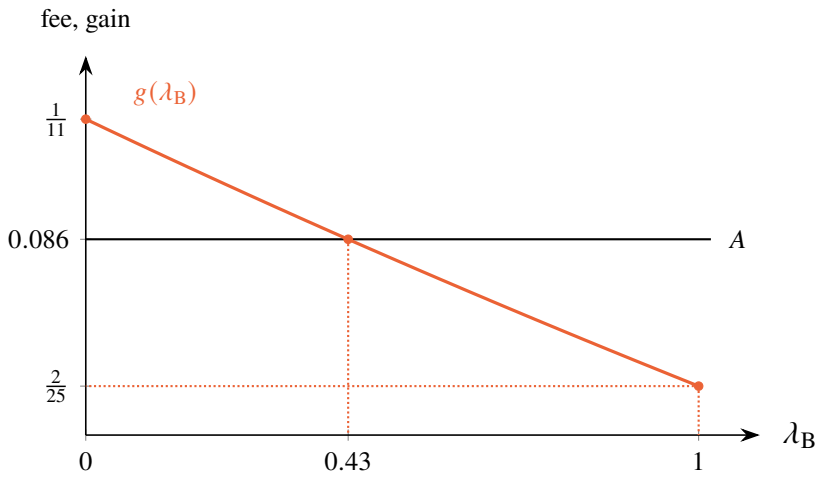


Figure 28.3: The brown seller's gain from applying, $g(\lambda_B)$, at $\lambda = \frac{1}{2}$ and $\tau = \frac{4}{5}$. It falls from $\frac{1}{11}$ to $\frac{2}{25}$, so each fee in this range meets it once.

Comparative static in precision. The effect of τ on λ_B is not monotone, because a more precise test pulls in two directions. It lowers the brown pass probability $\phi_B = \frac{1-\tau}{2}$, which discourages brown application, but it widens the price gap $\mu^* - \nu^*$, which encourages it. Either force can dominate, so λ_B can rise or fall with τ; these are the same two forces that Mason (2011) identifies for a lower brown pass probability. The gap itself, $\big(\overline{A}_{\text{pool}}(\tau), \underline{A}_{\text{sep}}(\tau)\big)$, also shifts and narrows with τ, so a fixed fee can move out of the partially separating region altogether as the test becomes more precise.

What this means for the green seller. The green seller applies for sure. At the brown seller's indifference, $\mu^* - \nu^* = \frac{A}{\phi_B}$, so the green seller's net gain from applying is

$$\phi_G\big(\mu^* - \nu^*\big) - A = A\,\frac{\phi_G - \phi_B}{\phi_B} = \frac{2A\tau}{1-\tau} \;>\; 0, \tag{28.18}$$

strictly positive and increasing in both the fee and the precision. The brown seller's mixing is what sets the label's value: it leaves the brown seller indifferent while the green seller, passing the test with a higher probability, still strictly profits. A more precise test widens this wedge, so the green seller's advantage from certification grows with τ even though the brown seller is held to zero rents.

Pooling PBE (no seller applies). When the fee clears the upper separating cutoff, $A > \phi_G(1 - \hat{\nu})$, not even a green seller applies. Both types pool on not applying, and the market reverts to the no-information benchmark of the first step.

Result 3 (from the fee to the equilibrium). As A rises from zero, the equilibria in which some seller applies move through three regions, in this order:

- the pooling PBE in which both sellers apply;
- the partially separating PBE;
- the separating PBE.

For $A > \phi_G(1 - \hat{\nu})$, no seller applies, and this pooling PBE is the unique equilibrium.

The cutoffs (28.9), (28.11), and (28.12) are ordered strictly, because $\phi_G > \phi_B$, the pooling posterior μ is strictly less than the separating posterior $\hat{\mu} = 1$, and the pooling posterior ν is strictly greater than the separating posterior $\hat{\nu}$. The pooling PBE in which no seller applies also exists at lower fees, and whether it does depends on what consumers believe after an unexpected label. If they believe that the label comes from a green seller, this PBE requires $A \geq \phi_G(1-\lambda)$, which lies below $\phi_G(1-\hat{\nu})$ because $\lambda > \hat{\nu}$; if they believe that it comes from a brown seller, it exists for every $A \geq 0$. As in Mason (2011), a second partially separating PBE, in which green sellers mix and brown sellers do not apply, exists for $\phi_G(1-\lambda) < A < \phi_G(1-\hat{\nu})$; like Mason, this chapter focuses on the one in which brown sellers mix.

The equilibrium map. At $\lambda = \frac{1}{2}$ the pooling posteriors reduce to $\mu = \frac{1+\tau}{2}$ and $\nu = \frac{1-\tau}{2}$, so

$$\mu - \nu = \tau, \qquad \hat{\nu} = \frac{1-\tau}{3-\tau}.$$

The three cutoffs of (28.9), (28.11), and (28.12) then become

$$\overline{A}_{\text{pool}}(\tau) = \frac{\tau(1-\tau)}{2}, \qquad \underline{A}_{\text{sep}}(\tau) = \frac{1-\tau}{3-\tau}, \qquad \overline{A}_{\text{sep}}(\tau) = \frac{1+\tau}{3-\tau}. \tag{28.19}$$

These are the single-unit counterparts of the cutoffs in Mason (2011), each a pass probability times a price gap. Plotting the three curves in the (τ, A) plane carves it into the four regions of Figure 28.4: the shaded separating region between $\underline{A}_{\text{sep}}$ and $\overline{A}_{\text{sep}}$, the pooling region (both apply) below $\overline{A}_{\text{pool}}$, the partially separating wedge between them, and the pooling region (no seller applies) above $\overline{A}_{\text{sep}}$. The vertical slice at $\tau = \frac{4}{5}$ gives the numerical example of the sixth step:

$$\overline{A}_{\text{pool}} = \frac{2}{25}, \qquad \underline{A}_{\text{sep}} = \frac{1}{11}, \qquad \overline{A}_{\text{sep}} = \frac{9}{11}.$$

28.8 Sixth Step: A Numerical Example

Take

$$\lambda = \frac{1}{2}, \qquad \tau = \frac{4}{5}, \qquad A = \frac{1}{20},$$

so from the precision test (28.2) the pass probabilities are $\phi_G = \frac{9}{10}$ and $\phi_B = \frac{1}{10}$ (a type I and a type II error rate of $\frac{1}{10}$ each).

Prices. The no-label price is $P_0 = \lambda = \frac{1}{2}$. At $\lambda = \frac{1}{2}$ the posteriors equal the pass probabilities, so

$$\mu = \frac{1+\tau}{2} = \frac{9}{10}, \qquad \nu = \frac{1-\tau}{2} = \frac{1}{10}, \qquad \mu - \nu = \frac{4}{5}.$$

Prices are the posteriors: $P_c = \frac{9}{10}$ and $P_u = \frac{1}{10}$, straddling $P_0 = \frac{1}{2}$.

Which equilibrium? The pooling cutoff (28.9) is

$$\overline{A}_{\text{pool}} = \phi_B(\mu - \nu) = \frac{1}{10} \cdot \frac{4}{5} = \frac{2}{25} = 0.08.$$

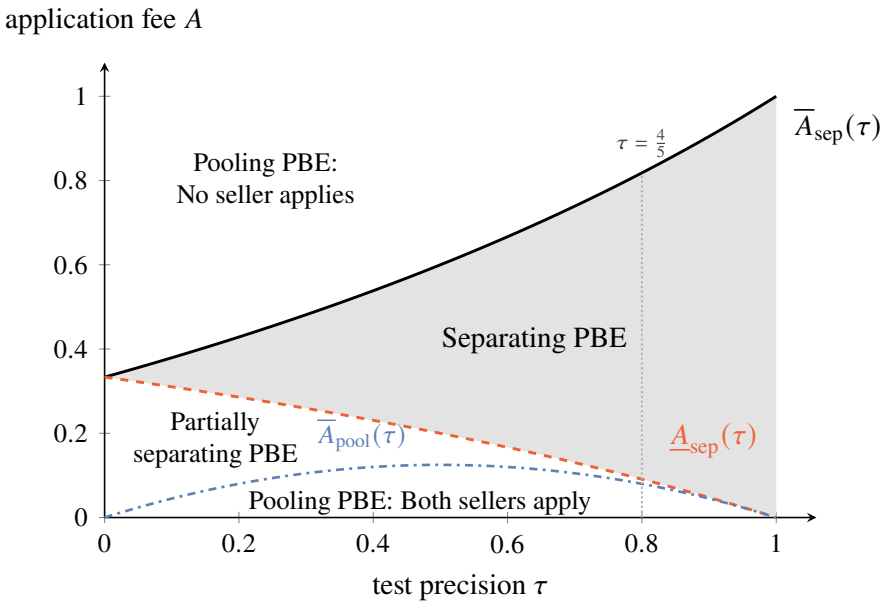


Figure 28.4: The equilibrium map in the (τ, A) plane, at $\lambda = \frac{1}{2}$. The three cutoffs of (28.19) split the plane into the four regions of the fifth step: a separating PBE (only green apply, shaded), a partially separating PBE (brown randomizes), and the pooling PBE in its two forms (both sellers apply at low fees, no seller applies at high fees). The separating region widens with precision; the both-apply pooling region is a thin sliver that is largest at intermediate precision. The dotted vertical slice at $\tau = \frac{4}{5}$ is the numerical example of the sixth step.

Since

$$A = \frac{1}{20} = 0.05 < 0.08 = \overline{A}_{\text{pool}},$$

we are in the pooling PBE: both types apply.

Application gains. From (28.7),

$$W_k = \phi_k(\mu - \nu) - A = \phi_k \cdot \frac{4}{5} - \frac{1}{20},$$

so

$$W_{\text{G}} = \frac{9}{10} \cdot \frac{4}{5} - \frac{1}{20} = \frac{36}{50} - \frac{1}{20} = \frac{72-5}{100} = \frac{67}{100}, \qquad W_{\text{B}} = \frac{1}{10} \cdot \frac{4}{5} - \frac{1}{20} = \frac{8-5}{100} = \frac{3}{100}.$$

Both are strictly positive, and $W_{\text{G}} > W_{\text{B}}$, as required. Raising the fee above $\frac{2}{25}$ would push the market into the partially separating region of Table 28.3, and above $\frac{1}{11}$ into a separating PBE.

28.9 Seventh Step: What the Toy Model Keeps and Drops

Every one of Mason's positive conclusions about the label survives the normalizations:

1. Bayesian sorting, $\nu < \lambda < \mu$, with gaps that increase in the precision τ (Result 1).

2. Ordered incentives, $W_G > W_B$, so brown sellers never apply unless green sellers do (Result 2).
3. Four equilibrium regions indexed by the fee: a pooling PBE in which both apply (low A), a partially separating PBE (intermediate A), a separating PBE (high A), and a pooling PBE in which no seller applies (prohibitive A), with closed-form cutoffs such as $\underline{A}_{sep} = \frac{1-\tau}{3-\tau}$ and $\overline{A}_{sep} = \frac{1+\tau}{3-\tau}$ (the single-unit counterparts of Mason's), shown in Figure 28.4 (Result 3).
4. The mixing probability is interior and well-behaved: it is unique, decreasing in the fee (Table 28.3 and Figure 28.3), and non-monotone in precision, and it leaves the green seller with a strictly positive gain $\frac{2A\tau}{1-\tau}$ that grows with τ (the fifth step).

What the toy model drops, relative to Mason (2011) and to the certifier-design literature, is threefold: (i) sellers here bring a single unit, so there is no choice of output and no producer-surplus welfare accounting; (ii) the two error rates are tied together by the symmetric test (28.2) at $\frac{1-\tau}{2}$, whereas Mason lets the two pass rates move independently; and (iii) the precision τ is fixed, rather than chosen by a certifier or a regulator who trades the value of information against the cost of a more accurate test. Each can be reinstated one at a time; the advantage of the toy model is that the sorting mechanism and the four equilibrium regions come through with almost no algebra.

28.10 Connections to Other Chapters

- Chapter 27 studies the opposite uncertainty: the label certifies a pass-fail standard without errors, but consumers do not know how demanding the standard is.
- In Chapter 29, brown products carry a green label because of fraud rather than test errors.
- The precision of the test in this chapter is an information structure, the object that the sender designs in Chapter 30.

References

Mason, C. F. (2011). "Eco-Labeling and Market Equilibria with Noisy Certification Tests." *Environmental and Resource Economics*, 48(4), pp. 537–560. https://doi.org/10.1007/s10640-010-9402-5

Chapter 29

Fraud in Green Markets: Collective Reputation and Eco-Certification

Companion note to Hamilton, S. F., and D. Zilberman (2006), "Green Markets, Eco-Certification, and Equilibrium Fraud," Journal of Environmental Economics and Management 52(3): 627–644.

29.1 Introduction

Many shoppers pay more for organic food or shade-grown coffee. They cannot check whether a product is truly green, not even after consuming it; they only see an eco-label that many firms share. A firm selling under that label can therefore pass off a conventional product, which we call a brown product, as green. How much of this fraud should we expect, and can certification fees or environmental taxes reduce it?

Hamilton and Zilberman (2006) answer these questions with a simple idea. Consumers cannot tell individual products apart, but they know the label's purity, that is, the share of truly green products among all products sold under the label, and they pay a premium that rises with purity. Each firm shares this premium with every other firm under the label, but it bears the full cost of its own green production. Each firm then has an incentive to free ride: it sells brown products disguised as green and lets the green production of its rivals keep the premium high.

The paper finds that fraud is more common when more firms share the label, since each firm captures a smaller part of the premium that its own green units create. A certification fee on every labeled unit reduces fraud and can raise firms' profits, while a tax on brown products raises fraud. These results matter for policy: a fee that looks like a burden on green firms can help them, and a tax that looks like support for green products can increase fraud.

The first step draws the paper's outcome regions for different numbers of firms. The toy model then studies a duopoly with linear demand and a convex cost of green production. The second step compares green and disguised units, the third step solves for equilibrium fraud, the fourth step

studies certification fees and brown taxes, the fifth step compares the equilibrium with the social optimum, the sixth step adds free entry, and the seventh step lists what the toy model keeps and drops.

29.2 First Step: The Paper's Outcome Regions and the Number of Firms

The paper's model. Consumers know only the share ρ of truly green products among those sold under the label, its purity. The label's price premium over the brown good is $\rho(1-\hat{y})$ when total sales under the label are $\hat{y}$ (Section 29.3 derives this premium from consumers' preferences). Each of n symmetric firms produces truly green units at a constant extra cost $\tilde{c}$ per unit, with $\tilde{c} \leq 1$, and can sell brown units disguised as green at an extra cost d per unit; as in the paper's Figure 1, there are no certification costs. In a symmetric equilibrium each firm sells $\frac{\hat{y}}{n}$, and the first-order conditions for disguised and green output, equations (15) and (16) of the paper, are

$$\rho\left(\frac{n-1}{n}-\hat{y}\right)-d \leq 0, \qquad \rho\left(\frac{n-1}{n}-\hat{y}\right)+\frac{1-\hat{y}}{n}-\tilde{c} \leq 0, \tag{29.1}$$

each holding with equality if the corresponding output is positive.

Three regions. With unit purity, $\rho = 1$, the second condition in (29.1) gives $\hat{y} = \frac{n(1-\tilde{c})}{n+1}$, and the first one then shows that no firm wants to disguise brown units if and only if $\tilde{c} \leq \tilde{c}^f$. Otherwise both conditions hold with equality; subtracting them gives $\hat{y} = 1 - n(\tilde{c} - d)$, which is positive if and only if $\tilde{c} < \tilde{c}^b$, where

$$\tilde{c}^f = \frac{1}{n^2} + \frac{n+1}{n}\,d, \qquad \tilde{c}^b = \frac{1}{n} + d. \tag{29.2}$$

Both loci rise from the vertical axis and meet on the line $\tilde{c} = 1$ at $d = \frac{n-1}{n}$. Below the 45° line, $\tilde{c} < d$, producing a truly green unit costs less than disguising a brown one.

Result 1 (outcome regions, Figure 1 and Proposition 2 of the paper). With n firms:

- region I, unit purity, arises if and only if $\tilde{c} \leq \tilde{c}^f$; it contains every cost pair below the 45° line;
- region II, a green market with disguised brown products, arises if and only if $\tilde{c}^f < \tilde{c} \leq \tilde{c}^b$;
- region III, no green market in a symmetric pure-strategy equilibrium, arises if and only if $\tilde{c} > \tilde{c}^b$;
- within the unit square of cost pairs $(d, \tilde{c})$, regions II and III have areas $\frac{(n-1)^2}{2n^3}$ and $\frac{(n-1)^2}{2n^2}$, and region I covers the rest.

Figure 29.1 depicts the three regions in the $(d, \tilde{c})$ plane, as in Figure 1 of the paper, for $n = 1, 2, 4$, and 100, and Table 29.1 reports their areas. With a monopolist, both loci lie on or above $\tilde{c} = 1$, so the whole square is region I. With two firms, $\tilde{c}^f$ starts at $\frac{1}{4}$ and $\tilde{c}^b$ at $\frac{1}{2}$, and both loci reach $\tilde{c} = 1$ at $d = \frac{1}{2}$. With four firms, they start at $\frac{1}{16}$ and $\frac{1}{4}$ and reach $\tilde{c} = 1$ at $d = \frac{3}{4}$. With 100 firms, they start at $\frac{1}{10000}$ and $\frac{1}{100}$ and run just above the 45° line, so unit purity survives essentially only where

producing green units costs less than disguising brown ones, and above that line the green market disappears.[1]

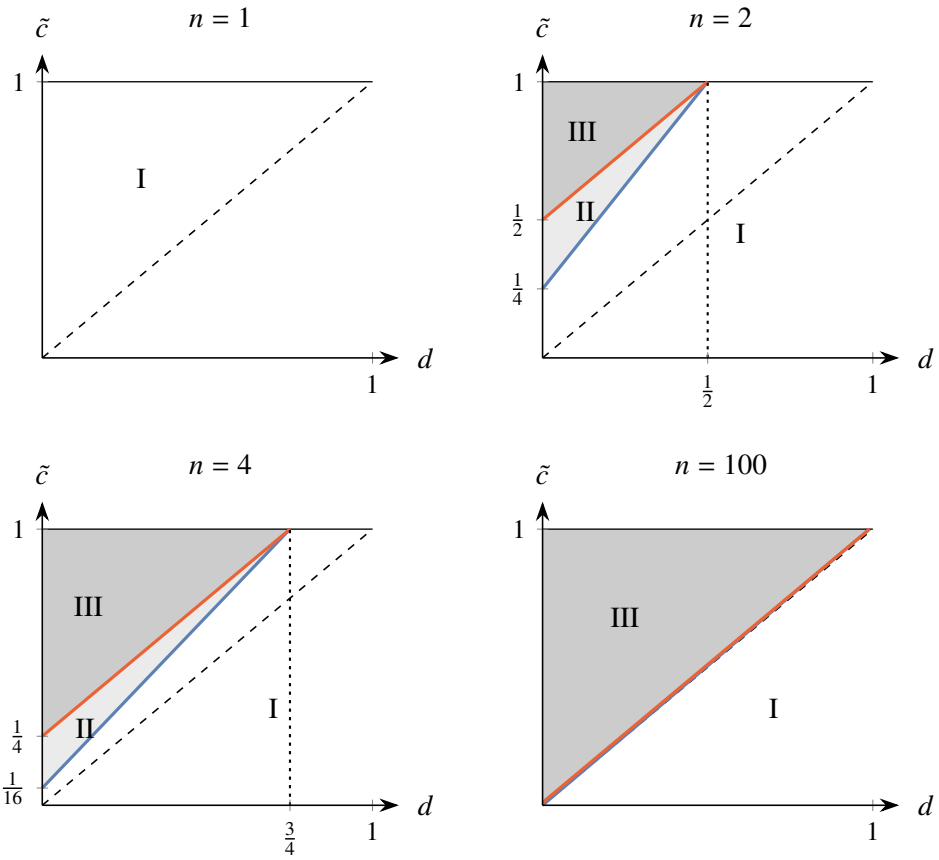


Figure 29.1: Outcome regions in the $(d, \tilde{c})$ plane for four numbers of firms: unit purity (I, white), a green market with disguised products (II, light gray), and no green market (III, dark gray). The blue line is $\tilde{c}^f$, the red line $\tilde{c}^b$, and the dashed line $\tilde{c} = d$.

Intuitively, each firm bears the whole cost of its green units but shares the label's reputation with its $n - 1$ rivals, so with more firms each firm captures a smaller part of the premium that its own green units create. Both loci fall toward the 45° line: the highest cost of green production compatible with unit purity falls toward d, the cost of disguising a brown unit, and the highest cost that keeps any green market alive falls with it. The fraud region is largest at $n = 3$, with area $\frac{2}{27}$, and then vanishes, so with many firms the market moves from unit purity directly to no green market as the cost of green production rises.

[1]With a constant cost $\tilde{c}$, the solution of (29.1) in region II is a saddle point of each firm's profit: the determinant of its Hessian with respect to green and disguised output is $-\frac{(n\hat{y}-n+1)^2}{n^2\hat{y}^2} \leq 0$, so a firm gains by deviating. Region II outcomes therefore require a convex cost of the green attribute, which the toy model of Section 29.3 assumes.

Number of firms, n	Region I (unit purity)	Region II (fraud)	Region III (no green market)
1	1	0	0
2	$\frac{13}{16}$	$\frac{1}{16}$	$\frac{1}{8}$
4	$\frac{83}{128}$	$\frac{9}{128}$	$\frac{9}{32}$
100	0.505	0.005	0.490

Table 29.1: Areas of the three regions in the unit square of cost pairs $(d, \tilde{c})$.

29.3 The Toy Model: A Convex Cost of Greenness

Consumers. A unit mass of consumers, indexed by θ uniformly distributed on $[0, 1]$, choose between a conventional (brown) good and a good sold under a green label. If a share ρ of the labeled goods is truly green, consumer θ values the label at $\theta\rho$ above the brown good. When sales under the label are $\hat{y}$, the consumers with $\theta \geq 1 - \hat{y}$ buy it, so the price premium of the label is

$$P = \rho\,(1 - \hat{y}). \tag{29.3}$$

Firms. The brown good is sold at its cost. Two firms sell under the label. Firm i produces y_i^g truly green units, at an extra cost

$$c(y_i^g) = \frac{y_i^g}{5} + 2\big(y_i^g\big)^2$$

above the brown cost, and y_i^f brown units disguised as green, at an extra cost $d = \frac{1}{10}$ per unit. Each labeled unit pays a net certification cost τ, which is a certification fee ϕ minus any tax t on the brown good, $\tau = \phi - t$, since a brown tax raises the price of both goods. Purity is

$$\rho = \frac{y_1^g + y_2^g}{\hat{y}}, \qquad \text{where} \qquad \hat{y} = y_1^g + y_1^f + y_2^g + y_2^f,$$

and firm i earns

$$\pi_i = [\rho(1 - \hat{y}) - \tau]\left(y_i^g + y_i^f\right) - d\,y_i^f - c(y_i^g). \tag{29.4}$$

Consumers cannot see which firm's units are green; they only know purity. Firms choose green and disguised output simultaneously.

29.4 Second Step: Green Units versus Disguised Units

In a symmetric equilibrium, each firm sells $\frac{\hat{y}}{2}$, and the first-order conditions for disguised and green output are

$$\rho\left(\tfrac{1}{2} - \hat{y}\right) - \tau - d \leq 0, \tag{29.5}$$

$$\rho\left(\tfrac{1}{2} - \hat{y}\right) + \frac{1 - \hat{y}}{2} - \tau - \frac{1}{5} - 2\rho\hat{y} = 0. \tag{29.6}$$

Condition (29.5) holds with equality when firms commit fraud. Both conditions share the term $\rho\left(\frac{1}{2}-\hat{y}\right)$, the marginal revenue from one more disguised unit. A green unit earns an extra $\frac{1-\hat{y}}{2}$ because it raises purity, and so the premium on the firm's half of the sales; the other half of this gain goes to the rival. This is the free-riding problem: the label's reputation is common, while the cost of green production is private.

29.5 Third Step: Equilibrium Fraud

Unit purity. If $\rho = 1$, (29.6) gives $\hat{y} = \frac{2(4-5\tau)}{35}$, and (29.5) shows that no firm wants to commit fraud if and only if $\tau \geq \frac{6}{25}$.

Fraud. If $\tau < \frac{6}{25}$, both conditions hold with equality. Subtracting them gives $\hat{y} = \frac{4}{5(1+4\rho)}$, and substituting into (29.5), purity solves

$$2\rho^2 - \left(\tfrac{7}{10} + 4\tau\right)\rho - \left(\tau + \tfrac{1}{10}\right) = 0. \tag{29.7}$$

Without any policy, $\tau = 0$, purity is

$$\rho = \frac{7+\sqrt{129}}{40} \approx 0.459,$$

so more than half of the labeled goods are disguised brown goods. The quadratic term in $c(y)$ makes each firm's problem concave, and every equilibrium in Sections 29.4 to 29.8 survives all unilateral deviations.

Result 2 (market structure, Proposition 2 of the paper). With $\tau = 0$:

- a monopolist never commits fraud, since its fraud condition is $-\rho\hat{y} - \tau - d < 0$;
- purity falls with the number of firms, from 1 with one firm to 0.459 with two, 0.291 with three, and 0.229 with four;
- with ten or more firms, no symmetric equilibrium with green production exists, since sales under the label would have to be zero.

Figure 29.2 depicts purity as a function of the number of firms, with $\tau = 0$ and treating n as continuous. Purity equals 1 until n reaches

$$\frac{5\sqrt{57}-35}{2} \approx 1.37,$$

where fraud starts, and falls afterward.

Intuitively, a monopolist bears the whole cost of lowering purity, since it sells every labeled unit. With more firms, each firm captures a smaller share of the premium its green production creates, so free riding grows, and with enough firms green production cannot be sustained.

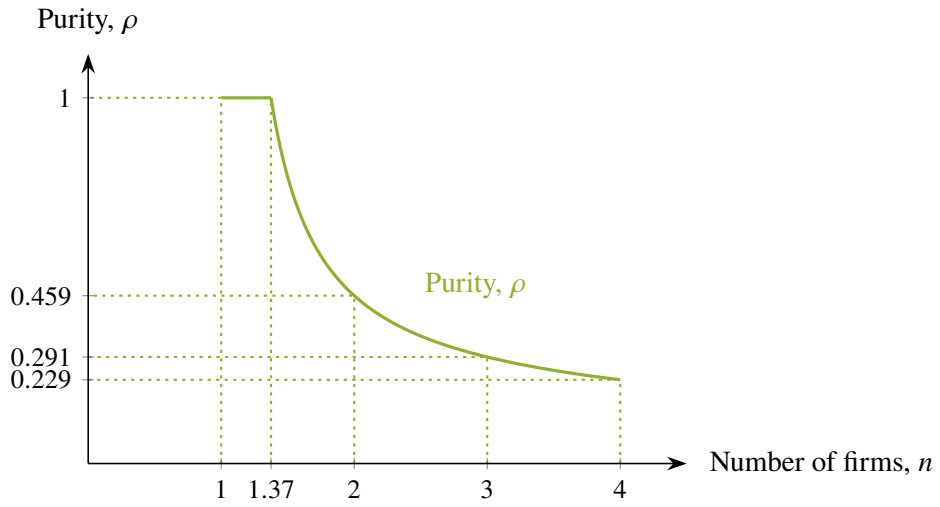


Figure 29.2: Purity of the label as a function of the number of firms, with $\tau = 0$.

29.6 Fourth Step: Certification Fees and Brown Taxes

Result 3 (eco-certification and environmental taxes, Section 4 of the paper). With two firms:

- a unit certification fee raises purity, from 0.459 at $\tau = 0$ to $\frac{3}{4}$ at $\tau = \frac{1}{8}$, and eliminates fraud at $\tau = \frac{6}{25}$;
- the fee also raises each firm's profit, from 0.0175 at $\tau = 0$ to $\frac{3}{160}$ at $\tau = \frac{1}{8}$ and $\frac{12}{625} = 0.0192$ at $\tau = \frac{6}{25}$, the fee that firms would choose for themselves;
- a tax on the brown good lowers τ and raises fraud: a tax of $\frac{1}{20}$ lowers purity to $\frac{5+\sqrt{65}}{40} \approx 0.327$.

Figure 29.3 depicts purity and sales under the label as functions of τ, and Figure 29.4 depicts profit per firm. Purity rises with τ until it reaches 1 at $\tau = \frac{6}{25}$, while sales fall from 0.347 at $\tau = -\frac{1}{20}$ to $\frac{4}{25}$ at $\tau = \frac{6}{25}$. Profit rises until fraud disappears and falls afterward, as with a standard cost increase.

Intuitively, a unit fee lowers the margin on every labeled unit by the same amount. For a disguised unit, that margin is all the firm gets, so firms cut disguised units until the premium has risen enough to cover the fee; a green unit also earns its share of the purity gain, $\frac{1-\hat{y}}{2}$, and the fee drops out of the comparison between green and disguised units, so green production does not fall: each firm's green output,

$$\frac{\rho\hat{y}}{2} = \frac{2\rho}{5(1+4\rho)},$$

rises with purity. Purity rises, and the higher premium more than compensates firms for the fee. A brown tax works in reverse, since it is equivalent to a subsidy to every labeled unit, green or not.

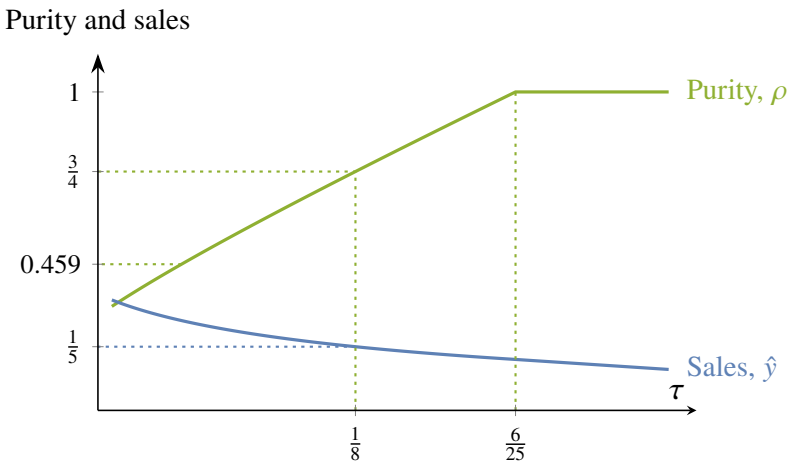


Figure 29.3: Purity (green) and sales under the label (blue) as functions of the net certification cost τ, with two firms.

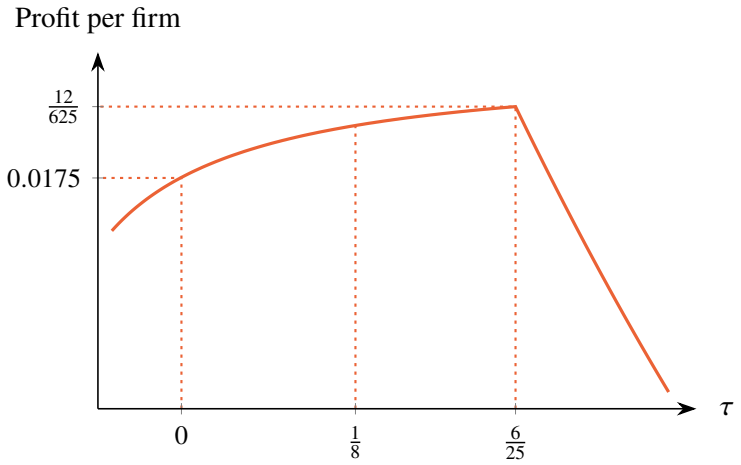


Figure 29.4: Profit per firm as a function of the net certification cost τ, with two firms.

29.7 Fifth Step: The Social Optimum

Total welfare is the consumers' valuation of the labeled goods minus the extra costs,

$$W = \rho\left(\hat{y} - \frac{\hat{y}^2}{2}\right) - d(1-\rho)\hat{y} - 2\,c\left(\frac{\rho\hat{y}}{2}\right).$$

Result 4 (social optimum, Propositions 1 and 3 of the paper). In the toy model:

- the first best has no fraud, $\rho^* = 1$, and sales $\hat{y}^* = \frac{4}{15}$, with welfare $\frac{8}{75}$;
- without fraud, firms would sell $\hat{y}^*$ only with a subsidy, $\tau^* = -\frac{2}{15}$, since market power restricts output;
- this subsidy keeps the market fraud-free only if the disguise cost is at least $\frac{11}{30}$; with $d = \frac{1}{10}$, firms gain from committing fraud, so the first best cannot be reached with unit fees alone.

Intuitively, the regulator would like a subsidy to expand output and a tax to deter fraud, but a unit fee applies equally to green and disguised units. Welfare in the toy model is highest at the fraud-free fee, $\frac{56}{625} \approx 0.090$, still below the first best $\frac{8}{75} \approx 0.107$.

29.8 Sixth Step: Extension to Free Entry

Consider that each firm also pays a fixed certification fee F and that firms enter until profit equals F, treating the number of firms n as continuous. With $\tau = \frac{1}{8}$, Table 29.2 reports the number of firms, purity, and the fixed fee that supports each number of firms.

Fixed fee, F	0.0324	$\frac{3}{160}$	0.0103	0.0061	0.0024
Number of firms, n	1.5	2	2.5	3	4
Purity, ρ	1	$\frac{3}{4}$	0.595	0.514	0.426

Table 29.2: Free entry with $\tau = \frac{1}{8}$: fixed fee, number of firms, and purity.

Result 5 (fixed certification fees, Section 4 of the paper). With free entry, a higher fixed certification fee reduces the number of firms and raises purity.

Intuitively, a fixed fee is an entry barrier, which would usually hurt consumers. Here, fewer firms means less free riding on the label's reputation, so purity, and with it the value of the label, rises.

29.9 Seventh Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Collective reputation: purity, not individual quality, sets the premium, so firms free ride and fraud grows with the number of firms.
- Unit certification fees that reduce fraud and raise profits, and brown taxes that raise fraud.
- The limits of fees alone in reaching the social optimum, and the effect of fixed fees under free entry.

What it drops.

- The brown market and its pollution externality, which the paper uses to set the brown tax.
- Monitoring and enforcement (Section 5 of the paper): with detection rates that do not depend on purity, monitoring targets fraud and supports the optimum, while detection that falls with purity does not help.
- The constant marginal cost of the paper's Figure 1, for the reason given in footnote 1.

29.10 Connections to Other Chapters

- In Chapter 28, brown products obtain a green label because the certification test makes errors, not because of fraud.
- Chapter 27 shows how consumers who do not know a label's standard draw inferences from the products that carry it.
- Chapter 23 studies another market where buyers cannot tell good products from bad ones, and the government revives it by buying the worst assets.

References

Hamilton, S. F., and D. Zilberman (2006). "Green Markets, Eco-Certification, and Equilibrium Fraud." *Journal of Environmental Economics and Management*, 52(3), pp. 627–644.

Part VII

Information Design and Persuasion

The two chapters in this part study a party that chooses how precise the information seen by another party is, rather than signaling what it already knows, as the firms of Chapters 24 through 26 do. Chapter 30 works through four papers. Kamenica and Gentzkow (2011) show that a Sender who commits to the way information is generated can change a Receiver's action to his advantage, even when the Receiver is a rational Bayesian who knows the Sender's goal. The other three papers make information costly to produce, let the two parties disagree about the prior, and let several Senders compete. Chapter 31 studies the opposite incentive. An agency examines applications, such as patents or merger requests, and applicants with bad projects make their applications harder to evaluate. Obfuscation is a costly way to make the agency's information less precise, so more examination resources can leave the agency worse off, because they raise the effort that applicants spend on obfuscation.

In both chapters, the precision of information is chosen by an interested party: the Sender of Chapter 30 designs it to persuade, and the applicant of Chapter 31 lowers it to avoid scrutiny.

Chapter 30

Bayesian Persuasion: Four Papers

Based on Kamenica and Gentzkow (2011) and related papers.

30.1 Introduction

A prosecutor wants the judge to convict. He cannot pay or threaten her. He can only choose which evidence to collect before she rules, and she knows his goal. Can he still raise the conviction rate when the judge is a rational Bayesian? Kamenica and Gentzkow (2011) ask this question for any Sender who controls what a Receiver learns, such as a firm that chooses how much to disclose about its environmental record to a green-minded consumer.

The judge convicts only when her belief that the defendant is guilty reaches one half, and without evidence she acquits. The prosecutor cannot raise her belief on average: averaged over all messages, her beliefs must return to her prior, a requirement called *Bayes plausibility*. He can, however, redistribute her beliefs. He sends a "convict" message for every guilty defendant and for just enough innocent ones that the judge is still willing to convict, and an "acquit" message otherwise. The acquit message costs him nothing, because the judge would acquit anyway.

The answer is therefore yes. When fewer than half of the defendants are guilty, the prosecutor convicts twice the share of guilty defendants: with 30 percent guilty, a rational judge convicts 60 percent of all defendants. The most he can earn is the *concave closure* of his payoff as a function of the judge's belief (the lowest concave curve above it), evaluated at the prior. With a firm as the Sender, the same result is a theory of greenwashing.

Steps 1 to 8 solve the prosecutor's problem by hand, and a figure draws the concave closure. The same picture then solves three extensions, each adding one ingredient and each returning the baseline at one parameter value: a costly investigation (Gentzkow and Kamenica 2014), different priors for the prosecutor and the judge (Alonso and Cámara 2016a), and a defense attorney who also designs evidence (Gentzkow and Kamenica 2017). A tour of the broader literature and of open questions closes the chapter.

30.2 Costless Persuasion (Kamenica and Gentzkow, 2011)

30.2.1 The model

States. There are two states of the world, $\omega \in \{G, I\}$, where G means the defendant is guilty and I means innocent. The prosecutor (Sender) and the judge (Receiver) share a common prior that the defendant is guilty. We call this prior μ_0 and keep it as a free parameter, so that every result below is a formula in μ_0; our running numerical illustration sets $\mu_0 = 0.3$. Because there are only two states, any belief is a single number, the probability it places on guilt, and we write $\mu \equiv \Pr(\text{guilty})$, so the prior places

$$\mu_0(G) = \mu_0 \qquad \text{and} \qquad \mu_0(I) = 1 - \mu_0.$$

We assume $\mu_0 < \frac{1}{2}$: a judge with no information would acquit, which is the case in which persuasion has any work to do. The parameter μ_0 is the one free parameter we carry from start to finish; each extension adds one more ingredient of its own, but the prior stays generic throughout.

Actions. The judge chooses one of two actions, $a \in \{\text{convict}, \text{acquit}\}$.

Payoffs. The judge wants to match the verdict to the truth. She earns 1 from the just action, convicting the guilty or acquitting the innocent, and 0 from the unjust action:

$$u(\text{convict}, G) = 1, \quad u(\text{acquit}, I) = 1, \quad u(\text{convict}, I) = 0, \quad u(\text{acquit}, G) = 0.$$

The prosecutor wants a conviction regardless of the truth. He earns 1 whenever the judge convicts and 0 whenever she acquits:

$$v(\text{convict}, \omega) = 1 \quad \text{and} \quad v(\text{acquit}, \omega) = 0 \quad \text{for both } \omega.$$

Notice that the prosecutor's payoff depends only on the judge's *action*, not on the state. This will simplify the geometric picture we draw later.

The instrument. The prosecutor chooses an *investigation*: a signal π that generates a message for the judge, with the message drawn from state-dependent distributions $\pi(\cdot \mid G)$ and $\pi(\cdot \mid I)$. He commits to the investigation before the evidence is in, and reports its outcome honestly; he cannot rewrite the rules once he sees what turns up. The judge observes the design of the investigation and its realized message, updates her belief by Bayes' rule, and then convicts or acquits.

When does the judge convict? At any belief μ, the judge convicts if the expected payoff from convicting is at least that of acquitting:

$$\underbrace{\mu \cdot 1 + (1-\mu)\cdot 0}_{\text{convict}} \;\geq\; \underbrace{\mu \cdot 0 + (1-\mu)\cdot 1}_{\text{acquit}} \quad \Longleftrightarrow \quad \mu \geq 1 - \mu \quad \Longleftrightarrow \quad \mu \geq \frac{1}{2}.$$

We record the judge's optimal action as a function of her belief,

$$\bar{a}(\mu) = \begin{cases} \text{convict} & \text{if } \mu \geq \dfrac{1}{2}, \\ \text{acquit} & \text{if } \mu < \dfrac{1}{2}, \end{cases}$$

where we assign a tie in the prosecutor's favor, consistent with the Sender-preferred equilibrium of the paper; without this convention the best conviction rate would be a value the prosecutor could approach but never quite reach.

30.2.2 Two benchmarks

Before designing anything clever, we compute what happens at the two extremes. These give the prosecutor a floor to improve upon.

Example 1 (No disclosure)**.** Consider that the investigation is completely uninformative, so the judge learns nothing and keeps her prior μ_0. Since $\mu_0 < \frac{1}{2}$, the judge acquits, and the prosecutor's payoff is

$$v_{\text{no info}} = 0.$$

We call *acquit* the *default action*, since it is what the judge does when she has no information beyond the prior. □

Example 2 (Full disclosure)**.** Consider that the investigation reveals the state with certainty. Then the judge's posterior is either $\mu = 1$ (and she convicts) or $\mu = 0$ (and she acquits). She convicts exactly when the defendant is truly guilty, which happens with probability μ_0. The prosecutor's expected payoff is

$$v_{\text{full info}} = \mu_0 \cdot 1 + (1 - \mu_0) \cdot 0 = \mu_0 \qquad (= 0.3 \text{ at } \mu_0 = 0.3\,).$$

Full disclosure dominates silence, raising the conviction rate from 0 to μ_0. The question is whether the prosecutor can do better still. □

Intuitively, revealing everything is a crude instrument. It hands the judge a conviction only on the guilty, and it throws away the possibility of convicting on the strength of merely suggestive evidence. We now show that a partial investigation does better than both benchmarks.

30.2.3 The one idea that makes the problem tractable

The prosecutor has infinitely many investigations to choose among, differing in how many messages they send and how each message correlates with the state. Searching over all of them directly looks hopeless. The insight of Kamenica and Gentzkow is that we do not need to. What the prosecutor ultimately cares about is the judge's *belief* when she acts, because the belief is what determines the verdict. So we can search over the beliefs an investigation can produce, rather than over the investigations themselves.

Each message the judge might receive leaves her with some posterior belief μ. An investigation therefore induces a *distribution of posteriors*: a list of the beliefs that can arise, together with the probability of each. The question becomes: which distributions of posteriors are attainable?

The answer is that a distribution of posteriors is attainable if and only if it is *Bayes plausible*, meaning that the beliefs average back to the prior:

$$\mathbb{E}[\mu] = \mu_0.$$

The requirement is nothing more than the law of iterated expectations. No matter what the judge might come to believe after seeing a message, her beliefs, weighted by how likely each message is, must return to where she started. No investigation can, on average, make the judge more convinced of guilt than the prior warrants; it can only redistribute her conviction across messages.

This is the whole content of Proposition 1 in the paper. It lets us restate the prosecutor's problem in a clean form:

$$\max_{\tau \text{ Bayes plausible}} \mathbb{E}_\tau\big[\,\hat{v}(\mu)\,\big],$$

where τ is the distribution of posteriors and $\hat{v}(\mu)$ denotes the prosecutor's payoff when the judge holds belief μ and acts optimally. In our example, since the prosecutor earns 1 whenever the judge convicts,

$$\hat{v}(\mu) = v\big(\bar{a}(\mu)\big) = \begin{cases} 1 & \text{if } \mu \geq \dfrac{1}{2}, \\ 0 & \text{if } \mu < \dfrac{1}{2}. \end{cases}$$

30.2.4 Solving the prosecutor's problem

We now solve the problem step by step. It is convenient to work with *straightforward* investigations: those whose messages are recommendations, "convict" or "acquit," which the judge is willing to obey. The paper shows this is without loss of generality, just as the revelation principle lets us restrict attention to truthful mechanisms. There are only two messages to design, so the problem reduces to two numbers. Two messages are all the prosecutor ever needs: because the judge takes only one of two actions, a signal with more realizations cannot raise his payoff.

Step 1. Name the choice variables. Let

$$\sigma_G = \pi(g \mid G) \qquad \text{and} \qquad \sigma_I = \pi(g \mid I)$$

be the probabilities that the investigation sends the message g ("recommend convict") in each state. The complementary probabilities, $1-\sigma_G$ and $1-\sigma_I$, are the chances of the message i ("recommend acquit").

Step 2. Write the objective. If the judge obeys, the prosecutor wins a conviction exactly when the message is g. That message arrives with total probability

$$\Pr(g) = \mu_0\,\sigma_G + (1-\mu_0)\,\sigma_I,$$

so the prosecutor maximizes $\mu_0\,\sigma_G + (1-\mu_0)\,\sigma_I$.

Step 3. Impose obedience on the conviction recommendation. For the judge to actually convict when she sees g, her posterior on guilt given g must be at least $\frac{1}{2}$. By Bayes' rule,

$$\mu_g \equiv \Pr(G \mid g) = \frac{\mu_0\,\sigma_G}{\mu_0\,\sigma_G + (1-\mu_0)\,\sigma_I} \;\geq\; \frac{1}{2}.$$

Cross-multiplying, this obedience condition is

$$\mu_0\,\sigma_G \;\geq\; (1-\mu_0)\,\sigma_I.$$

Step 4. Impose obedience on the acquittal recommendation. For the judge to acquit when she sees i, her posterior on guilt given i must be at most $\frac{1}{2}$:

$$\mu_i \equiv \Pr(G \mid i) = \frac{\mu_0 (1 - \sigma_G)}{\mu_0 (1 - \sigma_G) + (1 - \mu_0)(1 - \sigma_I)} \;\leq\; \frac{1}{2},$$

which simplifies to

$$\mu_0 (1 - \sigma_G) \;\leq\; (1 - \mu_0)(1 - \sigma_I).$$

Step 5. Argue that a guilty defendant is always recommended for conviction. The objective $\Pr(g)$ rises in σ_G, and raising σ_G also loosens the conviction-obedience condition in Step 3, since it puts more guilt mass behind the message g. There is thus no reason to hold back: the prosecutor sets

$$\sigma_G = 1.$$

Intuitively, evidence pointing at a guilty defendant can only help the case, so the optimal investigation never recommends acquittal for someone who is in fact guilty. This is the example's version of the general rule that when the judge takes the prosecutor's least-preferred action, she is certain that this action is right (Proposition 4 of the paper), which in our two-state example means certain that the defendant is innocent.

Step 6. Convict as many innocents as obedience allows. With $\sigma_G = 1$, the objective becomes $\mu_0 + (1 - \mu_0)\,\sigma_I$, which rises in σ_I. The prosecutor would like σ_I as large as possible, but Step 3 caps it at $\mu_0 \geq (1 - \mu_0)\,\sigma_I$. The constraint binds, so

$$\sigma_I = \frac{\mu_0}{1 - \mu_0}.$$

We check that Step 4 is satisfied:

$$\mu_0(1 - 1) = 0 \;\leq\; (1 - \mu_0)\Big(1 - \frac{\mu_0}{1 - \mu_0}\Big) = 1 - 2\mu_0,$$

which is positive because $\mu_0 < \frac{1}{2}$, so the judge is indeed willing to acquit on the message i.

Step 7. Read off the optimal investigation. Translating (σ_G, σ_I) back into the language of the two messages,

$$\pi(g \mid G) = 1, \qquad \pi(g \mid I) = \frac{\mu_0}{1 - \mu_0},$$

$$\pi(i \mid G) = 0, \qquad \pi(i \mid I) = \frac{1 - 2\mu_0}{1 - \mu_0}.$$

At the running prior $\mu_0 = 0.3$ this is $\big(1,\ \frac{3}{7},\ 0,\ \frac{4}{7}\big)$, precisely the investigation reported in equation (1) of Kamenica and Gentzkow (2011).

Step 8. Compute the conviction rate. The prosecutor's payoff equals the probability of the message g:

$$v^* = \Pr(g) = \mu_0 (1) + (1 - \mu_0)\,\frac{\mu_0}{1 - \mu_0} = \mu_0 + \mu_0 = 2\mu_0.$$

The optimal investigation therefore convicts with probability $2\mu_0$, twice the full-disclosure rate μ_0, whenever $\mu_0 < \frac{1}{2}$. At $\mu_0 = 0.3$ this is 0.6: a judge who knows that 70 percent of defendants are innocent convicts 60 percent of them, while fully aware that the investigation was built to maximize convictions.

Why the judge goes along

The prosecutor and the judge are not pure adversaries: both want to convict the guilty, so some messages serve them at once. And because the prosecutor commits to the investigation in advance and the judge knows its design, she can trust what a message implies rather than second-guess his motive. With that in place, it is worth reading the two posteriors the investigation produces, because they show there is no trick. On the message i, the posterior is

$$\mu_i = \frac{\mu_0 \cdot 0}{\mu_0 \cdot 0 + (1-\mu_0) \cdot \frac{1-2\mu_0}{1-\mu_0}} = 0,$$

so the judge is *certain* the defendant is innocent and acquits without hesitation. On the message g, the posterior is

$$\mu_g = \frac{\mu_0 \cdot 1}{\mu_0 \cdot 1 + (1-\mu_0) \cdot \frac{\mu_0}{1-\mu_0}} = \frac{\mu_0}{\mu_0 + \mu_0} = \frac{1}{2},$$

so the judge is indifferent between convicting and acquitting, and, breaking the tie in the prosecutor's favor, she convicts. The prosecutor has engineered the weakest belief consistent with a conviction: just enough to tip the judge over, and not a shade more.

We can also confirm Bayes plausibility directly. The message g arrives with probability $2\mu_0$ and leaves belief $\frac{1}{2}$; the message i arrives with probability $1 - 2\mu_0$ and leaves belief 0. Averaging,

$$2\mu_0 \cdot \frac{1}{2} + (1 - 2\mu_0) \cdot 0 = \mu_0,$$

so the two posteriors do return to the prior, as they must.

30.2.5 Graphical representation

The problem has a picture that gives the answer at once, and it generalizes far beyond this example. Plot the prosecutor's payoff $\hat{v}(\mu)$ against the judge's belief μ. As computed above, $\hat{v}$ is a step function: it is 0 for every belief below $\frac{1}{2}$, where the judge acquits, and jumps to 1 once the belief reaches $\frac{1}{2}$, where she convicts.

Now recall the prosecutor's problem: choose beliefs that average to the prior μ_0, so as to maximize the average height of $\hat{v}$. Mixing two beliefs traces the straight line joining the two points on the graph of $\hat{v}$. The best the prosecutor can do at any prior is therefore given by the *concave closure* of $\hat{v}$, written V: the lowest concave curve that lies weakly above $\hat{v}$ everywhere. The same concave-closure method appears earlier in Aumann and Maschler (1995), who use it to study repeated games with incomplete information. Figure 30.1 depicts $\hat{v}$, its concave closure V, and the construction at the prior, for $\mu_0 = 0.3$.

Reading the figure recovers everything from the algebra above, and faster. The concave closure is the line $V(\mu) = 2\mu$ for $\mu \le \frac{1}{2}$ and $V(\mu) = 1$ for $\mu > \frac{1}{2}$. Its value at the prior is

$$V(\mu_0) = 2\mu_0 \qquad (= 0.6 \text{ at } \mu_0 = 0.3\,),$$

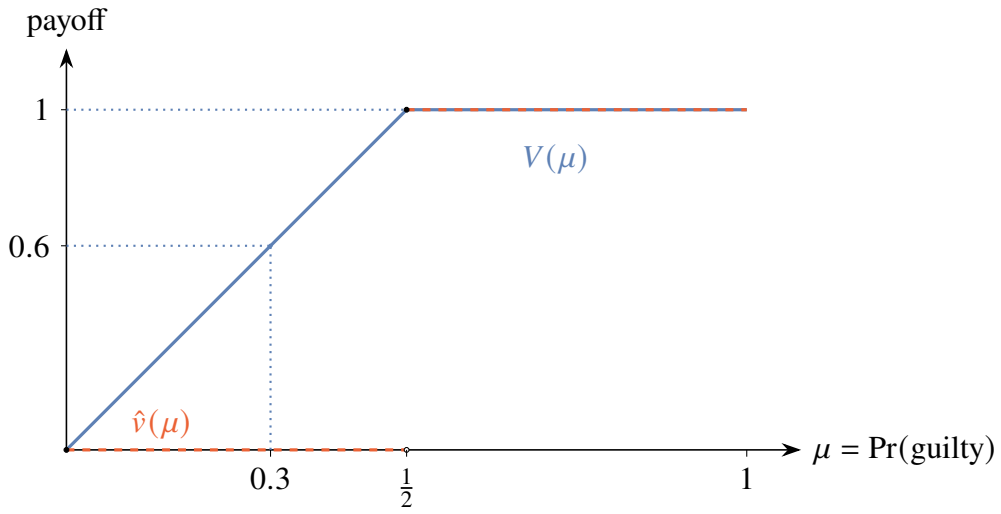


Figure 30.1: The step payoff $\hat{v}$ (dashed), its concave closure V (solid), and the value $V(0.3) = 0.6$ at the prior.

which is the optimal conviction rate. The two points the closure touches, $\mu = 0$ and $\mu = \frac{1}{2}$, are the two posteriors the optimal investigation induces. And the weights that place their average at μ_0 are $1 - 2\mu_0$ on $\mu = 0$ and $2\mu_0$ on $\mu = \frac{1}{2}$, the probabilities of the two messages. The gap between the curves at the prior,

$$V(\mu_0) - \hat{v}(\mu_0) = 2\mu_0 - 0 = 2\mu_0,$$

is the prosecutor's gain from persuasion. Wherever this gap is positive, the Sender benefits; wherever $\hat{v}$ already equals its own concave closure, no investigation helps.

30.2.6 Main points from the above example

The example teaches lessons that recur throughout what follows.

- Persuasion works when the payoff is not concave. The prosecutor gains precisely because $\hat{v}$ lies below its concave closure at the prior. If the prosecutor's payoff were concave in the judge's belief, no spreading of beliefs could raise its average, and silence would be optimal. If it were convex, full disclosure would be optimal. The interesting case, and the realistic one, is the step in between.
- The Sender loses nothing on the downside. Providing information sometimes moves the judge toward guilt and sometimes toward innocence. The move toward guilt earns a conviction; the move toward innocence costs the prosecutor nothing, because the judge was already acquitting by default. Persuasion benefits the Sender whenever the Receiver does not already take the Sender's preferred action, and her action stays fixed over a range of beliefs around the prior.
- The Sender induces the worst belief consistent with each action. When the judge acquits, she is made certain of innocence; when she convicts, she is made just indifferent. The prosecutor never leaves the judge more convinced than the action requires, because any leftover conviction

could be spent buying additional convictions elsewhere. This is the content of Propositions 4 and 5 in the paper, and the reader can see both at work in the two posteriors computed above.

30.2.7 An environmental application: green signaling and greenwashing

The courtroom is one dressing of a structure that recurs, and a second reading will run alongside us through the sections that follow. Read the Sender as a firm and the Receiver as a green-minded consumer, or a regulator. The state is whether the firm's product is genuinely clean or dirty; the prior μ_0 is now the share of clean products, and the consumer buys, or the regulator approves, only when she is convinced enough that the product is clean, say when her belief that it is clean is at least $\frac{1}{2}$. The firm, like the prosecutor, wants the favorable action regardless of the truth, and it chooses what to disclose about its environmental performance: which tests to run, which metrics to report, how to design a sustainability label.

The baseline result then reads as a theory of greenwashing. A firm sure to be exposed as dirty says nothing; a firm with a genuinely clean case reveals it; and in between the firm designs a label that raises a skeptical consumer's belief to the buying threshold and not a shade higher. The optimal green signal is therefore a partial, noisy label that clears the consumer at rate $2\mu_0$, twice the share of products that are actually clean, precisely the "convict 60 percent while only 30 percent are guilty" logic of the courtroom. The three extensions carry over intact: costly testing (Section 30.3) shrinks how much the firm discloses; a firm more optimistic about its own product than its audience is (Section 30.4) values the same label more without changing it; and a watchdog or a rival that can also disclose (Section 30.5) pushes the market toward full revelation of the firm's true type. This environmental reading is the subject of a large literature on certification and greenwashing (Lyon and Maxwell 2011).

30.2.8 A numerical summary

Table 30.1 reports the three investigations, no disclosure, full disclosure, and the optimal partial investigation, at the running prior $\mu_0 = 0.3$. Reading down the conviction column shows the prosecutor's problem in one glance: silence convicts no one, full disclosure convicts only the truly guilty, and the optimal partial investigation doubles the full-disclosure rate. The judge's own welfare, in the last column, is highest under full disclosure, where she always matches the verdict to the truth, and falls back to her uninformed level of 0.7 under the optimal investigation, since the prosecutor leaves her indifferent whenever she convicts.

Investigation	Conviction rate	Prosecutor payoff	Judge welfare
No disclosure	0.00	0.00	0.70
Full disclosure	0.30	0.30	1.00
Optimal (partial)	0.60	0.60	0.70

Table 30.1: The three investigations at the prior $\mu_0 = 0.3$.

How much the prosecutor gains depends on the prior. Since the value of persuasion is $V(\mu_0) =$

$\min(2\mu_0, 1)$ while full disclosure convicts at rate μ_0, the gain over full disclosure is μ_0 for every prior below $\frac{1}{2}$. Once $\mu_0 \geq \frac{1}{2}$ the judge convicts by default, silence already delivers the highest conviction rate, 1, and the gain from persuasion vanishes. Figure 30.2 depicts the two conviction rates as functions of the prior μ_0.

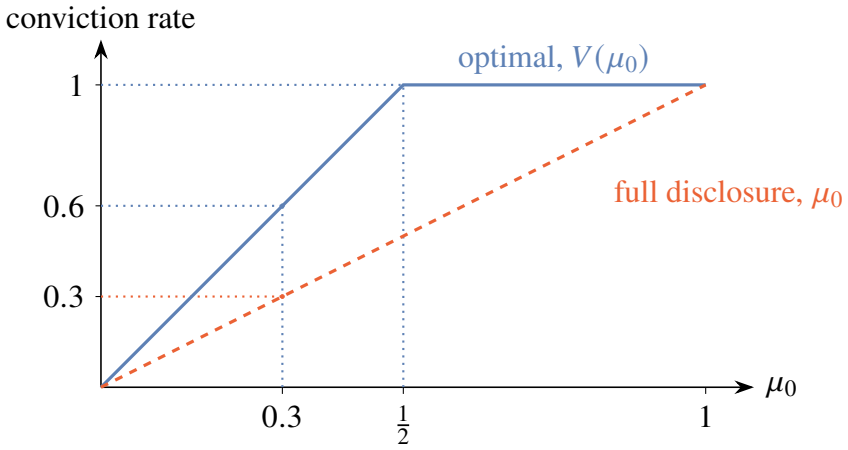


Figure 30.2: Conviction rate under the optimal investigation, $V(\mu_0) = \min(2\mu_0, 1)$, and under full disclosure, μ_0.

30.3 Costly Persuasion (Gentzkow and Kamenica, 2014)

30.3.1 From free information to costly information

The investigation in Section 30.2 was free. A prosecutor, though, does not gather evidence at no charge: forensic tests, expert witnesses, depositions, and subpoenas all consume time and money. It is natural to ask whether the elegant concavification of Section 30.2 survives once the Sender must pay for the information he generates.

At first sight the answer looks like no. The whole method rested on writing the Sender's payoff as a function of the Receiver's posterior alone, $\hat{v}(\mu)$, so that the problem reduced to averaging $\hat{v}$ over Bayes-plausible beliefs. Once signals are costly, this breaks: two investigations can leave the judge at the very same posterior μ while costing the prosecutor different amounts, one being a low-cost, coarse investigation and the other a costly, finely tuned one. The Sender's payoff at a belief then depends on *how* the belief was reached, not on the belief alone, and the concavification picture no longer applies.

Gentzkow and Kamenica (2014), henceforth GK, identify a family of cost functions for which the difficulty disappears. For this family the cost can be absorbed into the value function, producing a *cost-adjusted* payoff that again depends on the posterior alone, so the same picture goes through. And when the cost vanishes, the model returns to Kamenica and Gentzkow (2011), henceforth KG. The costless model of Section 30.2 is thus the special case of a one-parameter family, recovered when that parameter equals zero.

30.3.2 The model with costs

The states, actions, payoffs, and prior are those of Section 30.2: two states $\omega \in \{G, I\}$ with the same generic prior $\mu_0 = \Pr(\text{guilty}) < \frac{1}{2}$ (running value 0.3); the judge convicts if and only if her belief satisfies $\mu \geq \frac{1}{2}$; and the prosecutor earns 1 from a conviction and 0 from an acquittal. The only change is that running an investigation π now costs the prosecutor $c(\pi)$, so his payoff becomes

$$v(a, \omega) - c(\pi).$$

GK's main assumption is that the cost of an investigation is proportional to the *expected reduction in uncertainty* it produces. Let $H(\mu)$ be a measure of the uncertainty in a belief μ: a concave function that is largest when the belief is most spread out and zero when the belief is degenerate. Then the assumption is

$$c(\pi) = k\, \mathbb{E}\big[\, H(\mu_{\text{ref}}) - H(\mu_s) \,\big],$$

where μ_{ref} is a fixed reference belief, μ_s the posterior after message s, and $k \geq 0$ a cost parameter. Concavity of H guarantees, by Jensen's inequality, that information reduces uncertainty on average, so the cost is never negative, and it is zero for an uninformative investigation.

Two measures of uncertainty are standard. The first is Shannon entropy,

$$H(\mu) = -\sum_{\omega} \mu(\omega) \ln \mu(\omega),$$

which underlies the rational-inattention literature. The second is residual variance,

$$H(\mu) = \sum_{\omega} \mu(\omega)\big(1 - \mu(\omega)\big),$$

which in our two-state model is

$$H(\mu) = 2\mu(1 - \mu).$$

We adopt residual variance, and we place the reference belief at the prior, $\mu_{\text{ref}} = \mu_0$, so that the cost measures the uncertainty an investigation removes relative to what the judge knew at the outset. Both are choices GK explicitly allow. This pair keeps the algebra in closed form and lets us see every step; GK's own headline example instead uses entropy with a uniform reference belief, and we report that version in the numerical illustration below. With the reference at the prior the cost simplifies to

$$c(\pi) = k\big(\, H(\mu_0) - \mathbb{E}_\tau[H(\mu)] \,\big),$$

where τ is, as before, the distribution of posteriors the investigation induces.

30.3.3 The unifying result: concavify a cost-adjusted payoff

The Sender's problem is again to choose a Bayes-plausible distribution of posteriors, but now net of cost:

$$\max_{\tau \text{ Bayes plausible}} \; \mathbb{E}_\tau[\hat{v}(\mu)] - c(\pi) = \max_{\tau} \; \mathbb{E}_\tau[\hat{v}(\mu)] - k\big(H(\mu_0) - \mathbb{E}_\tau[H(\mu)]\big)$$

$$= \max_{\tau} \ \mathbb{E}_{\tau}\left[\hat{v}(\mu) + k\,H(\mu)\right] - k\,H(\mu_0).$$

The last line is the key. The cost has been absorbed into a single *cost-adjusted value function*

$$\hat{v}_k(\mu) \equiv \hat{v}(\mu) + k\,H(\mu),$$

which depends on the posterior alone, so the concavification method of Section 30.2 applies without change. Writing V_k for the concave closure of $\hat{v}_k$, the Sender's value is

$$V_k(\mu_0) - k\,H(\mu_0),$$

and the Sender benefits from persuasion if and only if $V_k(\mu_0) > \hat{v}_k(\mu_0)$. This is GK's Proposition 1, specialized to our cost.

The nesting. Set $k = 0$. Then $\hat{v}_k = \hat{v}$, the concave closure V_k coincides with V from Section 30.2, and the constant $k\,H(\mu_0)$ vanishes. The Sender's value is $V(\mu_0) = 2\mu_0$, the costless result of Section 30.2. Kamenica and Gentzkow (2011) is the $k = 0$ member of the family.

Intuitively, adding the concave term $k\,H(\mu)$ to $\hat{v}$ makes the payoff more concave. Since persuasion feeds on the gap between a function and its concave closure, and a more concave function leaves a smaller gap, information costs shrink the room for persuasion. The subtracted constant $k\,H(\mu_0)$ does not depend on the investigation, so it does not affect which investigation is optimal; it equals the cost of revealing the state fully.

30.3.4 The toy model, worked

Step 1. Write the cost-adjusted payoff. With $\hat{v}(\mu)$ the step function of Section 30.2 and $H(\mu) = 2\mu(1-\mu)$,

$$\hat{v}_k(\mu) = \begin{cases} 2k\,\mu(1-\mu) & \text{if } \mu < \dfrac{1}{2}, \\[2ex] 1 + 2k\,\mu(1-\mu) & \text{if } \mu \geq \dfrac{1}{2}. \end{cases}$$

Each piece is a downward parabola, and the two pieces are separated by an upward jump of size 1 at the judge's indifference belief $\mu = \frac{1}{2}$.

Step 2. The conviction belief stays at the threshold. To win a conviction the prosecutor must leave the judge with a posterior $\mu_h \geq \frac{1}{2}$. Raising μ_h above $\frac{1}{2}$ buys no extra conviction, since the judge already convicts, and it is costly, since H is falling above $\frac{1}{2}$ and a sharper belief removes more uncertainty. The prosecutor therefore leaves the judge indifferent, $\mu_h = \frac{1}{2}$, just as in Section 30.2.

Step 3. The acquittal belief now involves a trade-off. Let $\mu_l < \frac{1}{2}$ be the belief that leads the judge to acquit. In Section 30.2 the prosecutor drove μ_l all the way to 0, making the judge certain of innocence. That is no longer clearly best. Lowering μ_l toward 0 still raises the conviction probability, because Bayes plausibility then allows more weight on the conviction message; but it also raises the cost, because a belief nearer certainty carries less uncertainty and the cost rewards uncertainty left in place. The prosecutor must weigh the two.

Step 4. Solve for the acquittal belief. With $\mu_h = \frac{1}{2}$ and weight p on the conviction message, Bayes plausibility reads

$$p \cdot \frac{1}{2} + (1-p)\,\mu_l = \mu_0,$$

so

$$p = \frac{\mu_0 - \mu_l}{\frac{1}{2} - \mu_l}.$$

The prosecutor's net payoff, gross conviction probability p minus cost, is

$$N(\mu_l) = p - k\Big(H(\mu_0) - \big[\, p \cdot H(\tfrac{1}{2}) + (1-p)\,H(\mu_l)\,\big]\Big),$$

with $H(\frac{1}{2}) = \frac{1}{2}$ and $H(\mu_0) = 2\mu_0(1-\mu_0)$. Differentiating and evaluating at $\mu_l = 0$ gives the clean expression

$$\left.\frac{dN}{d\mu_l}\right|_{\mu_l=0} = 2(1-2\mu_0)\Big(\frac{k}{2} - 1\Big),$$

which at $\mu_0 = 0.3$ is $-0.8 + 0.4\,k$.

Step 5. Read off the three regimes. Because $1 - 2\mu_0 > 0$, the sign of the derivative is the sign of $\frac{k}{2} - 1$, so the threshold sits at $k = 2$ *for every prior* $\mu_0 < \frac{1}{2}$. This sorts the problem into three cases.

- If $k \le 2$, the derivative is not positive, so the prosecutor keeps $\mu_l = 0$. The optimal investigation is *identical* to the costless one of Section 30.2: the judge is made certain of innocence on the acquittal message and indifferent on the conviction message, and the conviction rate stays at $2\mu_0$. Only the prosecutor's net payoff falls, to

  $$N = 2\mu_0 - k\,\mu_0(1-2\mu_0) = \mu_0\big[\,2 - k(1-2\mu_0)\,\big],$$

  which at $\mu_0 = 0.3$ is $0.6 - 0.12\,k$.
- If $2 < k < k^*(\mu_0)$, the derivative at $\mu_l = 0$ is positive, so the prosecutor raises μ_l above 0. He no longer fully clears the innocent, the acquittal message stops implying certain innocence, and the conviction rate falls below $2\mu_0$.
- If $k \ge k^*(\mu_0)$, no investigation is worthwhile, and the prosecutor gathers no evidence.

The knife-edge at $k = 2$ is where a binary verdict first makes it worth economizing on information; the upper cutoff, where persuasion stops altogether, is the cost at which the optimal acquittal belief reaches the prior,

$$k^*(\mu_0) = \frac{2}{(1-2\mu_0)^2},$$

which depends on the prior and equals 12.5 at $\mu_0 = 0.3$.

Step 6. Recover the investigation at an interior cost. Return to the running prior $\mu_0 = 0.3$ and take $k = 3$, in the middle regime. Solving $N'(\mu_l) = 0$ yields $\mu_l = 0.0918$, hence $p = 0.510$. The two induced posteriors are $\mu_l = 0.0918$ with weight $1 - p = 0.490$ and $\mu_h = \frac{1}{2}$ with weight $p = 0.510$. We recover the investigation from

$$\pi(s \mid \omega) = \frac{\mu_s(\omega)\,\tau(\mu_s)}{\mu_0(\omega)},$$

which gives

$$\pi(g \mid G) = \frac{0.5 \times 0.510}{0.3} = 0.850, \qquad \pi(g \mid I) = \frac{0.5 \times 0.510}{0.7} = 0.364,$$
$$\pi(i \mid G) = \frac{0.0918 \times 0.490}{0.3} = 0.150, \qquad \pi(i \mid I) = \frac{0.9082 \times 0.490}{0.7} = 0.636.$$

The conviction rate is

$$\Pr(g) = 0.3(0.850) + 0.7(0.364) = 0.510,$$

down from 0.6. Notice that $\pi(i \mid G) = 0.150 > 0$: the prosecutor now recommends acquittal for some guilty defendants. This is the sharpest contrast with Section 30.2. When information was free, the acquittal message meant certain innocence (Proposition 4 of KG); once information is costly, buying that certainty is too expensive, so the prosecutor accepts an acquittal message that leaves the judge with lingering doubt.

30.3.5 What cost adds to the picture

Costly persuasion reshapes the baseline in ways worth drawing out.

- The costless model is a boundary, not a different model. Everything in Section 30.3 collapses to Section 30.2 at $k = 0$. The extension does not overturn Bayesian persuasion; it embeds it in a family and asks what happens as one moves away from the free-information boundary. This is the shape of a good extension: the earlier result sits inside the newer one at a single parameter value, and the new results describe the neighborhood around it.
- Costs shrink disclosure, and do so in stages. As k rises, the optimal investigation is first unchanged, then Blackwell less informative, then silent. The prosecutor first absorbs the cost while running the same investigation, then economizes by tolerating a noisier acquittal signal, and finally abandons persuasion. Lower costs always yield a weakly more informative investigation in this example, as GK find in their own version of it; GK also show that this comparative static need not hold in general.
- The Receiver captures none of the value of information. A short calculation shows the judge's expected payoff under the optimal investigation is

  $$p \cdot \tfrac{1}{2} + (1 - p)(1 - \mu_l) = 1 - \mu_0 \quad \text{for every } k,$$

  where we use $p \cdot \frac{1}{2} + (1 - p)\mu_l = \mu_0$. This is the judge's payoff with no information at all, $\Pr(\text{innocent}) = 1 - \mu_0$ (namely 0.7 at $\mu_0 = 0.3$). For every cost of investigation, the prosecutor leaves the judge no better off than her uninformed benchmark and keeps the entire value of information for himself.

30.3.6 A numerical summary

Table 30.2 traces the optimal investigation as the cost parameter k rises, under the residual-variance cost worked out above. The three regimes are visible in the columns. For $k \leq 2$ the acquittal belief

stays at $\mu_l = 0$ and the conviction rate holds at 0.6, the costless investigation; only the prosecutor's net payoff falls, linearly, by 0.12 per unit of k. For $2 < k < k^*$ the acquittal belief lifts off 0 and the conviction rate slides down. At $k^* = 12.5$ the prosecutor stops investigating. The judge's welfare, not shown, equals 0.7 in every row.

k	μ_l	Conviction rate	Prosecutor net	Regime
0	0.000	0.600	0.600	identical to KG
1	0.000	0.600	0.480	signal = KG
2	0.000	0.600	0.360	threshold
3	0.092	0.510	0.260	interior
4	0.146	0.434	0.189	interior
6	0.211	0.307	0.094	interior
10	0.276	0.106	0.011	interior
12.5	0.300	0.000	0.000	no disclosure

Table 30.2: Optimal investigation as the cost k rises, residual-variance cost with reference belief $\mu_0 = 0.3$.

Figure 30.3 depicts the case $k = 3$ the way Figure 30.1 depicted the costless problem. The cost bends each branch of the payoff into a parabola and lifts the whole picture, but the method is unchanged: the concave closure V_3 of the cost-adjusted payoff $\hat{v}_3$ still touches the graph at the two induced posteriors, now $\mu_l = 0.092$ and $\mu_h = \frac{1}{2}$. Reading the closure at the prior gives $V_3(0.3) = 1.52$, and the prosecutor's payoff is

$$V_3(0.3) - k\,H(\mu_0) = 1.52 - 1.26 = 0.26,$$

matching Table 30.2.

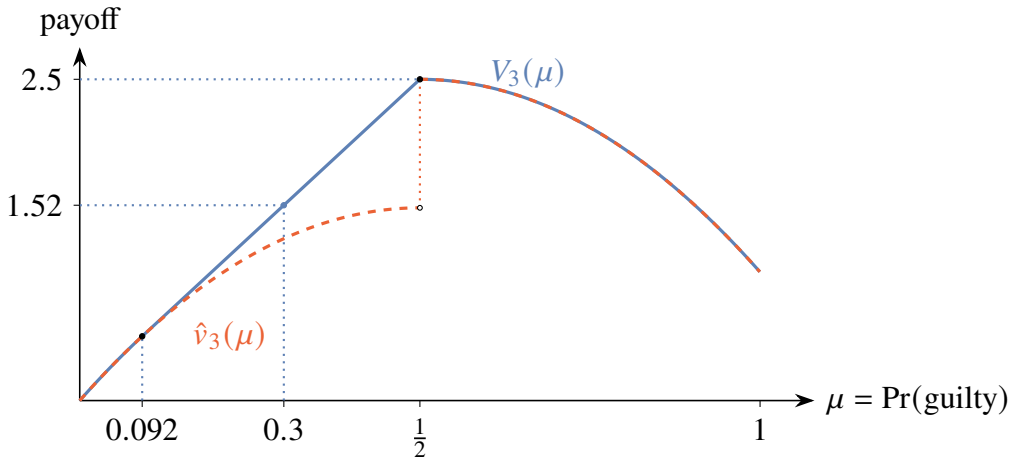


Figure 30.3: The cost-adjusted payoff $\hat{v}_3$ (dashed) and its concave closure V_3 (solid) at $k = 3$, with induced posteriors $\mu_l = 0.092$ and $\mu_h = \frac{1}{2}$.

Finally, Table 30.3 reports GK's own version of the example, which uses Shannon entropy as the measure of uncertainty and the uniform belief as the reference, rather than the residual variance

k	Conviction rate	μ_i	Prosecutor net
0	0.600	0.000	0.600
1	0.543	0.062	0.384
2	0.457	0.131	0.259
3	0.382	0.176	0.177
4	0.315	0.208	0.119
5	0.255	0.232	0.077
7	0.147	0.265	0.025
10	0.009	0.298	0.000

Table 30.3: GK's entropy version of the example, reference belief uniform.[1]

and prior reference used above. The pattern is the same, and $k = 0$ again returns the costless investigation of Section 30.2.

30.4 Persuasion under Disagreement (Alonso and Cámara, 2016a)

30.4.1 From a shared prior to disagreement

Sections 30.2 and 30.3 both assume that the prosecutor and the judge start from the same prior, μ_0. Often they do not. A prosecutor who has lived with a case for months may be more convinced of guilt than the judge who sees it for the first time; an advocate is typically more optimistic than the audience he addresses. Alonso and Cámara (2016a) allow the Sender and the Receiver to hold different priors, both common knowledge, so that the two agree to disagree.

We keep the judge's prior at the same generic value μ_0 (running value 0.3), so her behavior is as in Section 30.2, and we let the prosecutor's prior be a number $q \in (0, 1)$. When $q = \mu_0$ the two agree and we are back in Section 30.2. The question is what disagreement changes when $q \neq \mu_0$.

30.4.2 The model with two priors

The states, actions, and payoffs are those of Section 30.2. The judge holds the prior μ_0 and convicts if and only if her posterior is at least $\frac{1}{2}$, unchanged. The prosecutor holds a possibly different prior q. He designs a signal; the judge updates from her prior μ_0; but the prosecutor evaluates outcomes using his prior q, that is, using his own assessment of how likely each message, and each resulting conviction, is. Both priors are common knowledge.

[1]The published example reports a conviction rate of 0.43 at $k = 2$. Those numbers are reproduced when the entropy cost is averaged with the message probabilities under the common prior 0.3 rather than under the uniform reference belief; with the cost as defined in Gentzkow and Kamenica's (2014) Assumption 1, the optimum convicts at rate 0.457, shown here. The qualitative comparative statics are identical.

30.4.3 The unifying result: reweight the value function

The prosecutor and the judge read the same evidence but weigh the possible verdicts differently, because they disagree about how likely guilt is.

First, the judge's behavior still depends only on her posterior μ: she convicts if and only if $\mu \geq \frac{1}{2}$, and the prosecutor's gross payoff at that posterior is the same step function $\hat{v}(\mu)$ as in Section 30.2.

Second, the prosecutor weights each message not by the judge's probability but by his own. For a message that leaves the judge at posterior μ, the ratio of the prosecutor's probability of that message to the judge's probability of it is

$$\ell(\mu) = \mu \, \frac{q}{\mu_0} + (1 - \mu) \, \frac{1 - q}{1 - \mu_0},$$

which depends only on the belief the message induces, not on the signal that produced it. The reason is that, once the judge's posterior μ is known, the message's likelihood in each state relative to its average likelihood is fixed, and the two priors simply reweight those likelihoods.

Putting the two facts together, the prosecutor's expected payoff is $\mathbb{E}_\tau[\ell(\mu)\,\hat{v}(\mu)]$, where τ is the judge's Bayes-plausible distribution of posteriors. Define the reweighted value

$$w(\mu) \equiv \ell(\mu)\,\hat{v}(\mu).$$

The prosecutor's problem is

$$\max_{\tau \text{ Bayes plausible}} \mathbb{E}_\tau[w(\mu)] \quad \text{subject to } \mathbb{E}_\tau[\mu] = \mu_0,$$

solved once again by concavifying w and reading it at the judge's prior μ_0.

The nesting. When $q = \mu_0$, the reweighting is $\ell(\mu) = \mu + (1 - \mu) = 1$, so $w = \hat{v}$ and we recover Section 30.2. Heterogeneous priors is the family; the common-prior model of Kamenica and Gentzkow (2011) is its $q = \mu_0$ member.

This reweighting is close kin to the device behind costly persuasion. When the cost in Section 30.3 is measured from a reference belief other than the prior, Gentzkow and Kamenica (2014) translate between that reference belief and the prior in the same way that we here translate between the prosecutor's and the judge's priors, and they credit the device to Alonso and Cámara's work on heterogeneous priors. A single belief-map underlies both extensions, which is why the two share a shape.

30.4.4 The toy model, worked

Step 1. Reweight the payoff. With judge prior μ_0,

$$\ell(\mu) = \mu \, \frac{q}{\mu_0} + (1 - \mu) \, \frac{1 - q}{1 - \mu_0}, \qquad w(\mu) = \begin{cases} 0 & \text{if } \mu < \dfrac{1}{2}, \\ \ell(\mu) & \text{if } \mu \geq \dfrac{1}{2}. \end{cases}$$

Since ℓ is positive and affine in μ, the reweighted payoff w is still a single upward step, now with a sloped top.

Step 2. The optimal signal is unchanged. Because $\hat{v}$ is a step and ℓ is a positive affine function of μ, the concave closure of w touches the graph at the same two beliefs as in Section 30.2: $\mu_l = 0$ and $\mu_h = \frac{1}{2}$. Raising μ_h above $\frac{1}{2}$ lowers the ratio $\frac{\ell(\mu_h)}{\mu_h}$, and lowering μ_l to 0 frees weight for the conviction message, as before. So the prosecutor's optimal investigation is the one derived in Section 30.2, the same two messages and the same conviction rate of $2\mu_0$, for every prosecutor prior q. Intuitively, disagreement about the prior does not change which beliefs are least costly to induce, because the judge's Bayes plausibility and her indifference point are hers alone, and neither depends on q.

Step 3. What the disagreement changes. It changes how much the prosecutor values the unchanged investigation. His expected payoff is

$$\text{value} = 2\mu_0\,\ell(\tfrac{1}{2}) = q + \frac{\mu_0(1-q)}{1-\mu_0},$$

which rises in his own conviction q. At $q = \mu_0$ it equals $2\mu_0$, the Section 30.2 value; at the running $\mu_0 = 0.3$ the value is $\frac{3+4q}{7}$. A prosecutor more sure of guilt than the judge, $q > \mu_0$, values the same investigation more, because he believes convictions are more likely than the judge does.

Step 4. The prosecutor always benefits. With no investigation the judge keeps her prior $\mu_0 < \frac{1}{2}$ and acquits, so the prosecutor's payoff is 0 regardless of his own prior. Since the value in Step 3 is positive for every q, the prosecutor benefits from persuasion at every prior, even when he is nearly certain of guilt. In this courtroom the gain comes easily, because the judge acquits by default; Alonso and Cámara show more generally that, under disagreement, the Sender benefits from persuasion for almost every pair of priors, even in settings where persuasion has no value under a common prior.

30.4.5 A numerical summary

Table 30.4 reports the prosecutor's value as his prior q varies while the judge's prior stays at 0.3. The judge's conviction rate and her welfare do not move; only the prosecutor's valuation does, climbing along the line $\frac{3+4q}{7}$.

q	$\ell(\frac{1}{2})$	Prosecutor value	Judge conviction	Judge welfare
0.1	0.810	0.486	0.60	0.70
0.2	0.905	0.543	0.60	0.70
0.3	1.000	0.600	0.60	0.70
0.5	1.190	0.714	0.60	0.70
0.7	1.381	0.829	0.60	0.70
0.9	1.571	0.943	0.60	0.70

Table 30.4: Prosecutor value as the prosecutor's prior q varies, judge's prior fixed at $\mu_0 = 0.3$. The row $q = \mu_0 = 0.3$ is Section 30.2.

Figure 30.4 depicts the case $q = 0.5$ in the same way as Figures 30.1 and 30.3. The reweighted payoff w is again a single step, but its upper branch now slopes upward, because the prosecutor values a firmer conviction belief more than a marginal one. The concave closure still touches the graph at $\mu_l = 0$ and $\mu_h = \frac{1}{2}$, the very posteriors of Section 30.2, and reading it at the prior gives the prosecutor's value 0.71.

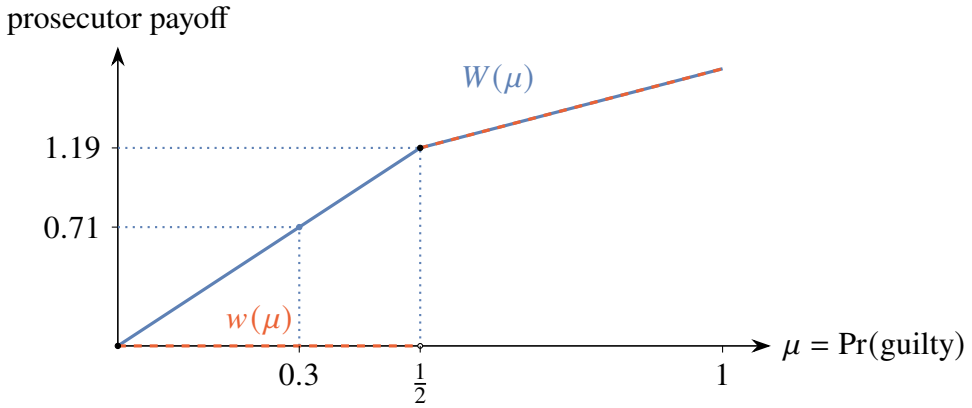


Figure 30.4: The reweighted payoff w (dashed) and its concave closure W (solid) at $q = 0.5$, with the same induced posteriors $\mu_l = 0$ and $\mu_h = \frac{1}{2}$ as in Figure 30.1.

30.5 Competition in Persuasion (Gentzkow and Kamenica, 2017)

30.5.1 Adding a second persuader

So far a single Sender controlled all the evidence. A courtroom has two. The prosecutor wants a conviction; the defense attorney wants an acquittal; each is free to commission investigations, and the judge sees the results of both. The old faith in the adversarial system holds that setting two interested persuaders against each other protects the audience. Gentzkow and Kamenica (2017) make the idea precise, and in our binary courtroom the answer is stark: two fully opposed persuaders reveal the whole truth.

30.5.2 The model with two senders

The states, actions, judge, and prior μ_0 are those of Section 30.2. Now two senders design evidence. The prosecutor earns $\hat{v}_P(\mu) = 1$ when the judge convicts ($\mu \geq \frac{1}{2}$) and 0 otherwise; the defense attorney earns the mirror image, $\hat{v}_D(\mu) = 1$ when the judge acquits ($\mu < \frac{1}{2}$) and 0 otherwise. Each sender chooses a signal; the judge pools everything she sees into a single posterior and acts; a profile of signals is an equilibrium when neither sender can gain by changing his own signal alone. With only the prosecutor present, this is Section 30.2.

30.5.3 The unifying result: both senders sit on their concave closure

The concavification of Section 30.2 told a lone sender where he would place the judge's belief. With two senders the same object characterizes equilibrium, through a short argument. Consider that an equilibrium leaves the judge, with positive probability, at a belief μ where the prosecutor's payoff is strictly below its concave closure, $\hat{v}_P(\mu) < V_P(\mu)$. The prosecutor could then add information at that belief, splitting it into two whose average is μ, and raise his expected payoff, since the concave closure is the most that splitting can reach. He would deviate, so this is not an equilibrium. The same argument applies to the defense attorney. Hence, in equilibrium, every belief the judge might hold lies where *both* senders sit on their own concave closure.

Write

$$M_i = \{\mu : \hat{v}_i(\mu) = V_i(\mu)\}$$

for the beliefs at which sender i cannot improve by splitting, and $M = M_P \cap M_D$. Gentzkow and Kamenica (2017) show that a Bayes-plausible distribution of posteriors is an equilibrium outcome if and only if its support lies in M.

The nesting. With a single sender, $M = M_P$, and the requirement is simply that the judge's beliefs sit on the prosecutor's concave closure, which is the concavification of Section 30.2. Kamenica and Gentzkow (2011) is the one-sender case; competition is what a second, opposed sender adds.

30.5.4 The toy model, worked

Step 1. The prosecutor's coincidence set. The prosecutor's payoff $\hat{v}_P$ is 0 below $\frac{1}{2}$ and 1 above, and its concave closure is $V_P(\mu) = \min(2\mu, 1)$. The two agree wherever $\hat{v}_P$ is already at its ceiling, the interval $[\frac{1}{2}, 1]$, together with the single belief $\mu = 0$ where both equal 0. Hence

$$M_P = \{0\} \cup [\tfrac{1}{2}, 1].$$

Step 2. The defense attorney's coincidence set. His payoff $\hat{v}_D$ is 1 below $\frac{1}{2}$ and 0 above, with concave closure $V_D(\mu) = \min(1, 2(1-\mu))$. The two agree on $[0, \frac{1}{2})$, where $\hat{v}_D$ is already at its ceiling, together with $\mu = 1$ where both equal 0. Hence

$$M_D = [0, \tfrac{1}{2}) \cup \{1\}.$$

Step 3. Intersect. The only beliefs in both sets are the two certainties,

$$M = M_P \cap M_D = \{0, 1\}.$$

At every other belief one of the two senders would add information.

Step 4. The equilibrium. The judge's beliefs must therefore be 0 or 1, so the only Bayes-plausible outcome is full disclosure: the evidence reveals the state, leaving belief 1 with probability μ_0 and belief 0 with probability $1 - \mu_0$. The judge convicts exactly the guilty, at rate μ_0, half the lone prosecutor's $2\mu_0$.

Intuitively, at any partial belief short of certainty, one side always holds evidence that pulls the judge its way: if the prosecutor tries to stop at the indifference belief, the defense attorney can push her toward innocence, and conversely. The only beliefs no one wishes to disturb are the two at which nothing more can be said. Competition unravels all the way to the truth.

30.5.5 What competition adds

Competition reveals more, and the judge gains. A lone prosecutor convicts $2\mu_0$ of defendants; two opposed persuaders convict only the guilty, μ_0, and the judge, now fully informed, reaches her first-best. This is the sharpest form of the general comparative static: when each sender can choose from a rich enough set of signals, competition never reduces the information revealed, and here it forces the maximum.

Full opposition is doing the work. The unraveling is this clean because the senders' interests are fully reversed, so their coincidence sets meet only at certainty. Gentzkow and Kamenica (2017) show, well beyond the binary case, that when each sender can choose from a rich enough set of signals (their Blackwell-connectedness condition), equilibrium outcomes are no less informative than what the senders would choose by colluding, regardless of their preferences; full disclosure is the extreme our courtroom lands on.

30.5.6 A numerical summary

Table 30.5 sets the one-sender and two-sender outcomes side by side at the running prior $\mu_0 = 0.3$. Adding the defense attorney halves the conviction rate and lifts the judge to full information.

Setting	Conviction rate	Prosecutor payoff	Judge welfare	Defense payoff
One sender (Section 30.2)	0.60	0.60	0.70	—
Two opposed senders	0.30	0.30	1.00	0.70

Table 30.5: One persuader versus two opposed persuaders, at $\mu_0 = 0.3$. In general the conviction rate falls from $2\mu_0$ to μ_0 and the judge's welfare rises from $1 - \mu_0$ to 1.

Figure 30.5 depicts the two senders' payoffs. The prosecutor's step rises at $\mu = \frac{1}{2}$; the defense attorney's falls there. Each payoff meets its concave closure only on the half where it is already at its ceiling, plus the far certainty, so the two coincidence sets overlap only at $\mu = 0$ and $\mu = 1$: the beliefs of full disclosure.

30.6 A Guided Tour of the Broader Literature

The four sections above share one move: write the Sender's payoff as a function of the Receiver's belief, and the best the Sender can do is the value of the concave closure of that payoff at the prior.

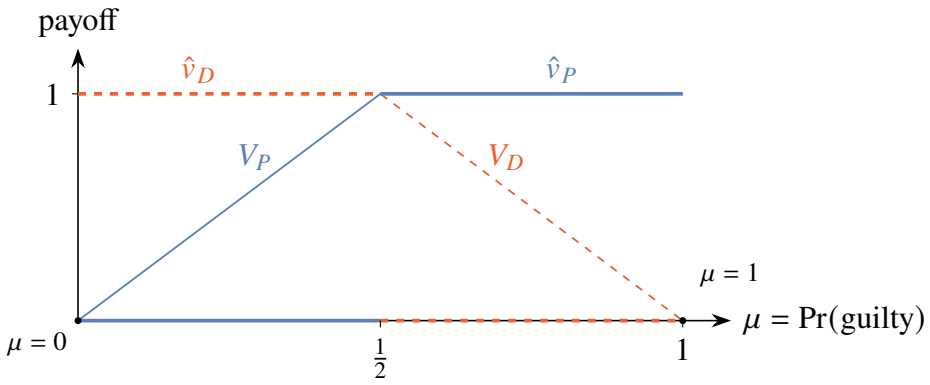


Figure 30.5: The prosecutor's payoff $\hat{v}_P$ (thick solid) and the defense attorney's $\hat{v}_D$ (thick dashed), with their concave closures V_P, V_D (thin). Each payoff meets its own closure on a different set of beliefs, and the two sets overlap only at $\mu = 0$ and $\mu = 1$, the full-disclosure beliefs.

Costs, disagreement, and competition enter only by reshaping that function or by intersecting two of them. The method works cleanly under three conditions: the state space is small enough to plot the value function, the Sender commits fully to the experiment, and there is a single Receiver whose action depends on her posterior. Each active branch of the literature relaxes one of these conditions. This section takes the branches in turn, keeping the prosecutor and judge in view. For each, we say what it changes relative to our baseline, describe concretely how the model would be modified to capture it, and point to representative papers, so that a student can move from our courtroom to the frontier without changing examples. The literature is large; we refer the reader to Kamenica (2019) and Bergemann and Morris (2019) for the complete bibliography.

30.6.1 Many receivers

What changes. There is now a panel rather than a single judge, and the Sender may address the panel with one public signal or with private, receiver-specific signals; the receivers may also care about one another's verdicts. *How the model changes.* Replace the judge with n jurors and a collective rule, say conviction by majority. If the prosecutor is restricted to a single *public* signal, nothing else changes: reinterpret $\bar{a}(\mu)$ as the panel's collective verdict when all jurors share the posterior μ, and concavify $\hat{v}$ as in Section 30.2. If instead he sends *private* signals and jurors care about the collective outcome, a juror's vote depends not only on her own belief but on what she thinks the others learned, and the belief-based value function no longer captures the problem. One then works in the language of Bayes correlated equilibria (Bergemann and Morris 2013, 2016) and information design (Taneva 2019; Mathevet, Perego, and Taneva 2020): characterize the outcomes attainable under *some* information structure, a set carved out by each juror's obedience constraints, and pick the best. Persuading voters is the workhorse case (Alonso and Cámara 2016b; Wang 2015; Bardhi and Guo 2018), and the same tools yield the map of feasible pairs of profit and consumer surplus under third-degree price discrimination in Bergemann et al. (2015).

30.6.2 Many senders: competition in persuasion

What changes. The single persuader becomes several, with possibly opposed preferences, each free to design evidence that the Receiver pools. *How the model changes.* Give sender i a value function $\hat{v}_i$ and its concave closure V_i, as we did for the prosecutor and the defense attorney; an equilibrium is a Bayes-plausible belief distribution supported only where *every* sender sits on his own concave closure, that is, on $\bigcap_i M_i$. Gentzkow and Kamenica (2017) show, well beyond the binary case, that when each sender can choose from a rich enough set of signals, equilibrium outcomes are no less informative than what the senders would choose by colluding; our full-disclosure result is the extreme this comparative static reaches when the two sides are fully opposed. The conclusion is not automatic once senders cannot freely correlate their signals: Board and Lu (2018) show that in a search market where sellers can observe buyers' beliefs or coordinate their disclosures, there is an equilibrium in which competing sellers provide only the monopoly level of information, while if buyers' beliefs are private there is an equilibrium with full information as search costs vanish. Li and Norman (2018) show that adding senders can reduce the information revealed once senders move sequentially, cannot freely correlate their signals, or are restricted to pure strategies.

30.6.3 Limited commitment

What changes. The prosecutor no longer commits fully; with some probability he can quietly alter the message he reports, which is the assumption that had separated persuasion from cheap talk (Crawford and Sobel 1982) and from verifiable disclosure (Grossman 1981; Milgrom 1981). *How the model changes.* Add one parameter, the probability ρ that the prosecutor can secretly replace an unfavorable message with a favorable one. At $\rho = 0$ we are back in Section 30.2; at $\rho = 1$ the prosecutor is a pure cheap talker, and the judge trusts nothing that is not self-serving to reveal. In between, the value is governed by an object that combines the *concavification* of full commitment with the *quasi*-concavification that solves the cheap-talk problem for a state-independent Sender (Lipnowski and Ravid 2020; Lipnowski, Ravid, and Shishkin 2022; Min 2021). A striking lesson is that the judge can be better off when the prosecutor is *less* credible, since a prosecutor who cannot be believed cannot mislead. A related relaxation lets the Sender observe something about the state before he designs the experiment, so that his very choice of investigation leaks information (Perez-Richet 2014; Hedlund 2017), and laboratory tests of how closely subjects follow these predictions are beginning to appear (Fréchette, Lizzeri, and Perego 2022).

30.6.4 A better-informed receiver

What changes. The judge is no longer a blank slate: she holds private information, about her own threshold for conviction or about the state, or she can gather her own evidence after seeing the prosecutor's. *How the model changes.* Give the judge a private type θ, say a privately known conviction threshold drawn from a known distribution, and redefine the prosecutor's value $\hat{v}$ by averaging over θ; the problem stays concavification-friendly, but the optimal investigation acquires structure, becoming a censorship rule that reveals the states below a cutoff and pools those above it (Kolotilin, Mylovanov, Zapechelnyuk, and Li 2017; Kolotilin 2018), or an interval rule

when the private information is about the state (Guo and Shmaya 2019). Rayo and Segal (2010), contemporaneously with Kamenica and Gentzkow (2011), already give the receiver a private outside option. Alternatively, let the judge pay a cost to run her own test after the prosecutor's message: the mere threat disciplines him into disclosing more, and it is without loss to look at equilibria in which she never actually investigates (Matysková and Montes 2023).

30.6.5 Dynamics

What changes. The problem is no longer static: the state evolves over time, or the receiver's past choices shape her future options, so the Sender chooses not only how much to reveal but when. *How the model changes.* Let the defendant's guilt evolve as a Markov process and let the judge act myopically each period on her current belief; the prosecutor now chooses a disclosure policy over time rather than a single experiment. Ely (2017) solves this case and finds that the optimal policy discloses good news with a deterministic delay. When the receiver's sunk actions change her opportunities, dynamic disclosure can sustain more of the action the Sender wants than any one-shot experiment (Ely and Szydlowski 2020). A separate dynamic branch puts beliefs directly into the utility function, modeling entertainment as the management of suspense and surprise (Ely, Frankel, and Kamenica 2015).

30.6.6 Large state spaces and solution methods

What changes. The state is a continuum and the Sender's payoff depends only on the *mean* of the receiver's posterior, so the value function can still be drawn but its concave closure no longer solves the problem, because not every distribution of posterior means with the right average is attainable. *How the model changes.* Let guilt be a continuous severity in $[0, 1]$ rather than a yes-or-no, with the judge's action depending on expected severity. Two reformulations restore tractability: represent each signal by a convex function of the posterior mean and search over that space (Gentzkow and Kamenica 2016), or cast the design as a linear program, equivalently a price-theoretic problem in which the Sender "buys" posterior means using the prior as his budget, with a simple dual test for optimality (Kolotilin 2018; Dworczak and Martini 2019). These methods say when the optimum is a monotone partition (Ivanov 2021; Mensch 2021) and now carry much of the applied work. When even they fail, computation is the fallback: the general problem is hard, but good approximation algorithms exist (Dughmi and Xu 2016; Dughmi 2017).

30.6.7 Applications

What changes. Nothing in the mechanism; only the names of the Sender, the Receiver, and the state. *How the model changes.* Rename the players. A regulator designing a bank stress test reveals just enough about solvency to avert a run (Goldstein and Leitner 2018; Inostroza and Pavan 2025); a school choosing a grading policy reveals just enough about ability to place its graduates (Ostrovsky and Schwarz 2010; Boleslavsky and Cotton 2015); a routing app gives drivers noisy advice to unclog roads (Das, Kamenica, and Mirka 2017); a police force randomizes speed traps

to deter speeding (Lazear 2006); and, in the environmental reading we return to throughout, a firm designs a sustainability label that raises a green-minded consumer's belief to the buying threshold and no further, which is precisely greenwashing (Lyon and Maxwell 2011). In each case the prosecutor-judge diagram is a faithful first picture, and the comparative statics of Sections 30.3–30.5 carry over: costlier testing means coarser disclosure, a more confident designer values the same disclosure more, and a watchdog forces the truth out.

30.6.8 A neighboring market: selling information

What changes. Information is no longer given away to move an action but sold as a good, and the designer, a data broker, a credit bureau, a rating agency, is paid in fees rather than in the Receiver's action (Admati and Pfleiderer 1990; Bergemann and Bonatti 2019). *How the model changes.* Replace the prosecutor's verdict payoff with a price: instead of valuing a conviction, the Sender posts a menu of experiments and prices and the Receiver buys one. The design question shifts from what to reveal to what to sell and at what price, and the difference between selling an information structure up front and selling its realization after the fact maps onto the products that brokers and publishers actually trade. Our courtroom becomes a pricing problem, and the two literatures meet in a shared object.

30.7 Directions for Future Work

The prosecutor and the judge are not only a teaching device; several of the field's open questions can be posed inside them, which is a useful discipline on speculation. We describe four directions where the model of this chapter runs into the research frontier.

Combining contracts with experiments. Incentives and information are usually studied in separate literatures, yet real designers set both at once: a firm chooses a wage schedule and an evaluation system, a court sets both penalties and rules of evidence, a platform sets both prices and what buyers are shown. Inside our model this means giving the prosecutor a second instrument, a transfer to the judge, alongside the investigation, and asking how the optimal experiment changes when the Sender can also move the Receiver's payoffs directly. Li (2017) takes a first step by adding transfers to the Bayesian persuasion model, but a general theory of the joint design of a contract and an experiment is missing, and it is the natural home for the question of when disclosure and payment are substitutes and when they are complements.

A theory, not a catalog, of many receivers. The papers on persuading voters, juries, and bidders each solve a particular case, but there is no organizing statement of when a designer prefers a single public signal to receiver-specific private ones, or of how that preference turns on whether the receivers' actions are strategic complements or substitutes. The machinery to pose the question exists, in the form of Bayes correlated equilibria (Bergemann and Morris 2016); what is missing is the comparative-statics theory built on top of it. In our setting this is the question of how the optimal investigation changes as the lone judge becomes a voting panel, and as the voting rule moves from unanimity to simple majority.

Designing information and options together, over time. A court controls both the evidence a jury sees and the order in which questions are put to it; a firm controls both an employee's feedback and the tasks that remain open to her after she acts. Treating the design of information and the design of the receiver's evolving choice set as one problem, rather than two, is young and promising (Ely and Szydlowski 2020), and the courtroom and green-label narratives used in this chapter are natural laboratories for it: let the prosecutor choose not only what the judge learns but when the verdict must be cast, and the timing of disclosure becomes a second lever.

Solving harder problems, and testing the benchmark. Most models past the smallest examples resist a closed form. The convex-analysis and linear-programming methods of Section 30.6 have widened the solvable class, but many applied designs, especially with several interacting receivers, remain out of analytic reach, and the algorithmic methods developed in computer science are only beginning to be imported into economics (Dughmi 2017). Running alongside is an empirical question the theory cannot answer on its own: do real receivers update as Bayesian persuasion assumes? Laboratory work has just begun to test this (Fréchette, Lizzeri, and Perego 2022), and systematic evidence would tell how far to trust the rational benchmark behind the concavification picture.

30.8 Connections to Other Chapters

- Chapter 24 studies a firm that, after learning its quality, discloses it at a cost or signals it through its price.
- In Chapter 28, the precision of a certification test plays the role of the information structure.
- In Chapter 31, applicants with bad projects obfuscate, which makes the agency's information about an application less precise.

References

Admati, A. R., and P. Pfleiderer (1990). Direct and Indirect Sale of Information. *Econometrica*, 58(4), 901–928.

Alonso, R., and O. Cámara (2016a). Bayesian Persuasion with Heterogeneous Priors. *Journal of Economic Theory*, 165, 672–706.

Alonso, R., and O. Cámara (2016b). Persuading Voters. *American Economic Review*, 106(11), 3590–3605.

Aumann, R. J., and M. B. Maschler (1995). *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press.

Bardhi, A., and Y. Guo (2018). Modes of Persuasion toward Unanimous Consent. *Theoretical Economics*, 13(3), 1111–1149.

Bergemann, D., and A. Bonatti (2019). Markets for Information: An Introduction. *Annual Review of Economics*, 11, 85–107.

Bergemann, D., B. Brooks, and S. Morris (2015). The Limits of Price Discrimination. *American Economic Review*, 105(3), 921–957.

Bergemann, D., and S. Morris (2013). Robust Predictions in Games with Incomplete Information. *Econometrica*, 81(4), 1251–1308.

Bergemann, D., and S. Morris (2016). Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium. *American Economic Review: Papers & Proceedings*, 106(5), 586–591.

Bergemann, D., and S. Morris (2019). Information Design: A Unified Perspective. *Journal of Economic Literature*, 57(1), 44–95.

Board, S., and J. Lu (2018). Competitive Information Disclosure in Search Markets. *Journal of Political Economy*, 126(5), 1965–2010.

Boleslavsky, R., and C. Cotton (2015). Grading Standards and Education Quality. *American Economic Journal: Microeconomics*, 7(2), 248–279.

Crawford, V. P., and J. Sobel (1982). Strategic Information Transmission. *Econometrica*, 50(6), 1431–1451.

Das, S., E. Kamenica, and R. Mirka (2017). Reducing Congestion through Information Design. In *55th Annual Allerton Conference on Communication, Control, and Computing*, 1279–1284.

Dughmi, S. (2017). Algorithmic Information Structure Design: A Survey. *SIGecom Exchanges*, 15(2), 2–24.

Dughmi, S., and H. Xu (2016). Algorithmic Bayesian Persuasion. *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, 412–425.

Dworczak, P., and G. Martini (2019). The Simple Economics of Optimal Persuasion. *Journal of Political Economy*, 127(5), 1993–2048.

Ely, J. C. (2017). Beeps. *American Economic Review*, 107(1), 31–53.

Ely, J. C., A. Frankel, and E. Kamenica (2015). Suspense and Surprise. *Journal of Political Economy*, 123(1), 215–260.

Ely, J. C., and M. Szydlowski (2020). Moving the Goalposts. *Journal of Political Economy*, 128(2), 468–506.

Fréchette, G. R., A. Lizzeri, and J. Perego (2022). Rules and Commitment in Communication: An Experimental Analysis. *Econometrica*, 90(5), 2283–2318.

Gentzkow, M., and E. Kamenica (2014). Costly Persuasion. *American Economic Review: Papers & Proceedings*, 104(5), 457–462.

Gentzkow, M., and E. Kamenica (2016). A Rothschild–Stiglitz Approach to Bayesian Persuasion. *American Economic Review: Papers & Proceedings*, 106(5), 597–601.

Gentzkow, M., and E. Kamenica (2017). Competition in Persuasion. *Review of Economic Studies*, 84(1), 300–322.

Goldstein, I., and Y. Leitner (2018). Stress Tests and Information Disclosure. *Journal of Economic Theory*, 177, 34–69.

Grossman, S. J. (1981). The Informational Role of Warranties and Private Disclosure about Product Quality. *Journal of Law and Economics*, 24(3), 461–483.

Guo, Y., and E. Shmaya (2019). The Interval Structure of Optimal Disclosure. *Econometrica*, 87(2), 653–675.

Hedlund, J. (2017). Bayesian Persuasion by a Privately Informed Sender. *Journal of Economic Theory*, 167, 229–268.

Inostroza, N. A., and A. Pavan (2025). Adversarial Coordination and Public Information Design. *Theoretical Economics*, 20(2), 763–813.

Ivanov, M. (2021). Optimal Monotone Signals in Bayesian Persuasion Mechanisms. *Economic Theory*, 72(3), 955–1000.

Kamenica, E. (2019). Bayesian Persuasion and Information Design. *Annual Review of Economics*, 11, 249–272.

Kamenica, E., and M. Gentzkow (2011). Bayesian Persuasion. *American Economic Review*, 101(6), 2590–2615.

Kolotilin, A. (2018). Optimal Information Disclosure: A Linear Programming Approach. *Theoretical Economics*, 13(2), 607–635.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017). Persuasion of a Privately Informed Receiver. *Econometrica*, 85(6), 1949–1964.

Lazear, E. P. (2006). Speeding, Terrorism, and Teaching to the Test. *Quarterly Journal of Economics*, 121(3), 1029–1061.

Li, C. (2017). A Model of Bayesian Persuasion with Transfers. *Economics Letters*, 161, 93–95.

Li, F., and P. Norman (2018). On Bayesian Persuasion with Multiple Senders. *Economics Letters*, 170, 66–70.

Lipnowski, E., and D. Ravid (2020). Cheap Talk with Transparent Motives. *Econometrica*, 88(4), 1631–1660.

Lipnowski, E., D. Ravid, and D. Shishkin (2022). Persuasion via Weak Institutions. *Journal of Political Economy*, 130(10), 2705–2730.

Lyon, T. P., and J. W. Maxwell (2011). Greenwash: Corporate Environmental Disclosure under Threat of Audit. *Journal of Economics & Management Strategy*, 20(1), 3–41.

Mathevet, L., J. Perego, and I. Taneva (2020). On Information Design in Games. *Journal of Political Economy*, 128(4), 1370–1404.

Matysková, L., and A. Montes (2023). Bayesian Persuasion with Costly Information Acquisition. *Journal of Economic Theory*, 211, 105678.

Mensch, J. (2021). Monotone Persuasion. *Games and Economic Behavior*, 130, 521–542.

Milgrom, P. R. (1981). Good News and Bad News: Representation Theorems and Applications. *Bell Journal of Economics*, 12(2), 380–391.

Min, D. (2021). Bayesian Persuasion under Partial Commitment. *Economic Theory*, 72(3), 743–764.

Ostrovsky, M., and M. Schwarz (2010). Information Disclosure and Unraveling in Matching Markets. *American Economic Journal: Microeconomics*, 2(2), 34–63.

Perez-Richet, E. (2014). Interim Bayesian Persuasion: First Steps. *American Economic Review: Papers & Proceedings*, 104(5), 469–474.

Rayo, L., and I. Segal (2010). Optimal Information Disclosure. *Journal of Political Economy*, 118(5), 949–987.

Taneva, I. (2019). Information Design. *American Economic Journal: Microeconomics*, 11(4), 151–185.

Wang, Y. (2015). Bayesian Persuasion with Multiple Receivers. *Available at SSRN 2625399*.

Chapter 31

Funding a Review Agency under Obfuscation

Companion note to Figueroa, N., and J. Lemus (2024), "The Benefit of the Doubt: Patent Examination under Strategic Obfuscation," Management Science 70(9): 5853–5870. Full reference at the end.

31.1 Introduction

How much should a patent office spend on examining applications? More examiners catch more bad patents, so the answer seems to be "as much as it can afford." Competition authorities and drug regulators face the same choice. Figueroa and Lemus (2024) show that the answer is less simple, because applicants react to the resources of the agency.

An applicant with a bad project is approved only if the agency fails to learn the truth. It can lower that chance by *obfuscation*: adding claims and documents that are costly to check. Extra resources screen out more bad projects, which helps the agency, but they also give bad applicants more to evade, so these applicants obfuscate more. Over a middle range the second force wins, so the agency's payoff from examination first rises with its resources, then falls, and then rises again.

This dip changes the budget question. When good projects are valuable enough relative to the cost of approving a bad one, the optimal budget is either full or modest, and it jumps from one to the other at a single unit cost of resources. A small change in how an agency is funded can therefore cause a large change in how much it examines.

The toy model fixes one set of numbers, so that each expression is a number rather than a formula. The first step solves a benchmark without obfuscation. The second to fourth steps derive the effort of a bad applicant and the approval decision of the agency, and show when more resources lower the agency's payoff. The fifth and sixth steps solve the budget problem, the seventh step draws two policy implications, and the eighth step varies the parameters.

31.2 The Toy Model

Applicant and agency. An applicant has a good project with probability $p = \frac{1}{2}$ and a bad project otherwise, and it privately observes which. The agency maximizes a welfare measure, such as consumer surplus. Approving a good project yields the agency a value $v = 6$, approving a bad project costs it $c = 2$, and rejection yields zero. Approval yields the applicant a gain that we normalize to one. In merger review the applicant is the pair of merging firms, who know whether the merger is good or bad for consumers; in patent examination the applicant knows whether granting a patent on its idea raises or lowers welfare. Two numbers summarize the stakes: the expected value of a good project is $pv = 3$ and the expected cost of a bad project is $(1 - p)c = 1$. An agency that cannot tell the projects apart gains $3 - 1 = 2 > 0$ from approving an application, so it approves when it is uncertain, unless obfuscation changes the pool of uncertain applications.

Obfuscation and examination. The agency invests resources $\rho \in [0, 1]$, which the applicant observes. The applicant then chooses an obfuscation effort $e \geq 0$ at cost e, which the agency does not observe. With probability $1 - \varphi(\rho, e)$ the agency learns the type of the project and approves if and only if it is good. With probability $\varphi(\rho, e)$ it remains uncertain, and it then approves with probability $\lambda \in [0, 1]$, which it chooses. We use the specification of Figueroa and Lemus (2024),

$$\varphi(\rho, e) = 1 - \rho\big(1 - \sqrt{2e}\big), \qquad e \in \big[0, \tfrac{1}{2}\big]. \tag{31.1}$$

More resources raise the probability of learning the type, and more obfuscation lowers it. With no obfuscation ($e = 0$) the agency stays uncertain with probability $1 - \rho$, and with full resources and no obfuscation it learns the type for sure, $\varphi(1, 0) = 0$.

Budget. Resources cost $K(\rho) = k\rho^2$, where $k \geq 0$ is the *unit cost* of resources. The agency's payoff is $U(\rho) - k\rho^2$, where $U(\rho)$ is its equilibrium payoff from examination, derived below.

Timing. (1) The agency chooses ρ, which the applicant observes. (2) The applicant learns its type and chooses e. (3) The agency learns the type with probability $1 - \varphi(\rho, e)$; otherwise it approves with probability λ. We solve for the perfect Bayesian equilibrium by backward induction.

Parameter values. We fix $p = \frac{1}{2}$, $v = 6$, and $c = 2$ in every step below. The values keep the arithmetic round, and they have the two properties that the analysis needs: the payoff from examination has a dip (see the fourth step), and the optimal budget jumps from full to modest (see the sixth step). The eighth step states the conditions on p, v, and c under which these properties hold and shows that other values give the same qualitative results.

31.3 First Step: The Benchmark without Obfuscation

Consider first the case in which applicants cannot obfuscate, so $e = 0$ and the agency stays uncertain with probability $1 - \rho$ about every project. An uncertain agency faces good and bad projects in equal proportions, and approving yields

$$\frac{1}{2} \cdot 6 - \frac{1}{2} \cdot 2 = 2 > 0$$

in expectation, so it approves every uncertain application, $\lambda = 1$. It collects the value 3 of the good projects and bears the cost 1 only when a bad project goes unexamined:

$$U_0(\rho) = 3 - (1 - \rho) = 2 + \rho. \tag{31.2}$$

This payoff is strictly increasing in ρ at the constant rate 1: a unit of resources lowers by one the probability that a bad project goes unexamined, which avoids the expected cost 1 of approving it. The agency maximizes $2 + \rho - k\rho^2$ over $[0, 1]$, so

$$\rho_0^* = \min\left\{\frac{1}{2k},\ 1\right\}. \tag{31.3}$$

Result 1 (benchmark). Without obfuscation the optimal budget is continuous and weakly decreasing in the unit cost: it equals one up to $k = \frac{1}{2}$ and falls gradually afterward.

The budget is one for $k \leq \frac{1}{2}$, $\frac{1}{2}$ at $k = 1$, and $\frac{1}{4}$ at $k = 2$. There is no jump. The second to sixth steps show that obfuscation changes this picture.

31.4 Second Step: The Applicant with a Bad Project

Consider the applicant's choice of effort in stage 2 for given resources ρ and for an anticipated approval probability λ. An applicant with a good project is approved with probability $1 - \varphi(1 - \lambda)$, which is weakly decreasing in φ, so it never spends effort that raises φ at a positive cost. An applicant with a bad project is approved only when the agency remains uncertain and then approves, which has probability $\lambda\varphi$, so it maximizes

$$\lambda\varphi(\rho, e) - e = \lambda\big(1 - \rho + \rho\sqrt{2e}\big) - e. \tag{31.4}$$

This function is concave in e. For $\lambda\rho > 0$ the first-order condition is

$$\frac{\lambda\rho}{\sqrt{2e}} = 1,$$

and it gives

$$e = \frac{(\lambda\rho)^2}{2} \leq \frac{1}{2}. \tag{31.5}$$

Substituting (31.5) into (31.1), the probability that a bad project remains uncertain is

$$\varphi^*(\lambda) = 1 - \rho + \lambda\rho^2. \tag{31.6}$$

Equation (31.6) drives the rest of the analysis. Without obfuscation the bad project stays uncertain with probability $1 - \rho$; obfuscation adds the term $\lambda\rho^2$, which rises with the resources ρ that the agency invests and with the probability λ that an uncertain application is approved. Obfuscation is worthwhile only if approval is possible and only if there is examination to evade: if $\lambda = 0$ or $\rho = 0$, the effort is zero.

31.5 Third Step: The Agency's Approval Decision and Payoff

Only applicants with bad projects obfuscate, so, when $\lambda\rho > 0$, uncertain applications contain a larger share of bad projects than the prior. An uncertain agency puts weight $\frac{1}{2}(1-\rho)$ on a good project (a good applicant stays uncertain with probability $1-\rho$) and weight $\frac{1}{2}\varphi^*(\lambda)$ on a bad project, so approving yields

$$\frac{1}{2}(1-\rho)\cdot 6 - \frac{1}{2}\varphi^*(\lambda)\cdot 2$$

in expectation. It is willing to approve if and only if

$$\Delta(\lambda) \equiv 3(1-\rho) - \left(1-\rho+\lambda\rho^2\right) = 2(1-\rho) - \lambda\rho^2 \geq 0. \tag{31.7}$$

The function Δ is strictly decreasing in λ for $\rho > 0$: the more the agency approves, the more the bad applicants obfuscate, and the worse the uncertain pool becomes. Three cases follow.

1. If $\Delta(1) \geq 0$, then $\Delta(\lambda) > 0$ for every $\lambda < 1$, so the agency strictly prefers to approve unless $\lambda = 1$, and the only equilibrium is $\lambda = 1$. Rearranging, $\Delta(1) \geq 0$ is equivalent to $\rho^2+2\rho-2 \leq 0$, which holds if and only if $\rho \leq \widehat{\rho}$, where

$$\widehat{\rho} \equiv \sqrt{3} - 1 = 0.732 \tag{31.8}$$

   is the positive root of that quadratic.

2. If $\Delta(1) < 0$ and $\rho < 1$, then $\lambda = 1$ is not a best reply, and $\Delta(0) = 2(1-\rho) > 0$, so $\lambda = 0$ is not either. The agency mixes, which requires $\Delta(\lambda) = 0$, and the unique solution is

$$\lambda = \frac{2(1-\rho)}{\rho^2} \in (0,1). \tag{31.9}$$

3. At $\rho = 1$ a good project is never uncertain, $\Delta(\lambda) = -\lambda$, and the agency rejects, $\lambda = 0$: if it approved with positive probability, bad applicants would obfuscate, uncertain applications would come only from bad projects, and approving would be strictly unprofitable.

The agency's payoff from examination is

$$\begin{aligned} U &= 3[\rho + (1-\rho)\lambda] - \lambda\varphi^*(\lambda) \\ &= 3\rho + \lambda\Delta(\lambda). \end{aligned}$$

For $\lambda = 1$ this equals $3 - (1-\rho+\rho^2)$, and when $\lambda < 1$ either $\lambda = 0$ or $\Delta(\lambda) = 0$, so $U = 3\rho$:

$$U(\rho) = \begin{cases} 2+\rho-\rho^2 & \text{if } \rho \leq \widehat{\rho}, \\ 3\rho & \text{if } \rho > \widehat{\rho}. \end{cases} \tag{31.10}$$

Result 2 (equilibrium for given resources; Figueroa and Lemus, 2024, specialized). The equilibrium outcome is unique, and only bad applicants obfuscate, with effort $e = \frac{(\lambda\rho)^2}{2}$. The agency

- approves every uncertain application if $\rho \le \widehat{\rho}$;
- approves uncertain applications with the probability (31.9) if $\widehat{\rho} < \rho < 1$;
- rejects every uncertain application if $\rho = 1$.

Figure 31.1 depicts the approval probability λ and the effort e of a bad applicant against the resources ρ. Above $\widehat{\rho}$, substituting (31.9) into (31.5) gives

$$e = \frac{2(1-\rho)^2}{\rho^2},$$

so the effort peaks at $\widehat{\rho}$, where $e = \frac{\widehat{\rho}^2}{2} = 0.268$.

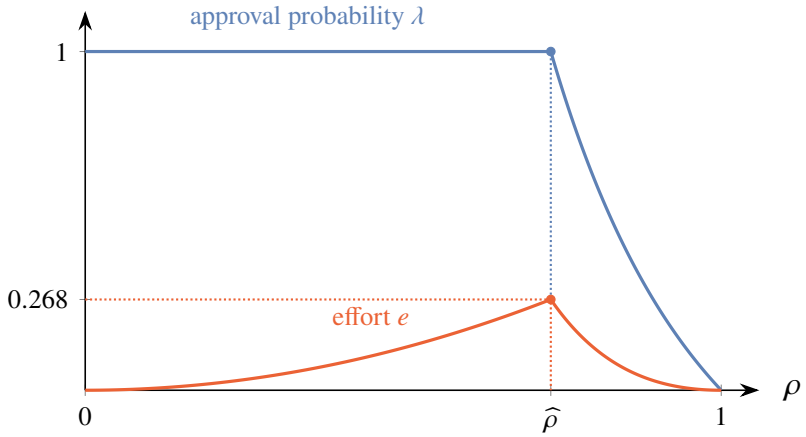


Figure 31.1: Approval probability of an uncertain application and effort of a bad applicant against the resources ρ.

Intuitively, below $\widehat{\rho}$ the agency approves every uncertain application, so more resources leave more examination to evade and raise the effort. Above $\widehat{\rho}$ the agency approves with a lower probability as resources rise, just enough to remain indifferent, which lowers the return to obfuscation, so the effort falls and vanishes at full resources.

Compared with (31.2), $U = U_0 - \rho^2$ whenever $\rho \le \widehat{\rho}$, so ρ^2 is the loss that obfuscation causes. Above $\widehat{\rho}$ the agency is indifferent when it is uncertain, so obfuscation no longer affects its payoff, and $U = 3\rho$ is the value of the good projects that the agency identifies, because uncertain applications yield zero in expectation.

31.6 Fourth Step: When More Resources Lower the Payoff

Differentiating (31.10),

$$U'(\rho) = 1 - 2\rho \ \text{ for } \rho \le \widehat{\rho}, \qquad U'(\rho) = 3 \ \text{ for } \rho > \widehat{\rho}, \tag{31.11}$$

and U is continuous at $\widehat{\rho}$ because $\Delta(1) = 0$ there. The slope below $\widehat{\rho}$ has a simple decomposition. With $\lambda = 1$, the probability that a bad project stays uncertain is $\varphi^* = 1 - \rho + \rho^2$, so

$$\frac{d\varphi^*}{d\rho} = -1 + 2\rho.$$

A unit of resources lowers the probability by one, which gains 1, but it raises obfuscation, so the probability rises by 2ρ, which costs 2ρ. Since every good project is approved anyway when $\rho \le \widehat{\rho}$, extra resources screen out only bad projects, and the net effect $1 - 2\rho$ is negative once ρ exceeds one half.

The dip exists because $\widehat{\rho} = 0.732 > \frac{1}{2}$. Evaluating $\Delta(1)$ at $\rho = \frac{1}{2}$ gives

$$2 \cdot \frac{1}{2} - \frac{1}{4} = \frac{3}{4} > 0,$$

so the agency still approves every uncertain application at one half, and since $\Delta(1)$ is strictly decreasing in ρ, $\widehat{\rho}$ lies above one half. At $\rho = \frac{1}{2}$ the payoff equals

$$U(\tfrac{1}{2}) = 2 + \frac{1}{2} - \frac{1}{4} = 2.25,$$

which can be written as $3\widetilde{\rho}$ with

$$\widetilde{\rho} \equiv \frac{3}{4} \in (\widehat{\rho}, 1). \tag{31.12}$$

Result 3 (non-monotone payoff; Figueroa and Lemus, 2024, specialized). The payoff U is

- strictly increasing on $[0, \frac{1}{2}]$,
- strictly decreasing on $[\frac{1}{2}, \widehat{\rho}]$, and
- strictly increasing on $[\widehat{\rho}, 1]$.

Moreover, $U(\rho) < U(\frac{1}{2})$ for every $\rho \in (\frac{1}{2}, \widetilde{\rho})$, and $U(1) = 3 > U(\frac{1}{2})$.

Figure 31.2 depicts $U(\rho)$, with the benchmark (31.2) as a dashed line. The payoff rises to a first peak $U(\frac{1}{2}) = 2.25$, falls to $U(\widehat{\rho}) = 2.196$ at $\widehat{\rho} = 0.732$, where the agency starts to mix, and then rises at the rate 3. It exceeds its first peak only for $\rho > \widetilde{\rho} = \frac{3}{4}$. The two peaks, the interval between them, and the kink at $\widehat{\rho}$ organize the rest of this chapter.

31.7 Fifth Step: The Optimal Budget

Throughout this step resources cost $k\rho^2$, and only the unit cost $k \ge 0$ and the budget ρ are free. By Result 3, the payoff from examination in Figure 31.2 rises to a first peak $U(\frac{1}{2}) = 2.25$, falls until $\widehat{\rho} = 0.732$, and rises again to $U(1) = 3$. The agency chooses $\rho \in [0, 1]$ to maximize its net payoff

$$\Pi_k(\rho) \equiv U(\rho) - k\rho^2, \tag{31.13}$$

which is continuous, so a maximizer exists. Figure 31.3 depicts Π_k for two unit costs. Each curve has two peaks: a modest one, below one half, and a high one, at full resources.

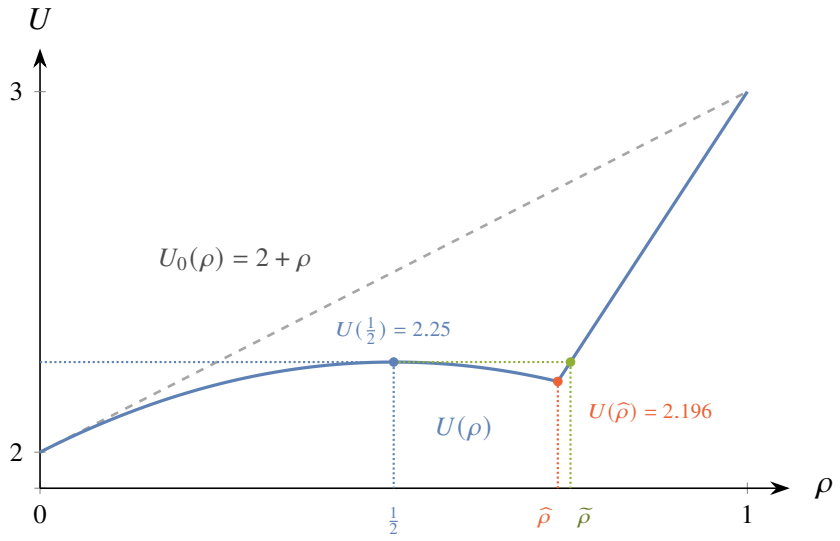


Figure 31.2: Payoff from examination $U(\rho)$ and the benchmark without obfuscation $U_0(\rho) = 2 + \rho$.

- At $k = \frac{1}{2}$ the high peak is higher: $\Pi_k = 2.500$ at $\rho = 1$ against 2.167 at $\rho = \frac{1}{3}$.
- At $k = 1$ the ranking reverses: $\Pi_k = 2.000$ at $\rho = 1$ against 2.125 at $\rho = \frac{1}{4}$.

The following three parts show that this picture holds for every unit cost.

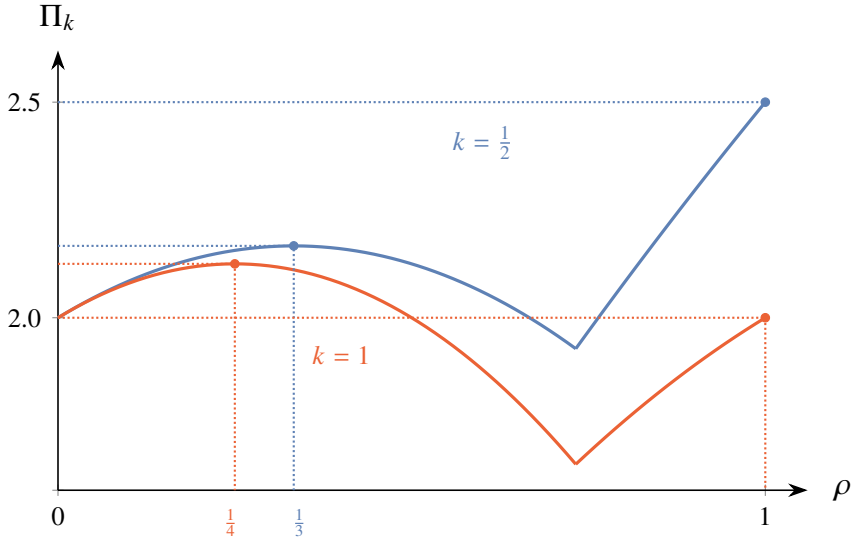


Figure 31.3: Net payoff $\Pi_k(\rho) = U(\rho) - k\rho^2$ for $k = \frac{1}{2}$ and $k = 1$.

(i) Budgets where more resources hurt are never chosen. On $[\frac{1}{2}, \widehat{\rho}]$, U is strictly decreasing and $k\rho^2$ is weakly increasing, so Π_k is strictly decreasing there, and $\Pi_k(\frac{1}{2}) > \Pi_k(\rho)$ for every $\rho \in (\frac{1}{2}, \widehat{\rho}]$. Cutting such a budget to one half raises the payoff from examination and lowers the cost of resources. The optimal budget therefore lies on the *modest* branch $[0, \frac{1}{2}]$ or on the *high*

branch $(\widehat{\rho}, 1]$, whose values are

$$L(k) \equiv \max_{\rho \in [0, \frac{1}{2}]} \Pi_k(\rho), \qquad H(k) \equiv \max_{\rho \in [\widehat{\rho}, 1]} \Pi_k(\rho). \tag{31.14}$$

(ii) The budget falls weakly as the unit cost rises. Let ρ be optimal at k and ρ' be optimal at $k' > k$. Optimality gives

$$U(\rho) - k\rho^2 \geq U(\rho') - k\rho'^2 \qquad \text{and} \qquad U(\rho') - k'\rho'^2 \geq U(\rho) - k'\rho^2.$$

Adding the two inequalities, the terms in U cancel and $(k' - k)(\rho^2 - \rho'^2) \geq 0$, so $\rho' \leq \rho$.

(iii) The budget jumps once. Let $D(k) \equiv H(k) - L(k)$ be the advantage of the high branch, which is continuous in k. At $k = 0$ each branch attains its peak,

$$H(0) = U(1) = 3 \qquad \text{and} \qquad L(0) = U(\tfrac{1}{2}) = 2.25,$$

so $D(0) = 0.75 > 0$. When the unit cost rises from k to k', the value of a branch falls by $k' - k$ times the squared budget that it uses, and the high branch uses at least $\widehat{\rho}^2 = 0.536$ while the modest branch uses at most $\frac{1}{4}$. Formally, take a maximizer ρ_ℓ of Π_k on the modest branch and a maximizer ρ'_u of $\Pi_{k'}$ on the high branch. Then

$$\begin{aligned} H(k') &= \Pi_k(\rho'_u) - (k' - k)\rho'^2_u \leq H(k) - (k' - k)\rho'^2_u, \\ L(k') &\geq \Pi_{k'}(\rho_\ell) = L(k) - (k' - k)\rho_\ell^2, \end{aligned}$$

so

$$D(k') \leq D(k) - (k' - k)\left(\rho'^2_u - \rho_\ell^2\right) < D(k), \tag{31.15}$$

because $\rho'^2_u \geq \widehat{\rho}^2 > \frac{1}{4} \geq \rho_\ell^2$. Hence D is strictly decreasing. Finally,

$$\begin{aligned} H(k) &\leq U(1) - k\widehat{\rho}^2 = 3 - k\widehat{\rho}^2, \\ L(k) &\geq \Pi_k(0) = U(0) = 2, \end{aligned}$$

so $D(k) \leq 1 - k\widehat{\rho}^2 < 0$ for every $k > \frac{1}{\widehat{\rho}^2} = 1.866$. Therefore D has a unique zero $\widehat{k}$ between 0 and 1.866: the high branch is better below $\widehat{k}$ and the modest branch above it. Figure 31.3 places $\widehat{k}$ between $\frac{1}{2}$ and 1, and the sixth step shows that $\widehat{k} = 0.866$.

Result 4 (the optimal budget). Let the cost of resources be $k\rho^2$.

(i) For every $k \geq 0$, no budget in $(\frac{1}{2}, \widehat{\rho}]$ is optimal.

(ii) The optimal budget is weakly decreasing in k.

(iii) There is a unique $\widehat{k} > 0$ such that every optimal budget lies in $(\widehat{\rho}, 1]$ if $k < \widehat{k}$ and in $[0, \frac{1}{2}]$ if $k > \widehat{k}$, and both branches contain an optimal budget at $k = \widehat{k}$.

Intuitively, the payoff has a first peak at one half and a second peak at full resources, and the cost of resources penalizes the second peak more because it is farther out. The value of the high branch therefore falls faster in k than the value of the modest branch, and the two values cross once. At $\widehat{k}$

the budget jumps from above $\widehat{\rho} = 0.732$ to at most one half, a fall of more than 0.232. Compare with Result 1: the benchmark payoff U_0 is linear in ρ, so the net payoff $U_0(\rho) - k\rho^2$ is strictly concave and its maximizer moves continuously with k.

Beyond the example. Parts (i) to (iii) used only two facts:

- the payoff U rises to a first peak a, falls until some $b > a$, and rises again to $U(1) > U(a)$;
- the cost of resources is k times a strictly increasing function $\psi(\rho)$ with $\psi(0) = 0$.

Replacing $\frac{1}{2}$ and $\widehat{\rho}$ by a and b, and ρ^2 by $\psi(\rho)$, every step goes through unchanged. Result 4 therefore holds for any strictly increasing cost, such as the power costs $k\rho^m$, and for any values of p, v, and c for which the payoff has the dip of Result 3; the eighth step studies both.

31.8 Sixth Step: Closed Form with a Quadratic Cost

By Result 4, the optimal budget lies on the modest branch $[0, \frac{1}{2}]$ or on the high branch $(\widehat{\rho}, 1]$. On the modest branch the objective is $2 + \rho - (1 + k)\rho^2$, a strictly concave quadratic with the vertex

$$\rho_1 \equiv \frac{1}{2(1+k)} \leq \frac{1}{2},$$

so

$$L(k) = 2 + \frac{1}{4(1+k)}. \tag{31.16}$$

On the high branch the objective is $3\rho - k\rho^2$, whose derivative $3 - 2k\rho$ is positive for $\rho < 1$ whenever $k \leq \frac{3}{2}$, so the maximizer is $\rho = 1$ and

$$H(k) = 3 - k \qquad (k \leq \tfrac{3}{2}). \tag{31.17}$$

The difference is

$$D(k) = 1 - k - \frac{1}{4(1+k)} = \frac{3 - 4k^2}{4(1+k)}, \tag{31.18}$$

which has the unique positive zero

$$\overline{k} \equiv \frac{\sqrt{3}}{2} = 0.866. \tag{31.19}$$

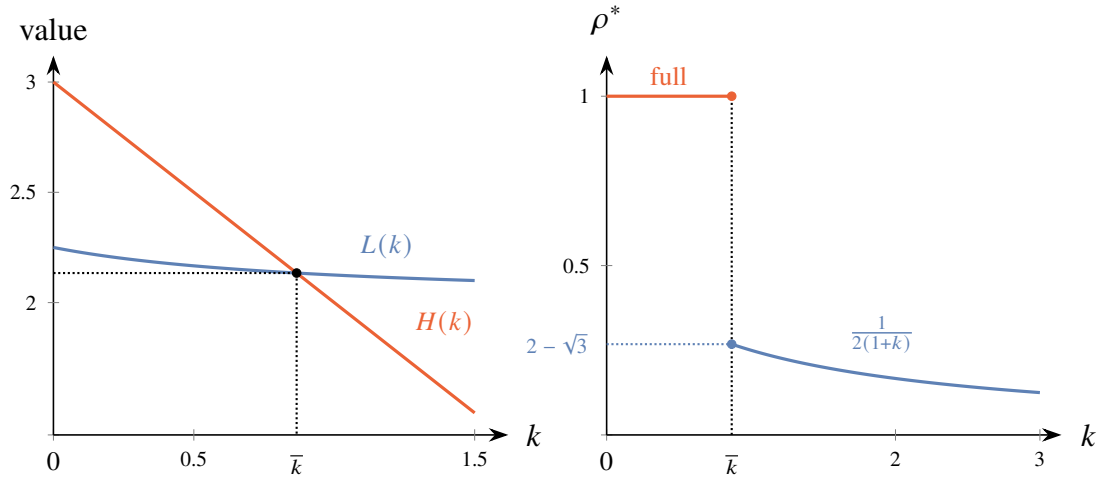
The zero is a valid jump point because $\overline{k} \leq \frac{3}{2}$.

Result 5 (closed form). The jump point is $\widehat{k} = \overline{k}$, and

- if $k < \overline{k}$, the optimal budget is $\rho^* = 1$;
- if $k > \overline{k}$, the optimal budget is $\rho^* = \frac{1}{2(1+k)} < 2 - \sqrt{3}$;
- if $k = \overline{k}$, both $\rho = 1$ and $\rho = 2 - \sqrt{3}$ are optimal.

The set of optimal budgets over all $k \geq 0$ is therefore $(0, 2 - \sqrt{3}] \cup \{1\}$.

Here $2 - \sqrt{3} = 0.268$. Figure 31.4 depicts the two values in panel (a) and the optimal budget in panel (b). In panel (a), the value $H(k) = 3 - k$ of full resources falls faster than the value $L(k)$ of the modest budget, and the lines cross at $\overline{k}$. In panel (b), the budget equals one up to $\overline{k}$, falls discontinuously to 0.268, and then falls gradually, to 0.222 at $k = 1.25$ and to 0.167 at $k = 2$.



(a) Values of the two branches (b) Optimal budget

Figure 31.4: Values of the two branches and optimal budget against the unit cost k.

Why the value of good projects does not matter for the cutoff. The cutoff (31.19) comes from (31.18), which contains the cost 1 of a bad project and no trace of the value 3 of good projects. At full resources the agency learns every type and approves every good project, and at the modest budget it approves every uncertain application and therefore every good project too. Both branches collect the same 3, so it cancels from $D(k)$; the choice between them turns on the cost of the bad projects that the modest budget lets through. The eighth step shows how the cutoff changes with that cost.

The equilibrium behind the budget. Table 31.1 reports the equilibrium outcome for k from 0.25 to 3 in steps of 0.25. The columns show the optimal budget ρ^*; the effort e^* of a bad applicant; the probability $1 - \varphi(\rho^*, e^*)$ that the agency learns the type of that applicant; the approval probability λ^* of an uncertain application; the expected value V^* obtained from good projects; the expected cost C^* of the bad projects that it approves; the payoff $U = V^* - C^*$; and the payoff $U - k(\rho^*)^2$ net of the cost of resources, where

$$V^* \equiv 3[\rho^* + (1 - \rho^*)\lambda^*] \qquad \text{and} \qquad C^* \equiv \lambda^* \varphi(\rho^*, e^*).$$

Table 31.1 splits into two regimes:

- For $k \leq 0.75$, below $\overline{k} = 0.866$, the agency uses full resources, learns every type, and obtains $U = 3$ without any obfuscation.
- From $k = 1$ on, the budget is modest and falls with k, from 0.250 to 0.125. Bad applicants now obfuscate, but their effort falls with the budget, and the agency learns their type with a lower

k	ρ^*	e^*	$1-\varphi(\rho^*, e^*)$	λ^*	V^*	C^*	U	$U - k(\rho^*)^2$
0.25	1.000	0.000	1.000	0.000	3.000	0.000	3.000	2.750
0.50	1.000	0.000	1.000	0.000	3.000	0.000	3.000	2.500
0.75	1.000	0.000	1.000	0.000	3.000	0.000	3.000	2.250
1.00	0.250	0.031	0.188	1.000	3.000	0.813	2.188	2.125
1.25	0.222	0.025	0.173	1.000	3.000	0.827	2.173	2.111
1.50	0.200	0.020	0.160	1.000	3.000	0.840	2.160	2.100
1.75	0.182	0.017	0.149	1.000	3.000	0.851	2.149	2.091
2.00	0.167	0.014	0.139	1.000	3.000	0.861	2.139	2.083
2.25	0.154	0.012	0.130	1.000	3.000	0.870	2.130	2.077
2.50	0.143	0.010	0.122	1.000	3.000	0.878	2.122	2.071
2.75	0.133	0.009	0.116	1.000	3.000	0.884	2.116	2.067
3.00	0.125	0.008	0.109	1.000	3.000	0.891	2.109	2.062

Table 31.1: Equilibrium outcome as the unit cost k rises.

probability, from 0.188 to 0.109. The agency approves every uncertain application, so $V^* = 3$ in every row, and the cost C^* of the bad projects that it approves rises from 0.813 to 0.891, which lowers U from 2.188 to 2.109.

The jump lowers the payoff from examination from 3 to 2.188, because the agency now approves bad projects that it would have rejected, but the saving in the cost of resources makes the net payoff fall only from 2.250 (at $k = 0.75$) to 2.125 (at $k = 1$), and the net payoff falls continuously in k afterward.

31.9 Seventh Step: Two Policy Implications

A budget imposed from outside can lower the agency's payoff. By Result 4(i), the agency never chooses a budget in $(\frac{1}{2}, \widehat{\rho}]$. An outside budget can, and it need not help. At $k = 1$ the agency chooses $\rho^* = 0.25$ and obtains the net payoff 2.125. A floor of $\rho \geq \frac{1}{2}$, such as a hiring target, leaves the agency with at most 2.000: on $[\frac{1}{2}, \widehat{\rho}]$ the net payoff falls from 2.000, and on $[\widehat{\rho}, 1]$ it rises back to 2.000 at $\rho = 1$. In contrast, a lower unit cost leads the agency to choose a high budget on its own: at $k = \frac{1}{2}$ it chooses $\rho^* = 1$ without any floor. A fee that is independent of the resources the agency spends is a lump sum and leaves k, and hence the budget, unchanged.

A small change in how the agency is funded can move its budget from full to modest. At $k = 0.85$ the agency chooses $\rho^* = 1$ and obtains the net payoff 2.150. At $k = 0.90$, an increase of 6 percent in the unit cost, it chooses $\rho^* = 0.263$, which is 74 percent lower, and obtains 2.132. Differences in budgets across agencies or years therefore need not reflect different valuations of examination, and a small difference in the cost at which agencies obtain resources can produce a large difference in how much they examine.

31.10 Eighth Step: Other Parameter Values, and What the Toy Model Keeps and Drops

Other values of p, v, and c. Write $G \equiv pv$ for the expected value of a good project and $B \equiv (1-p)c$ for the expected cost of a bad project, so that the numbers of Section 31.2 are $G = 3$ and $B = 1$. Repeating the first to sixth steps with general $G > B$ gives $U_0(\rho) = G - B(1-\rho)$,

$$\widehat{\rho} = \frac{-(G-B) + \sqrt{(G-B)^2 + 4B(G-B)}}{2B},$$

and $U(\rho) = G - B(1 - \rho + \rho^2)$ for $\rho \le \widehat{\rho}$ and $G\rho$ above. Two conditions reproduce the results:

- the payoff has a dip if and only if $G > \frac{3B}{2}$: $\Delta(1)$ at $\rho = \frac{1}{2}$ equals

$$\frac{G}{2} - \frac{3B}{4}, \qquad \text{and} \qquad \widetilde{\rho} = 1 - \frac{3B}{4G}.$$

  If $G \le \frac{3B}{2}$, U is strictly increasing on $[0, 1]$ and Result 3 does not hold;

- the budget jumps from full to modest if $G \ge \sqrt{3}\, B$, with

$$L(k) = G - B + \frac{B^2}{4(B+k)}, \qquad H(k) = G - k, \qquad \overline{k} = \frac{\sqrt{3}}{2} B, \tag{31.20}$$

  and $\rho^* = \frac{B}{2(B+k)}$ above $\overline{k}$.

The cutoff $\overline{k}$ is proportional to the expected cost B of a bad project and does not depend on G. Any values that satisfy the two conditions, such as $G = 2$ and $B = 1$, or the values $G = 1.4$ and $B = 0.3$ of Figueroa and Lemus (2024), give the same qualitative results.

The condition $G \ge \sqrt{3}\, B$ cannot be relaxed. Evaluating $D(k) = G - k - L(k)$ at $k = \frac{G}{2}$ gives

$$D(\tfrac{G}{2}) = \frac{3B^2 - G^2}{2(G + 2B)},$$

which is positive if $G < \sqrt{3}\, B$. Then $\widehat{k} > \frac{G}{2}$, and for $k \in (\frac{G}{2}, \widehat{k})$ the high branch is still better but its maximizer is interior, $\rho = \frac{G}{2k} < 1$, so the agency uses neither full resources nor a modest budget. For $G = 1.6$ and $B = 1$, $\widehat{k} = 0.873$ exceeds $\frac{G}{2} = 0.8$, and the optimal budget just below $\widehat{k}$ is 0.917.

A common unit cost can lead to full examination of high-stakes cases and modest examination of low-stakes cases. Because $\overline{k} = \frac{\sqrt{3}}{2} B$ rises with the expected cost of a bad project and does not depend on G, applications that differ in B have different cutoffs. Consider two kinds of applications, both with $G = 3$, and a common unit cost $k = 1$:

- low-stakes applications, with $B = \frac{1}{2}$ (for example $c = 1$), have the cutoff 0.433, so at $k = 1$ the agency examines them with the modest budget

$$\frac{B}{2(B+k)} = \frac{1}{6} = 0.167;$$

- high-stakes applications, with $B = \frac{3}{2}$ (for example $c = 3$), have the cutoff 1.299, so at $k = 1$ the agency examines them with full resources.

A rise in the unit cost therefore first cuts examination where bad approvals are less costly.

Four further variations show how much of the jump survives.

Obfuscation as waste. Let the agency's objective also subtract the effort of bad applicants, at the expected cost $(1 - p)e$. This lowers the payoff by $w\rho^2$, with $w \equiv \frac{1-p}{2}$, when $\rho \leq \widehat{\rho}$, by less between $\widehat{\rho}$ and one, and by nothing at full resources, where nobody obfuscates. The payoff keeps the two-peak shape used in the fifth step, with the first peak at $\frac{B}{2(B+w)}$, and the cutoff and the modest budget become

$$\overline{k}_w = \frac{\sqrt{(w+B)(w+3B)} - w}{2}, \qquad \rho^* = \frac{B}{2(B+w+k)} \quad (k > \overline{k}_w), \tag{31.21}$$

provided $\overline{k}_w \leq \frac{G}{2}$. The cutoff satisfies $\overline{k} < \overline{k}_w < B$, so waste raises it by less than the factor $\frac{2}{\sqrt{3}} = 1.155$. With $p = \frac{1}{2}$, $w = \frac{1}{4}$: the cutoff rises from 0.866 to 0.883, the budget just above it is 0.234 instead of 0.268, and at $k = 1.25$ the budget falls from 0.222 to 0.200. The waste falls only on the modest budget, which induces obfuscation, whereas full resources induce none.

Other cost functions. Result 4 holds for every strictly increasing cost (see the end of the fifth step), but the size of the jump depends on the cost. Table 31.2 reports, for $K(\rho) = k\rho^m$, the unit cost $\widehat{k}$ at which the optimal budget moves from the high branch to the modest branch, and the budgets just below and just above $\widehat{k}$:

- with a linear cost ($m = 1$), the agency jumps from full resources to no resources at $\widehat{k} = B = 1$: the modest branch is $\rho = 0$ when $k \geq B$, so $L = G - B$, $H = G - k$, and $D = B - k$;
- with $m = 2$ and $m = 3$, the agency jumps from full resources to a modest budget;
- with steeper costs ($m \geq 4$), the budget just below $\widehat{k}$ is already below one, because the last units of resources are very costly, but the budget still falls discontinuously at $\widehat{k}$.

The budget just above the jump rises with the convexity of the cost.

m	1	2	3	4	5	6
Jump point $\widehat{k}$	1.000	0.866	0.807	0.781	0.800	0.865
Budget just below $\widehat{k}$	1	1	1	0.987	0.931	0.896
Budget just above $\widehat{k}$	0	0.268	0.351	0.400	0.431	0.451

Table 31.2: Jump point and budgets on either side of it for the cost $K(\rho) = k\rho^m$, with $G = 3$ and $B = 1$.

A dip is sufficient for the jump but not necessary. Let the applicant's gain from approval be $g \in (\frac{1}{2}, 1]$ instead of one, so that the effort becomes $e = \frac{(g\lambda\rho)^2}{2}$. The payoff has a dip if and only if

$$G > \frac{B(4g-1)}{2(2g-1)}.$$

With $G = 3$ and $B = 1$:

- for $g = 0.75$ the threshold is $2 < 3$, so the payoff has a dip; the cutoff is 0.843 and the budget after the jump is 0.314;
- for $g = 0.55$ the threshold is $6 > 3$, so the payoff has no dip. Yet the optimal budget still jumps, from 1 to 0.366 at $k \approx 0.817$, because the payoff has a convex kink at $\widehat{\rho}$, where the agency switches from approving every uncertain application to mixing.

The jump therefore does not require the dip: the convex kink at which the agency's behavior changes is enough.

Capacity. If the feasible resources are capped at $\overline{\rho} < 1$, so that examination is never perfect, the cutoff depends on G: for $B = 1$ and $\overline{\rho} = 0.95$ it is 0.787 at $G = 3$ and 0.603 at $G = 6$, and the jump disappears for $G \geq \frac{3B}{4(1-\overline{\rho})} = 15$. Independence of the cutoff from G is therefore a property of full capacity.

What the toy model drops. Figueroa and Lemus (2024) derive the equilibrium for a general probability φ and write down the budget problem for a convex cost of resources, but they do not characterize its solution. The toy model uses the single specification (31.1) and a quadratic (or power) cost, which yield closed forms. It has one applicant, a one-shot decision, and resources that the applicant observes before choosing its effort. The agency pays for resources through a cost that does not depend on who applies, and no fee changes the pool of applicants. Each of these features can be relaxed one at a time. A natural extension is a self-funded office, in which a fee sets who applies and the unit cost of resources at the same time.

31.11 Connections to Other Chapters

- Chapter 30 studies senders who design the information that a receiver sees. Obfuscation in this chapter is a costly way to make the information about a project less precise.
- Chapter 11 studies merging firms that spend resources to produce verifiable evidence of a merger's efficiencies.

References

Figueroa, N., and J. Lemus (2024). "The Benefit of the Doubt: Patent Examination under Strategic Obfuscation." *Management Science*, 70(9), pp. 5853–5870.

Part VIII

Advertising

The three chapters in this part study advertising, and each focuses on a different question: whom an ad reaches, what it says, and who pays for it. Chapter 32 asks whom firms should reach. Two firms compete in prices and in informative advertising, and targeting lets each firm drop the ads that reach its rival's loyal consumers and advertise to comparison shoppers only part of the time, which softens price competition and raises profits. Chapter 33 asks what an ad should say. Hype, which praises features that every consumer values, shifts demand outward, while real information, which lets each consumer learn how well the product fits his tastes, rotates demand. Profits are U-shaped in the spread of valuations, so a firm chooses either a mass-market or a niche posture. Chapter 34 turns to the media that carry the ads. Free television and radio sell their viewers' attention to advertisers, and viewers dislike commercials.

In Chapters 32 and 33, advertising informs consumers and the firm that advertises pays for it. In Chapter 34, advertising is a nuisance that viewers accept in exchange for free programs, and whether the market delivers too many or too few ads depends on how much viewers dislike them.

Chapter 32

The Targeting of Advertising: Loyal Consumers and Comparison Shoppers

Companion note to Iyer, G., D. Soberman, and J. M. Villas-Boas (2005), "The Targeting of Advertising," Marketing Science 24(3): 461–476.

32.1 Introduction

Consider two pizza shops in the same town. Some customers always order from the same shop, while others compare prices and order from the cheaper one. No customer learns about a shop's offer unless one of its ads reaches her. With online platforms and loyalty-card data, each shop can now choose *who* sees its ads, not only how many ads to buy. How does this ability change advertising, prices, and profits? Iyer et al. (2005) answer this question.

Without targeting, an ad reaches the whole market, including the rival's loyal customers, so part of every ad is wasted. With targeting, each firm always advertises to its own loyal customers but advertises to the comparison shoppers only part of the time. When a firm skips the comparison shoppers, they see only the rival's ad, and the rival sells to them as a monopolist. Advertising less to comparison shoppers thus creates differentiation that the products themselves lack, which softens price competition.

Targeting therefore raises profits for two reasons: it softens price competition, and it removes the ads that are wasted on the rival's loyal customers. Targeting raises advertising spending when advertising is costly and lowers it when advertising is inexpensive. The ability to target prices, in contrast, does not change profits: loyal customers pay more and comparison shoppers pay less. The value of targeting comes from choosing who sees the ads, not from choosing who pays which price.

The toy model fixes the size of each segment (a quarter of consumers loyal to each firm and half comparison shoppers), so only the cost of advertising varies. The first step explains why prices are random, as in Varian (1980). The second and third steps solve uniform and targeted advertising,

and the fourth step compares advertising spending and prices. The fifth step adds targeted pricing, the sixth step lets firms invest in targeting, and the seventh step lists what the toy model keeps and drops.

32.2 The Toy Model

Consumers. Two firms, 1 and 2, sell a product at zero marginal cost to a unit mass of consumers. Each consumer buys at most one unit and values it at 1. The market has three segments:

- a mass $h = \frac{1}{4}$ of consumers loyal to firm 1, who buy only from firm 1;
- a mass $h = \frac{1}{4}$ of consumers loyal to firm 2, who buy only from firm 2;
- a mass $s = 1 - 2h = \frac{1}{2}$ of comparison shoppers, who buy from the firm with the lower price.

A larger h means that the products are more differentiated; with $h = 0$, the model collapses to Bertrand competition.

Advertising. Consumers do not know that a product exists unless an ad of that firm reaches them. A consumer buys only from the firms whose ads she received, and only if the price is at most 1. Advertising to the whole market costs A. When a firm can target, advertising to a segment costs A times the size of the segment: $\frac{A}{4}$ for a loyal segment and $\frac{A}{2}$ for the comparison shoppers. We let $A \in \left(0, \frac{3}{4}\right)$; the upper bound, $A < (1 - h) \cdot 1$, guarantees that firms advertise in every case below.

Timing. Firms simultaneously choose whether to advertise to each segment they can reach and their prices, without observing the rival's advertising. We write ϕ for the probability that a firm advertises to a segment, which we can interpret as its advertising frequency within a period.

32.3 First Step: Why Prices Are Random

Consider that both firms advertise to the comparison shoppers and charge the same price $p > 0$. Each firm would then undercut slightly and win all the comparison shoppers. Undercutting stops only at a price so low that charging 1 to the loyal consumers alone is better, and from there the rival wants to raise its price again. Hence, as in Varian (1980), no equilibrium in pure price strategies exists, and each firm draws its price from a distribution. Every symmetric equilibrium in the second to fifth steps has the same form: the price p of an advertising firm has c.d.f.

$$F(p) = 1 - c\,\frac{1-p}{p} \qquad \text{for } p \in \left[\frac{c}{1+c},\, 1\right], \tag{32.1}$$

where $c > 0$ is a constant that depends on the case. The lower bound $\frac{c}{1+c}$ solves $F(p) = 0$, and the expected price is

$$E[p] = \int_{\frac{c}{1+c}}^{1} p\,\frac{c}{p^2}\,dp = c\ln\left(\frac{1+c}{c}\right), \tag{32.2}$$

which increases in c. A larger c thus shifts the whole price distribution up.

32.4 Second Step: Uniform Advertising

Without targeting, an ad reaches the whole market, including the rival's loyal consumers, who never buy. Consider that each firm advertises with probability ϕ and, when it does, draws its price from F. An advertising firm that charges p earns

$$\pi(p) = \frac{1}{4}\,p + \frac{1}{2}\,p\left[(1-\phi) + \phi\,(1-F(p))\right] - A, \tag{32.3}$$

since it sells to its loyal consumers and to the comparison shoppers whenever the rival does not advertise or charges more. At $p = 1$ the firm sells to the comparison shoppers only when the rival does not advertise, so

$$\pi(1) = \frac{1}{4} + \frac{1}{2}(1-\phi) - A.$$

- If $A \leq \frac{1}{4}$, advertising and charging 1 to the loyal consumers alone is profitable, so firms always advertise, $\phi = 1$, and earn $\frac{1}{4} - A$.
- If $A > \frac{1}{4}$, firms must be indifferent between advertising and earning zero, so $\pi(1) = 0$, which gives $\phi^U = \frac{3}{2} - 2A$.

Setting $\pi(p)$ equal to the equilibrium profit for every p in the support gives (32.1) with $c = \frac{1}{2}$ when $A \leq \frac{1}{4}$ and

$$c^U = \frac{4A}{3-4A} \qquad \text{when } A > \frac{1}{4},$$

whose lowest price is $\frac{4A}{3}$.

Result 1 (uniform advertising, Lemma 1 and Proposition 1 of the paper). With uniform advertising:

- if $A \leq \frac{1}{4}$, firms always advertise and each earns $\frac{1}{4} - A$;
- if $\frac{1}{4} < A < \frac{3}{4}$, firms advertise with probability $\phi^U = \frac{3}{2} - 2A$, which decreases in A, and earn zero profit;
- if $A > \frac{1}{4}$, the price distribution has $c^U = \frac{4A}{3-4A}$ in (32.1), which increases in A, so prices rise with A.

Intuitively, a firm can always secure the profit of a monopolist over its loyal consumers, $\frac{1}{4} - A$, by charging 1, and competition for the comparison shoppers drives profit down to that level. When $A > \frac{1}{4}$ this level is negative, so firms advertise only part of the time. A quarter of every uniform ad is wasted, since it reaches the consumers loyal to the rival.

32.5 Third Step: Targeted Advertising

With targeting, a firm never advertises to the rival's loyal consumers. Advertising only to its own loyal consumers and charging 1 yields

$$\frac{1}{4} - \frac{A}{4} = \frac{1-A}{4} > 0,$$

so each firm always advertises to them. Consider that each firm also advertises to the comparison shoppers with probability ϕ. A firm that advertises to both segments and charges p earns

$$\pi(p) = \frac{1}{4}\,p + \frac{1}{2}\,p\left[(1-\phi) + \phi\,(1-F(p))\right] - \frac{3A}{4}. \tag{32.4}$$

Indifference with advertising to the loyal segment only, $\pi(1) = \frac{1-A}{4}$, gives

$$\frac{1}{2}(1-\phi) = \frac{A}{2},$$

that is, $\phi^T = 1 - A$. Setting $\pi(p) = \frac{1-A}{4}$ on the support gives (32.1) with

$$c^T = \frac{1+2A}{2(1-A)}, \tag{32.5}$$

whose lowest price is $\frac{1+2A}{3}$.

Result 2 (targeted advertising, Proposition 2 of the paper). With targeted advertising:

- each firm advertises to its loyal consumers with probability one and to the comparison shoppers with probability $\phi^T = 1 - A$;
- each firm earns $\Pi^T = \frac{1-A}{4}$, above its profit with uniform advertising for every A;
- the value of targeting, Π^T minus the profit with uniform advertising, is $\frac{3A}{4}$ if $A \leq \frac{1}{4}$ and $\frac{1-A}{4}$ if $A > \frac{1}{4}$, largest at $A = \frac{1}{4}$.

Figure 32.1 depicts the advertising probabilities. Under targeting, firms advertise more to their loyal consumers (probability one) than to comparison shoppers ($1 - A$). The two curves cross at $A = \frac{1}{2}$: below it, uniform advertisers reach comparison shoppers with a higher probability than targeting firms do; above it, with a lower probability. Figure 32.2 depicts profits: targeted advertising earns more than uniform advertising at every cost, and the gap peaks at $A = \frac{1}{4}$, where it equals $\frac{3}{16}$.

Intuitively, targeting raises profits through two channels. First, it eliminates the ads that fall on the rival's loyal consumers, a direct cost saving. Second, by advertising to comparison shoppers only part of the time, each firm leaves them to the rival with probability A, who then sells to them as a monopolist. Advertising less to comparison shoppers thus creates differentiation that the products themselves lack, which softens price competition and lets each firm keep the profit $\frac{1-A}{4}$ it earns from its loyal consumers.

32.6 Fourth Step: Advertising Spending and Prices

Spending. With uniform advertising, a firm spends $A\,\phi^U$, that is, A if $A \leq \frac{1}{4}$ and $\frac{A(3-4A)}{2}$ otherwise. With targeting, it spends

$$\frac{A}{4} + \frac{A}{2}\,\phi^T = \frac{A(3-2A)}{4}.$$

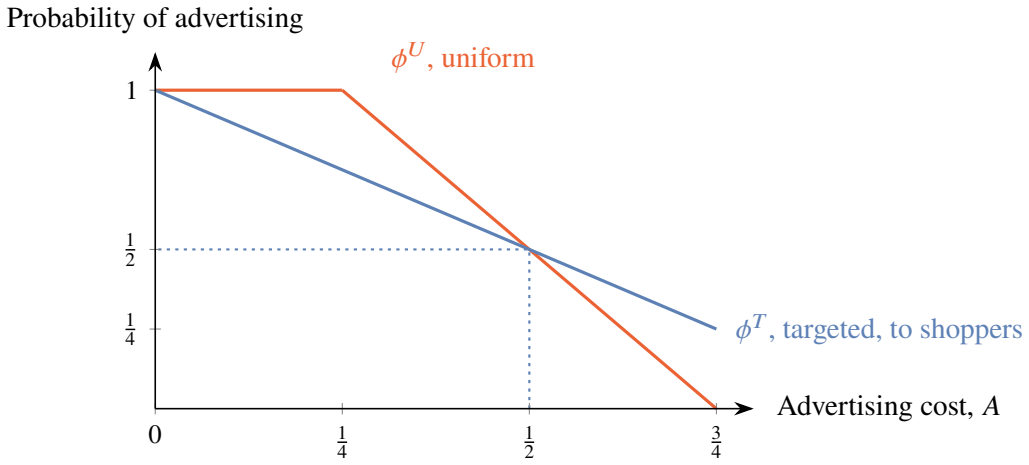


Figure 32.1: Probability of advertising with uniform advertising (red) and of advertising to comparison shoppers with targeted advertising (blue).

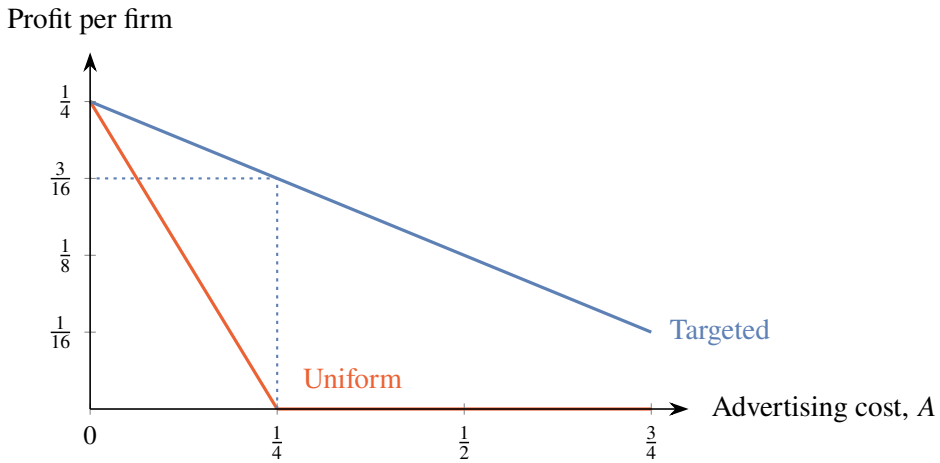


Figure 32.2: Profit per firm with uniform advertising (red) and with targeted advertising (blue).

For $A > \frac{1}{4}$ the difference is

$$\frac{A(3-4A)}{2} - \frac{A(3-2A)}{4} = \frac{3A(1-2A)}{4}. \tag{32.6}$$

Result 3 (spending and prices, Proposition 3 of the paper). Compared with uniform advertising, targeted advertising:

- lowers advertising spending if $A < \frac{1}{2}$ and raises it if $A > \frac{1}{2}$;
- shifts each firm's price distribution up if $A < \frac{1}{2}$ and down if $A > \frac{1}{2}$, since $c^T > c^U$ if and only if $A < \frac{1}{2}$ (for $A \le \frac{1}{4}$, the distribution with uniform advertising has $c = \frac{1}{2} < c^T$).

Figure 32.3 depicts advertising spending. Spending with uniform advertising peaks at $A = \frac{3}{8}$, where it equals $\frac{9}{32}$, and falls to zero as A approaches $\frac{3}{4}$, since the probability of advertising then approaches zero. Spending with targeting keeps rising and crosses the uniform curve at $A = \frac{1}{2}$. At $A = \frac{3}{8}$, for instance, targeting cuts spending from $\frac{9}{32}$ to $\frac{27}{128}$; at $A = \frac{5}{8}$, it raises spending from $\frac{5}{32}$ to $\frac{35}{128}$.

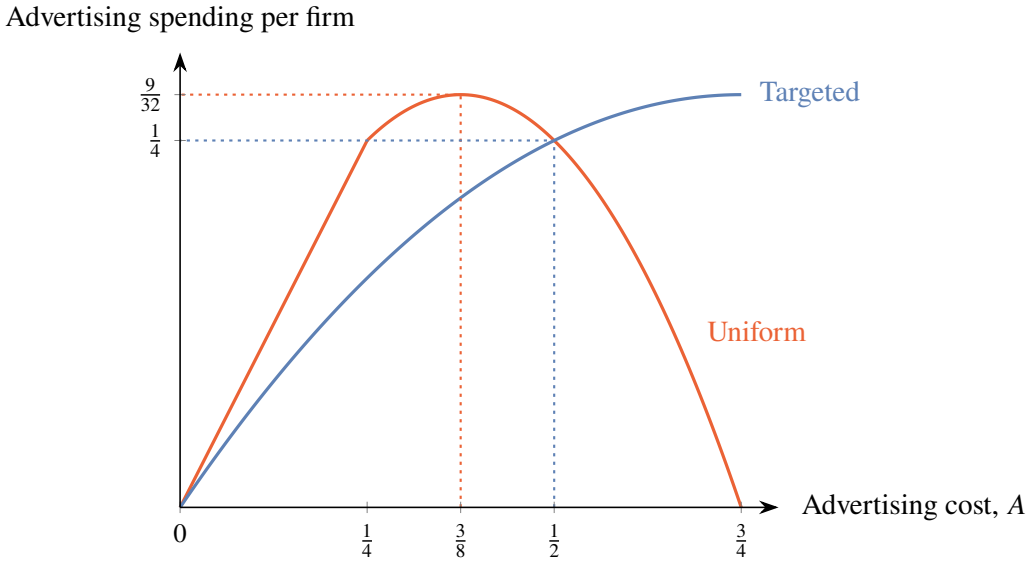


Figure 32.3: Advertising spending per firm with uniform advertising (red) and with targeted advertising (blue).

Figure 32.4 depicts the price distributions at $A = \frac{3}{8}$. With uniform advertising, $c^U = 1$, so prices lie in $\left[\frac{1}{2}, 1\right]$ and average $\ln 2 \approx 0.693$. With targeting, $c^T = \frac{7}{5}$, so prices lie in $\left[\frac{7}{12}, 1\right]$ and average

$$\frac{7}{5} \ln \frac{12}{7} \approx 0.755.$$

At $A = \frac{5}{8}$ the ranking flips: the lowest price is $\frac{5}{6}$ with uniform advertising and $\frac{3}{4}$ with targeting.[1]

[1]The paper notes that targeting raises average prices. In the toy model, this holds for $A < \frac{1}{2}$.

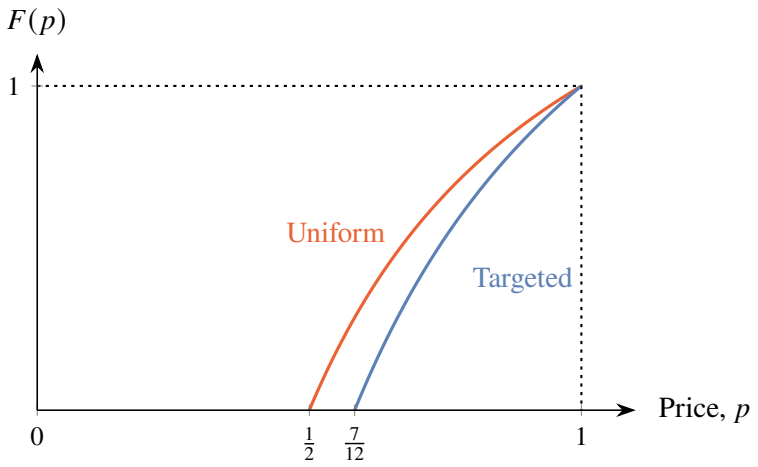


Figure 32.4: Price distributions at $A = \frac{3}{8}$ with uniform advertising (red) and with targeted advertising (blue).

Intuitively, targeting has two effects on spending. It eliminates wasted ads, which lowers spending, and it makes each ad more effective, which raises the incentive to advertise. When advertising is inexpensive, uniform advertisers already advertise with a high probability and waste many ads, so the first effect dominates. When advertising is costly, uniform advertisers advertise with a low probability, and targeting lets firms reach the consumers who are likely to buy, so the second effect dominates. The same comparison drives prices. Comparison shoppers receive ads with probability $\phi^T = 1 - A$ under targeting and $\phi^U = \frac{3}{2} - 2A$ under uniform advertising, and $\phi^T - \phi^U = A - \frac{1}{2}$. If $A < \frac{1}{2}$, targeting firms reach comparison shoppers with a lower probability, so competition for them is softer and prices are higher; if $A > \frac{1}{2}$, uniform advertisers are the ones less likely to meet a rival, and their prices are higher. In contrast with this competitive setting, Esteban et al. (2001) find that targeting lowers the advertising spending of a monopolist.

32.7 Fifth Step: Targeted Pricing

Consider that firms can also charge different prices to their loyal consumers and to comparison shoppers. Each firm then charges 1 to its loyal consumers, and competition for comparison shoppers no longer affects the price that loyal consumers pay. With targeted advertising, a firm that advertises to comparison shoppers and charges them p earns

$$\frac{1}{2}\, p \left[(1-\phi) + \phi(1-F(p))\right] - \frac{A}{2}$$

from that segment, and indifference with not advertising to them gives $\phi = 1 - A$ and $c = \frac{A}{1-A}$ in (32.1). With uniform advertising, the advertising probability is again $\phi^U = \frac{3}{2} - 2A$ when $A > \frac{1}{4}$, and the price to comparison shoppers has $c = \frac{4A-1}{3-4A}$.

Result 4 (targeted pricing, Proposition 4 and Table 1 of the paper). Adding targeted pricing, with either form of advertising:

- leaves the advertising probabilities and profits unchanged;
- lowers the prices paid by comparison shoppers and raises the price paid by loyal consumers to 1.

Table 32.1 reports the four combinations at $A = \frac{3}{8}$. Only the targeting of advertising changes profits, from 0 to $\frac{5}{32}$. Targeted pricing lowers the average price paid by comparison shoppers from 0.693 to 0.462 with uniform advertising and from 0.755 to 0.588 with targeted advertising.

Table 32.1: Outcomes at $A = \frac{3}{8}$. Prices are averages conditional on the firm advertising to the comparison shoppers.

Advertising	Uniform	Targeted	Uniform	Targeted
Pricing	Uniform	Uniform	Targeted	Targeted
Probability of advertising to own loyal consumers	$\frac{3}{4}$	1	$\frac{3}{4}$	1
Probability of advertising to comparison shoppers	$\frac{3}{4}$	$\frac{5}{8}$	$\frac{3}{4}$	$\frac{5}{8}$
Average price to loyal consumers	0.693	0.755	1	1
Average price to comparison shoppers	0.693	0.755	0.462	0.588
Advertising spending per firm	$\frac{9}{32}$	$\frac{27}{128}$	$\frac{9}{32}$	$\frac{27}{128}$
Profit per firm	0	$\frac{5}{32}$	0	$\frac{5}{32}$

Intuitively, targeted pricing has two opposite effects. It lets each firm charge 1 to its loyal consumers, but it also removes the cost of cutting prices for comparison shoppers, since a price cut no longer applies to loyal consumers. Competition for comparison shoppers becomes fiercer, and in equilibrium the two effects cancel. Targeted advertising, instead, softens that competition, which is why the ability to target advertising is more valuable to competing firms than the ability to target prices.

32.8 Sixth Step: Extension to Investing in Targeting

One targeting firm. Consider that firm 1 can target its advertising while firm 2 can only advertise uniformly, both with uniform pricing. An outcome in which both firms advertise to the comparison shoppers with probability one is not an equilibrium, since firm 1 would then gain by skipping them part of the time (the paper's Lemma 2).

Result 5 (one targeting firm, Proposition 5 of the paper). Firm 1 always advertises to its loyal consumers, and:

- if $A < \frac{1}{2}$, firm 2 always advertises and firm 1 advertises to the comparison shoppers with probability $1 - A$; firm 1 earns $\frac{1-A}{4}$ and firm 2 earns $\frac{1-2A}{4}$;
- if $A > \frac{1}{2}$, firm 1 always advertises to the comparison shoppers and firm 2 advertises with probability $\frac{3}{2} - 2A$; firm 1 earns $\frac{A}{4}$, which increases in A, and firm 2 earns zero.

Intuitively, if $A > \frac{1}{2}$, a costlier ad hurts firm 1 directly but makes its uniform rival advertise with a lower probability, and the second effect dominates.[2]

Investment. Before competing, each firm decides whether to pay a fixed cost f to acquire the targeting technology. Table 32.2 reports the profits, gross of f, in the competition stage, where Π^U is the profit when both firms advertise uniformly, Π^T when both target, Π^a the profit of the only targeting firm, and Π^d the profit of its uniform rival.

Table 32.2: Profits in the competition stage, gross of the investment cost f.

Advertising cost	Π^U	Π^a	Π^d	Π^T
$0 < A \leq \frac{1}{4}$	$\frac{1}{4} - A$	$\frac{1-A}{4}$	$\frac{1-2A}{4}$	$\frac{1-A}{4}$
$\frac{1}{4} < A < \frac{1}{2}$	0	$\frac{1-A}{4}$	$\frac{1-2A}{4}$	$\frac{1-A}{4}$
$\frac{1}{2} < A < \frac{3}{4}$	0	$\frac{A}{4}$	0	$\frac{1-A}{4}$

Both firms invest if investing is a best response to an investing rival, $\Pi^T - f > \Pi^d$, and neither invests if not investing is a best response to a rival that does not invest, $\Pi^U > \Pi^a - f$. Between the two cutoffs, only one firm invests.

Result 6 (investing in targeting, Proposition 6 of the paper). Both firms invest if f is below a lower cutoff $\underline{f}$, neither invests if f is above an upper cutoff $\overline{f}$, and only one firm invests if $\underline{f} < f < \overline{f}$, where:

- if $A \leq \frac{1}{4}$, $\underline{f} = \frac{A}{4}$ and $\overline{f} = \frac{3A}{4}$;
- if $\frac{1}{4} < A < \frac{1}{2}$, $\underline{f} = \frac{A}{4}$ and $\overline{f} = \frac{1-A}{4}$;
- if $\frac{1}{2} < A < \frac{3}{4}$, $\underline{f} = \frac{1-A}{4}$ and $\overline{f} = \frac{A}{4}$.

Figure 32.5 depicts the three regions in the (A, f) plane.[3] The interval where only one firm invests is widest at $A = \frac{1}{4}$ and as A approaches $\frac{3}{4}$, and it closes at $A = \frac{1}{2}$, where both cutoffs equal $\frac{1}{8}$. At $A = \frac{3}{8}$, for instance, both firms invest if $f < \frac{3}{32}$, only one invests if $\frac{3}{32} < f < \frac{5}{32}$, and neither invests if $f > \frac{5}{32}$.

Intuitively, ex ante identical firms can choose different technologies because targeting is worth more against a uniform rival, $\Pi^a - \Pi^U$, than against a targeting rival, $\Pi^T - \Pi^d$. A uniform rival either advertises with probability one, so the targeting firm can step back from the comparison shoppers, or advertises with a low probability, so the targeting firm faces little competition for them. The uniform rival does not lose when the other firm targets, since $\Pi^d \geq \Pi^U$, and it gains if $A < \frac{1}{2}$: a firm that advertises less to comparison shoppers softens competition for both.

[2] For $\frac{1}{4} < A < \frac{1}{2}$, the paper also lists the second type of equilibrium. In the toy model, firm 1 would earn $\frac{A}{4}$ in it but $\frac{1-A}{4}$ by advertising to its loyal consumers only, so only the first type survives. The paper selects the first type in this range, so the analysis below is unaffected.

[3] The paper states $\overline{f} = A(1-h)$, that is, $\frac{3A}{4}$, for all $A < \frac{1}{2}$. That cutoff uses $\Pi^U = \frac{1}{4} - A$, which is the uniform profit only when $A \leq \frac{1}{4}$. For $A > \frac{1}{4}$, $\Pi^U = 0$ by Result 1, which gives $\overline{f} = \Pi^a = \frac{1-A}{4}$.

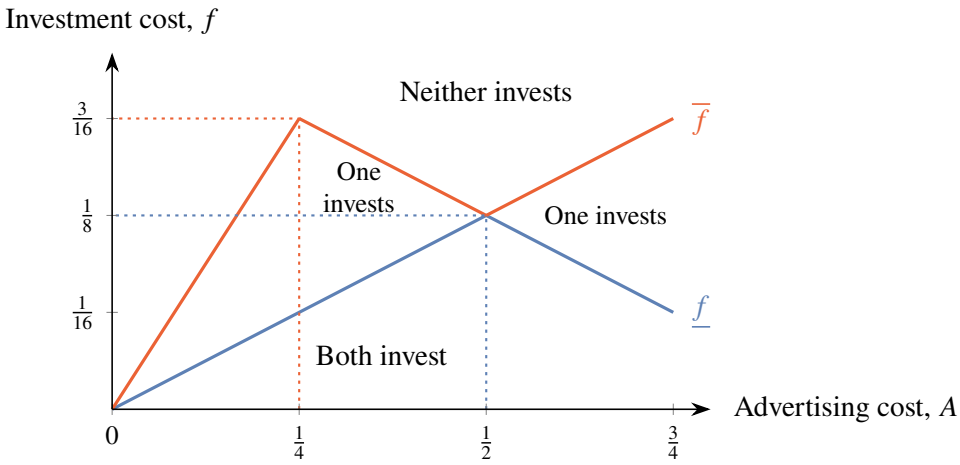


Figure 32.5: Equilibrium investment in targeting: both firms invest below $\underline{f}$ (blue), neither invests above $\overline{f}$ (red), and only one firm invests in between.

32.9 Seventh Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Mixed advertising and pricing strategies, and the result that targeting firms advertise more to their loyal consumers than to comparison shoppers.
- The profit gain from targeting, through less wasted advertising and softer price competition.
- The reversal of advertising spending at $A = \frac{1}{2}$, and the irrelevance of targeted pricing for profits.
- The investment game, with an interval of fixed costs where only one firm acquires targeting.

What it drops.

- General segment sizes and reservation prices: the toy model fixes $h = \frac{1}{4}$ and a value of 1, so the paper's cutoffs on A become $\frac{1}{4}$ and $\frac{1}{2}$, and the comparative statics in h are lost.
- Consumers who know the products without advertising (Section 3.5 of the paper), where firms may advertise less to their loyal consumers than to comparison shoppers.
- Imperfect targeting, where ads leak to other segments and lower profits (Section 3.6).
- Markets with more than two firms, and pricing after observing the rival's advertising.

32.10 Connections to Other Chapters

- Chapter 35 studies ads sold by platforms through auctions, where the ads determine which products consumers discover.
- Chapter 36 studies firms that target their prices with data on consumers' positions in a social network.

- Chapter 4 studies firms that learn their customers' preferences from past purchases and use these data to target prices.
- Chapter 33 studies what an advertisement says rather than who sees it: hype shifts demand, while real product information rotates it.

References

Esteban, L., Gil, A., and Hernández, J. M. (2001). Informative advertising and optimal targeting in a monopoly. *Journal of Industrial Economics*, 49(2), 161–180.

Iyer, G., Soberman, D., and Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3), 461–476.

Varian, H. R. (1980). A model of sales. *American Economic Review*, 70(4), 651–659.

Chapter 33

Hype, Real Information, and the Preference for Extremes

Companion note to Johnson, J. P., and D. P. Myatt (2006), "On the Simple Economics of Advertising, Marketing, and Product Design," American Economic Review 96(3): 756–784.

33.1 Introduction

Advertising does not always make consumers like a product more. Consider a car maker with two possible campaigns. One praises the car's reliability, a feature that every buyer values. The other offers detailed reviews and test drives, so that each buyer learns whether the car suits his own tastes. How much should a firm tell consumers about its product, and how unusual should it make the product?

Johnson and Myatt (2006) call the first kind of advertising hype and the second real information. Hype raises every consumer's willingness to pay, so it shifts the demand curve outward. Real information spreads valuations apart without changing their average: the most eager consumers become more eager and the least eager become less eager. Therefore, it rotates the demand curve around the consumer with the average valuation. A mass-market seller, which sells to most consumers, loses from this rotation, because its marginal buyer is below average. A niche seller, which sells only to the enthusiasts, gains from it, because its marginal buyer is above average.

As a result, profits are U-shaped in the spread of valuations, and the firm prefers one of two extremes. It either takes a mass-market posture, with a plain design that offers something for everyone and advertising that is pure hype, or a niche posture, with an idiosyncratic design and advertising full of real information. Design and advertising content go together: the firm never pairs a plain design with detailed information.

The toy model uses uniformly distributed valuations, for which every price, output, and cutoff has a simple closed form. The first step shows that hype shifts demand while real information rotates it,

and the second step derives the U-shaped profit of a monopolist. The third step shows that the firm provides either no real information or full information, and the fourth step adds product design. The fifth step studies hype over a product's life, the sixth step extends the analysis to Cournot competition, the seventh step studies welfare, and the eighth step lists what the toy model keeps and drops.

33.2 The Toy Model

Consumers and the product. A unit mass of consumers each buy at most one unit of a product. A consumer's true match value with the product is

$$\omega = \mu + \kappa\,\varepsilon, \qquad \varepsilon \sim U[-1, 1], \tag{33.1}$$

drawn independently across consumers, where $\mu > 0$ is the average match value and $\kappa > 0$ measures how idiosyncratic the design is. A plain design (small κ) pleases every consumer about equally, while an idiosyncratic design (large κ) is loved by some consumers and disliked by others. Until the fourth step, $\kappa = 1$.

Real information. Before buying, a consumer does not know ω, but he sees an advertisement x; the same structure describes a review, a demonstration, or a test drive. As in Lewis and Sappington (1994), whose specification the paper uses, the advertisement is "truth or noise": with probability s it reveals the match value, $x = \omega$, and with probability $1 - s$ it is an independent draw from the distribution of ω. The parameter $s \in [s_0, 1]$ is the real-information content of advertising, where $s_0 \geq 0$ is the information consumers obtain anyway, from word of mouth or independent reviews. Since x has the same distribution whether it is truth or noise, a consumer who sees x still believes that it is truth with probability s, so his expected match value is

$$E[\omega \mid x] = s\,x + (1 - s)\,\mu.$$

Hype. Hype announces the product or praises a feature that every consumer values, such as reliability. It raises every consumer's willingness to pay by $h \geq 0$. The fifth step adds the other role of hype, telling consumers that the product exists.

Valuations and demand. A risk-neutral consumer who sees x is willing to pay

$$\theta = h + s\,x + (1 - s)\,\mu.$$

Since x is uniform on $[\mu - \kappa, \mu + \kappa]$, valuations are uniform as well,

$$\theta \sim U[\mu + h - \sigma,\ \mu + h + \sigma], \qquad \text{where} \quad \sigma \equiv s\,\kappa \tag{33.2}$$

is the spread of valuations (half the width of their support). Hype moves the center of the distribution, while real information and design move its spread. The average valuation $\mu + h$ does not depend on σ: information changes how consumers' expectations are spread, not their average.

Firm. A monopolist produces at constant marginal cost $c \geq 0$ and charges a single price. Let

$$a \equiv \mu + h - c > 0 \tag{33.3}$$

denote its margin on a consumer with the average valuation. If $\sigma > 0$ and the firm sells $z \in [0, 1]$ units, the price at which z consumers buy is

$$P(z) = \mu + h + \sigma\,(1 - 2z), \tag{33.4}$$

and if $\sigma = 0$ every consumer is willing to pay $\mu + h$. The firm first chooses its design κ, its information content s, and its hype; consumers then see the advertisement, the firm sets its price, and consumers buy. In the figures, $\mu = \frac{1}{2}$ and $h = c = 0$, so $a = \frac{1}{2}$.

33.3 First Step: Hype Shifts Demand, Real Information Rotates It

Hype and real information change the inverse demand (33.4) in different ways. Hype raises the price at which every quantity sells, $\frac{\partial P}{\partial h} = 1$. A larger spread raises the price of the first units and lowers the price of the last ones, $\frac{\partial P}{\partial \sigma} = 1 - 2z$, which is positive for $z < \frac{1}{2}$ and negative for $z > \frac{1}{2}$. Marginal revenue, $\mu + h + \sigma(1 - 4z)$, responds to σ in the same way around $z = \frac{1}{4}$.

Result 1 (shifts and rotations). In the toy model:

- an increase in hype h shifts the inverse demand (33.4) upward by the same amount at every quantity;
- an increase in the spread σ rotates the inverse demand clockwise around the rotation point $(z^\dagger, p^\dagger) = (\frac{1}{2}, \mu + h)$; equivalently, demand

$$D(p) = \frac{\mu + h + \sigma - p}{2\sigma}$$

  increases in σ at prices above $\mu + h$ and decreases at prices below $\mu + h$;
- marginal revenue rotates clockwise around $z^\ddagger = \frac{1}{4} < z^\dagger$.

Figure 33.1 depicts both changes. In panel (a), the spread rises from $\frac{1}{4}$ to $\frac{1}{2}$ to 1 and every inverse demand passes through the rotation point $(\frac{1}{2}, \frac{1}{2})$. In panel (b), with $\sigma = \frac{1}{2}$, hype $h = \frac{1}{4}$ moves the inverse demand from $1 - z$ to $\frac{5}{4} - z$.

Intuitively, real information does not change the average valuation; it only spreads valuations around it. A consumer whose advertisement is good news raises his valuation, and a consumer whose advertisement is bad news lowers it, so the most eager consumers become more eager and the least eager become less eager. The consumer in the middle of the distribution, whose valuation is $\mu + h$, is not affected, which is why every inverse demand curve in panel (a) passes through the same point. Hype, instead, makes every consumer more eager by the same amount. A more idiosyncratic design, a larger κ, has the same effect as more real information, since valuations depend on the two only through $\sigma = s\kappa$.

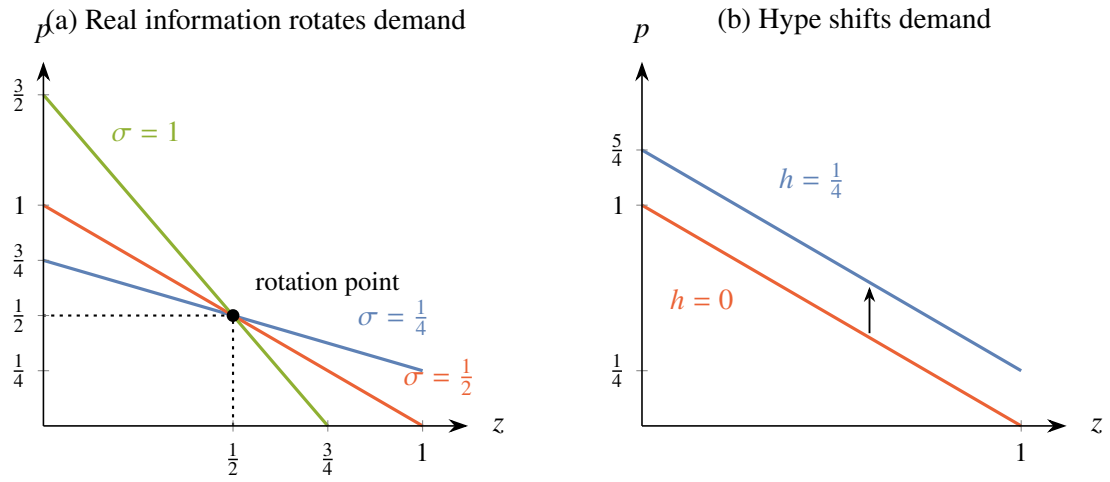


Figure 33.1: Inverse demand $P(z) = \mu + h + \sigma(1 - 2z)$ with $\mu = \frac{1}{2}$: (a) $h = 0$ and three spreads; (b) $\sigma = \frac{1}{2}$, without and with hype.

33.4 Second Step: The Monopolist's Preference for Extremes

Consider first a given spread $\sigma > 0$. The firm chooses z to maximize

$$z\,[P(z) - c] = z\,[a + \sigma(1 - 2z)].$$

The first-order condition, $a + \sigma - 4\sigma z = 0$, gives

$$z^* = \frac{a + \sigma}{4\sigma}, \tag{33.5}$$

which is at most one if and only if $\sigma \geq \frac{a}{3}$. With a smaller spread, every consumer values the product close to the average, and the firm sells to all of them at the valuation of the least eager consumer, $\mu + h - \sigma$. Substituting (33.5) into the objective,

$$\pi(\sigma) = \begin{cases} a - \sigma & \text{if } \sigma \leq \frac{a}{3}, \\ \dfrac{(a+\sigma)^2}{8\sigma} & \text{if } \sigma \geq \frac{a}{3}, \end{cases} \qquad \text{with} \qquad \frac{d}{d\sigma}\frac{(a+\sigma)^2}{8\sigma} = \frac{(a+\sigma)(\sigma - a)}{8\sigma^2}. \tag{33.6}$$

The two pieces of (33.6) meet at $\sigma = \frac{a}{3}$ with the same value, $\frac{2a}{3}$, and the same slope, -1. Following the paper, the firm is a *mass-market* seller if it sells more than the rotation quantity $z^\dagger = \frac{1}{2}$, and a *niche* seller if it sells less.

Result 2 (a preference for extremes). Given the spread σ, the monopolist:

- serves every consumer if $\sigma \leq \frac{a}{3}$, at price $\mu + h - \sigma$, and serves $z^* = \frac{a+\sigma}{4\sigma}$ consumers if $\sigma \geq \frac{a}{3}$, at price $p^* = c + \frac{a+\sigma}{2}$, earning $\pi(\sigma)$ in (33.6);
- is a mass-market seller if $\sigma < a$ and a niche seller if $\sigma > a$;
- earns a profit $\pi(\sigma)$ that is convex in σ, decreases for $\sigma < a$, increases for $\sigma > a$, and reaches its minimum, $\frac{a}{2}$, at $\sigma = a$;

- gains $\frac{d\pi}{dh} = z^*$ from a marginal increase in hype, which equals one when it serves every consumer and decreases in σ otherwise.

Figure 33.2 depicts output and profit for $a = \frac{1}{2}$. In panel (a), output equals one until $\sigma = \frac{1}{6}$, crosses the rotation quantity $\frac{1}{2}$ at $\sigma = a = \frac{1}{2}$, and falls to $\frac{3}{8}$ at $\sigma = 1$. In panel (b), profit falls from $\frac{1}{2}$ at $\sigma = 0$ to $\frac{1}{3}$ at $\sigma = \frac{1}{6}$ and to its minimum $\frac{1}{4}$ at $\sigma = \frac{1}{2}$, and then rises to $\frac{9}{32}$ at $\sigma = 1$.

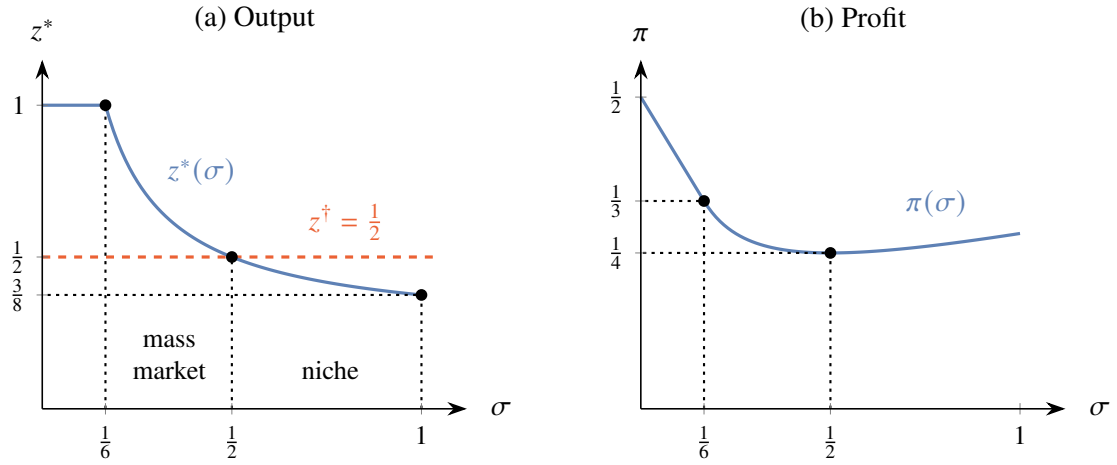


Figure 33.2: Monopoly output and profit as functions of the spread σ, for $a = \frac{1}{2}$.

Intuitively, a mass-market seller's marginal consumer is below average. A larger spread lowers his willingness to pay, and with it the price of every unit the firm sells, so profit falls. A niche seller's marginal consumer is above average, and a larger spread raises his willingness to pay, so profit rises. Once the spread is large enough for the firm to become a niche seller, a larger spread makes it sell even less, so it never returns to the mass market; this is why profit first falls and then rises. The paper shows that this argument works whenever the rotation quantity weakly increases with dispersion; in the toy model the rotation quantity is fixed at $\frac{1}{2}$, because information does not change the average valuation. Output falls with σ throughout, since z^* always exceeds marginal revenue's rotation point $\frac{1}{4}$. The last bullet says that hype is worth more to a mass-market seller, which applies the higher willingness to pay to more consumers.

33.5 Third Step: All or Nothing: Real Information in Advertising

Consider now that the firm chooses the information content $s \in [s_0, 1]$ of its advertising, with $\kappa = 1$, so $\sigma = s$. Providing real information costs nothing beyond the campaign itself. Since $\pi(\sigma)$ is quasi-convex, the firm picks an endpoint, and it compares $\pi(s_0)$ with $\pi(1) = \frac{(a+1)^2}{8}$. Define the threshold

$$\hat{s}(a) = \begin{cases} a - \dfrac{(a+1)^2}{8} & \text{if } a \le \frac{1}{3}, \\ a^2 & \text{if } \frac{1}{3} \le a \le 1. \end{cases} \tag{33.7}$$

For $a \le \frac{1}{3}$, $\hat{s}(a)$ solves

$$a - s_0 = \frac{(a+1)^2}{8},$$

which compares the two endpoints when the firm serves every consumer at s_0; this threshold is negative for $a < 3 - 2\sqrt{2}$. For $\frac{1}{3} \le a \le 1$, $\hat{s}(a)$ solves

$$\frac{(a+s_0)^2}{8s_0} = \frac{(a+1)^2}{8},$$

whose roots are $s_0 = a^2$ and $s_0 = 1$. The two pieces meet at $\hat{s}(\frac{1}{3}) = \frac{1}{9}$.

Result 3 (all or nothing). With $\sigma = s \in [s_0, 1]$:

- the firm chooses either $s = s_0$, adding no real information to what consumers already have (its advertising is pure hype), or $s = 1$, full information;
- if $a < 1$, it chooses full information if and only if $s_0 > \hat{s}(a)$; if $a \ge 1$, it never adds information;
- with $s_0 = 0$, it chooses full information if and only if $\frac{(a+1)^2}{8} > a$, that is, $a < 3 - 2\sqrt{2} \approx 0.172$;
- the threshold $\hat{s}(a)$ increases in a.

Figure 33.3 depicts the threshold $\hat{s}(a)$. Above the curve the firm provides full information and becomes a niche seller; below it the firm adds no information. For $a < 3-2\sqrt{2}$, the firm provides full information at every level of outside information. For $a = \frac{1}{2}$, the firm switches to full information once outside information exceeds $\hat{s}(\frac{1}{2}) = \frac{1}{4}$.

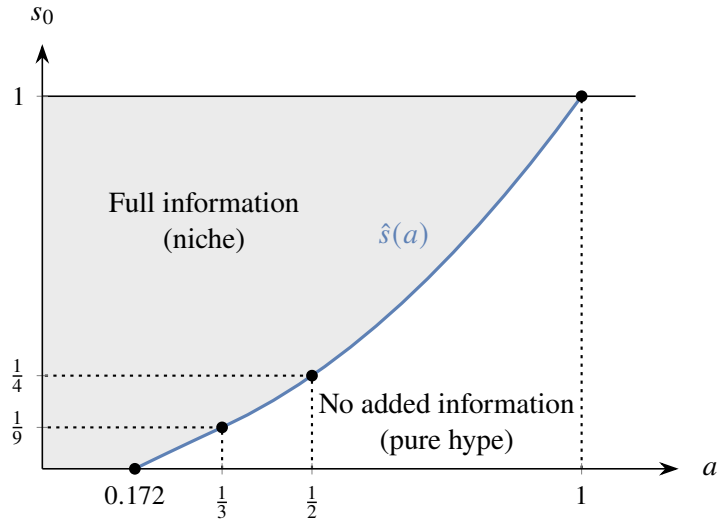


Figure 33.3: The firm's choice of real information, as a function of its margin a and of the outside information s_0.

Intuitively, consider first $s_0 = 0$. Without real information, every consumer has the same valuation, and the firm sells to all of them and extracts all their surplus, earning a. With full information, the firm sells at a high price only to the consumers who learn that they like the product. When the

margin a is small, selling to everyone at the average valuation barely covers the cost, and selling only to the enthusiasts is better. Outside information makes the pure-hype option less attractive, because it already spreads valuations, which is why the firm chooses full information over a wider range of margins as s_0 rises. The last bullet carries two predictions. A higher marginal cost lowers a and pushes the firm toward the niche, as the paper argues for families of distributions with a constant mean; and more hype raises a and pushes the firm toward the mass market, so hype and real information are substitutes. The paper notes that this all-or-nothing choice reproduces the result of Lewis and Sappington (1994), and explains it as a consequence of demand rotation.

33.6 Fourth Step: Product Design and Advertising Content

Consider now that the firm chooses both its design $\kappa \in [\kappa_L, \kappa_H]$, with $0 < \kappa_L < \kappa_H$, and its information content $s \in [s_0, 1]$, with $0 < s_0 < 1$. By (33.2), valuations depend on the two choices only through the spread $\sigma = s\kappa$, which ranges from $s_0\kappa_L$ to κ_H. If $s_0\kappa_L \le \frac{a}{3}$, the plain design with no added information earns $a - s_0\kappa_L$, and the idiosyncratic design with full information earns $\frac{(a+\kappa_H)^2}{8\kappa_H}$ if $\kappa_H \ge \frac{a}{3}$. The two are equal when

$$\kappa_H^2 - 2(3a - 4s_0\kappa_L)\kappa_H + a^2 = 0,$$

whose larger root is

$$\bar{\kappa} \equiv 3a - 4s_0\kappa_L + 2\sqrt{2\,(a - 2s_0\kappa_L)(a - s_0\kappa_L)}. \tag{33.8}$$

Result 4 (design and information as complements). When the firm chooses $(\kappa, s) \in [\kappa_L, \kappa_H] \times [s_0, 1]$:

- it chooses either the plain design with no added information, (κ_L, s_0), or the idiosyncratic design with full information, $(\kappa_H, 1)$; it never pairs full information with the plain design, nor the idiosyncratic design with no added information;
- if $s_0\kappa_L \le \frac{a}{3}$, it chooses $(\kappa_H, 1)$ if and only if $\kappa_H > \bar{\kappa}$ in (33.8).

Figure 33.4 depicts the feasible choices for $a = \frac{1}{2}$, $s_0 = \frac{1}{4}$, $\kappa_L = \frac{1}{2}$, and $\kappa_H = 2$, together with three curves of constant spread $s\kappa$. Profit depends only on the curve on which a choice lies; it is lowest on the curve $s\kappa = a = \frac{1}{2}$ and rises as the firm moves away from it in either direction. Corner A lies on the lowest feasible curve, $s\kappa = \frac{1}{8}$, and earns $\frac{3}{8}$. Corner B lies on the highest one, $s\kappa = 2$, and earns $\frac{25}{64}$. Here

$$\bar{\kappa} = 1 + \frac{\sqrt{3}}{2} \approx 1.866 < 2,$$

so the firm chooses B. The other two corners, with $s\kappa = \frac{1}{2}$, earn the minimum profit $\frac{1}{4}$.

Intuitively, design idiosyncrasy and real information are complements. Real information is worth providing only when there is something to learn, and an idiosyncratic design is worth building only when consumers can find out whether it suits them. A mass-market firm therefore sells a "plain vanilla" product promoted by pure hype, while a niche firm sells a "love it or hate it" product

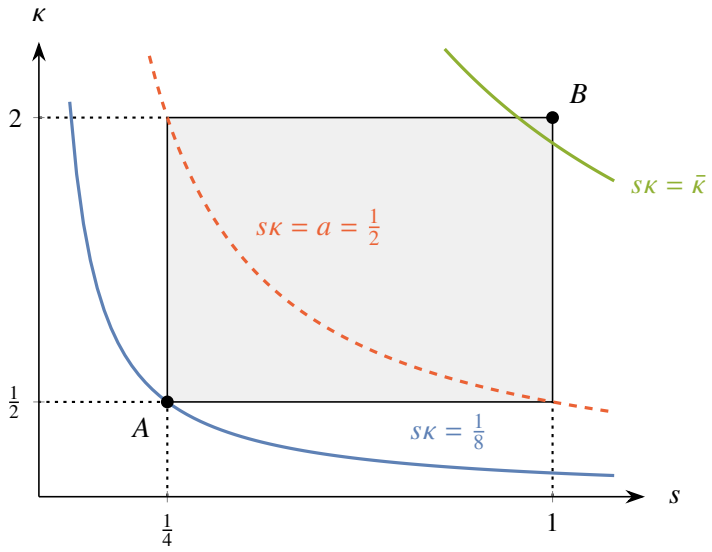


Figure 33.4: Feasible design and information choices (shaded) for $s_0 = \frac{1}{4}$, $\kappa_L = \frac{1}{2}$, and $\kappa_H = 2$. Corner $A = (s_0, \kappa_L)$ is the plain design with pure hype; corner $B = (1, \kappa_H)$ is the idiosyncratic design with full information.

with detailed advertising. The paper obtains the same pairing with normally distributed valuations and risk-neutral consumers: the firm never advertises a plain-vanilla product informatively. With risk-averse consumers, the paper shows that the pairing can fail, because information then also raises the average valuation by lowering the risk premium, while an idiosyncratic design lowers it.

33.7 Fifth Step: Hype Intensity and the Advertising Life Cycle

So far, hype raised every valuation. Its other role is to tell consumers that the product exists. Consider that only a fraction $n \in [0, 1]$ of consumers is aware of the product, that reaching a fraction n costs $\frac{k\,n^2}{2}$, and that k is large enough for the optimal n to be below one. If Π denotes the profit per aware consumer, the firm maximizes $n\,\Pi - \frac{k\,n^2}{2}$, so its hype intensity and its total profit are

$$n^* = \frac{\Pi}{k}, \qquad n^*\,\Pi - \frac{k\,(n^*)^2}{2} = \frac{\Pi^2}{2k}. \tag{33.9}$$

Since total profit increases in Π, the firm's choice of information content is the one in Result 3, and $\Pi = \max\{\pi(s_0), \pi(1)\}$. Over a product's life, outside information s_0 grows as reviews and word of mouth accumulate.

Result 5 (hype intensity over the life cycle). With $\kappa = 1$:

- hype intensity is $n^* = \frac{1}{k}\max\{\pi(s_0), \pi(1)\}$;
- with $a = \frac{1}{2}$ and $k = 1$, as s_0 rises from 0 to 1, the firm adds no information and n^* falls from $\frac{1}{2}$

to $\frac{9}{32}$ while $s_0 < \frac{1}{4}$, with

$$n^* = \begin{cases} \dfrac{1}{2} - s_0 & \text{if } s_0 \le \frac{1}{6}, \\ \dfrac{(\frac{1}{2} + s_0)^2}{8 s_0} & \text{if } \frac{1}{6} \le s_0 \le \frac{1}{4}; \end{cases}$$

for $s_0 > \frac{1}{4}$, the firm provides full information and $n^* = \frac{9}{32}$;

- if the firm could not add information, its hype intensity would be $\frac{\pi(s_0)}{k}$, which is U-shaped in s_0, with minimum $\frac{1}{4}$ at $s_0 = \frac{1}{2}$, and returns to $\frac{9}{32}$ at $s_0 = 1$.

Figure 33.5 depicts the two paths of hype intensity. The solid curve is the firm's choice; the dashed curve is the path without added information.

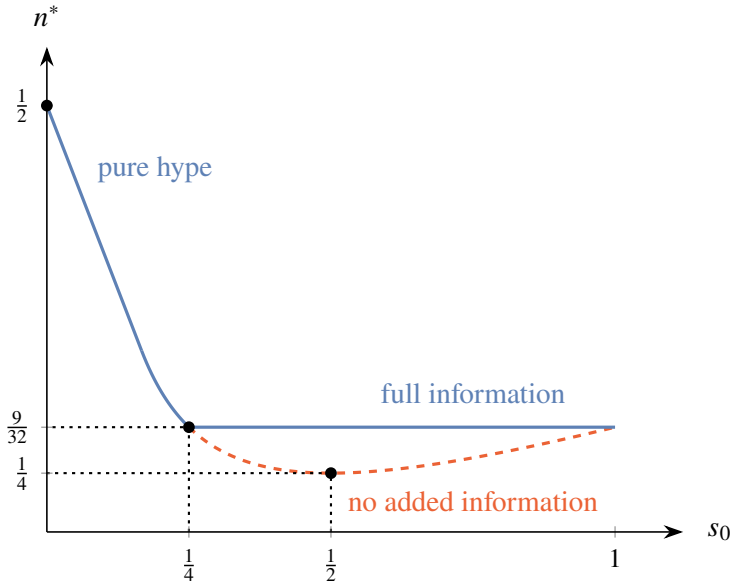


Figure 33.5: Hype intensity as outside information s_0 grows, for $a = \frac{1}{2}$ and $k = 1$.

Intuitively, hype is worth more when each aware consumer is worth more, so hype intensity follows the profit curve of the second step. Early in a product's life, consumers know little, the firm is a mass-market seller, and its advertising is intense pure hype. As outside information grows, the firm's profit per aware consumer falls, and so does its hype intensity. Once outside information is large enough, the firm switches to a niche posture with full information, which mirrors the paper's description of an advertising life cycle. The paper argues that hype intensity then rises again. In the toy model this rise appears only along the dashed path, because the firm can move to full information at once, after which its profit no longer depends on s_0.

33.8 Sixth Step: Cournot Competition

Consider now $m \ge 1$ identical firms that sell the same product and compete in quantities, with $\kappa = 1$ and a common spread σ, since reviews and word of mouth reach the customers of all firms.

If total output Q exceeds the mass of consumers, the extra units cannot be sold. For $Q \le 1$, firm i earns $q_i\,[a + \sigma(1 - 2Q)]$, so the first-order condition is

$$a + \sigma - 2\sigma Q - 2\sigma q_i = 0,$$

and in the symmetric equilibrium each firm produces $\frac{a+\sigma}{2\sigma(m+1)}$. Total output is at most one if and only if $\sigma \ge \frac{ma}{m+2}$. For a smaller spread, every firm sells $\frac{1}{m}$, and none wants to cut output, since its marginal revenue, $a - \sigma - \frac{2\sigma}{m}$, is not negative. In this region, asymmetric splits of the market are also equilibria; we focus on the symmetric one.

Result 6 (Cournot competition). In the symmetric equilibrium with m firms:

- if $\sigma \le \frac{ma}{m+2}$, each firm sells $\frac{1}{m}$ at price $\mu + h - \sigma$ and earns $\frac{a-\sigma}{m}$; if $\sigma \ge \frac{ma}{m+2}$, each firm sells $\frac{a+\sigma}{2\sigma(m+1)}$ with margin $\frac{a+\sigma}{m+1}$ and earns

$$\pi_m(\sigma) = \frac{(a+\sigma)^2}{2\sigma(m+1)^2};$$

- for every m, each firm's profit is convex in σ, decreases for $\sigma < a$, and increases for $\sigma > a$; at $\sigma = a$, industry output is $\frac{m}{m+1}$;
- full information ($\sigma = 1$) gives each firm a higher profit than no information ($\sigma = 0$) if and only if

$$a < a_m \equiv B_m - \sqrt{B_m^2 - 1}, \qquad \text{where} \quad B_m = \frac{m^2 + m + 1}{m};$$

  $a_1 = 3 - 2\sqrt{2}$, as in Result 3, and a_m decreases in m and tends to zero.

Table 33.1 reports the cutoff a_m and industry output at the profit-minimizing spread for several numbers of firms.

Table 33.1: Cournot competition with uniform valuations.

Number of firms m	1	2	3	4	5	10
Cutoff a_m for full information	0.172	0.146	0.117	0.096	0.081	0.045
Industry output at $\sigma = a$	$\frac{1}{2}$	$\frac{2}{3}$	$\frac{3}{4}$	$\frac{4}{5}$	$\frac{5}{6}$	$\frac{10}{11}$

Intuitively, competition adds a strategic effect to the direct effect of the second step. With $m \ge 2$ firms, industry output at $\sigma = a$ exceeds the rotation quantity $\frac{1}{2}$, so the marginal consumer is below average and a larger spread lowers the price at given outputs. But a larger spread also steepens demand, which leads rivals to cut their output and raises each firm's profit. With uniform valuations the two effects cancel at $\sigma = a$ for every m. The paper shows that the strategic effect can break the U shape for general rotations, and that with constant-elasticity demand the profit-minimizing dispersion rises with the number of firms. With uniform valuations the minimum stays at $\sigma = a$, but competition still favors the mass market: without information, firms sell to every consumer at his common valuation and competition does not lower the margin, while with full information firms compete for the enthusiasts and the margin falls to $\frac{a+1}{m+1}$. The range of margins for which full information is better therefore shrinks as m grows.

33.9 Seventh Step: Welfare

The paper leaves welfare aside; the toy model lets us compute it. Return to the monopolist, with $\kappa = 1$, $s_0 = 0$, and $a < 1$, and measure consumer surplus with valuations (33.2); since consumers are risk neutral, this equals their expected surplus from the true match values. With no information, the firm sells to every consumer at his valuation, so consumer surplus is zero and total surplus equals profit, a. With full information, the firm sells $z^* = \frac{a+1}{4}$ units along the inverse demand $\mu + h + 1 - 2z$, so consumer surplus is

$$z^* \cdot \frac{2z^*}{2} = \frac{(a+1)^2}{16}$$

and total surplus is

$$\frac{(a+1)^2}{8} + \frac{(a+1)^2}{16} = \frac{3(a+1)^2}{16}.$$

A planner who also prices at marginal cost sells to every consumer with $\theta \geq c$ and obtains $\frac{(a+1)^2}{4}$ with full information.

Result 7 (welfare). With $\kappa = 1$, $s_0 = 0$, and $a < 1$:

- full information raises consumer surplus from 0 to $\frac{(a+1)^2}{16}$;
- full information raises total surplus, from a to $\frac{3(a+1)^2}{16}$, if and only if $a < \frac{1}{3}$;
- the firm provides full information if and only if $a < 3 - 2\sqrt{2} \approx 0.172$; for $3 - 2\sqrt{2} < a < \frac{1}{3}$, it withholds information that would raise total surplus, and for $\frac{1}{3} < a < 1$, withholding information raises total surplus but lowers consumer surplus;
- with marginal-cost pricing, full information maximizes total surplus, since

$$\frac{(a+1)^2}{4} - a = \frac{(1-a)^2}{4} > 0.$$

Figure 33.6 depicts profit, consumer surplus, and total surplus with and without information as functions of the margin a. Profit with full information crosses the no-information line at $a = 3 - 2\sqrt{2} \approx 0.172$, and total surplus with full information crosses it at $a = \frac{1}{3}$.

Intuitively, the firm's gain from real information differs from society's gain for two reasons. First, with full information the consumers who learn that they like the product keep part of the surplus, which the firm does not count; this is why the firm withholds information for margins between 0.172 and $\frac{1}{3}$. Second, with full information the firm restricts output to sell at a high price, which wastes surplus; this is why, under monopoly pricing, information lowers total surplus when $a > \frac{1}{3}$. Without monopoly pricing, information is always efficient: it lets the consumers who would not enjoy the product avoid buying it.

33.10 Eighth Step: What the Toy Model Keeps and Drops

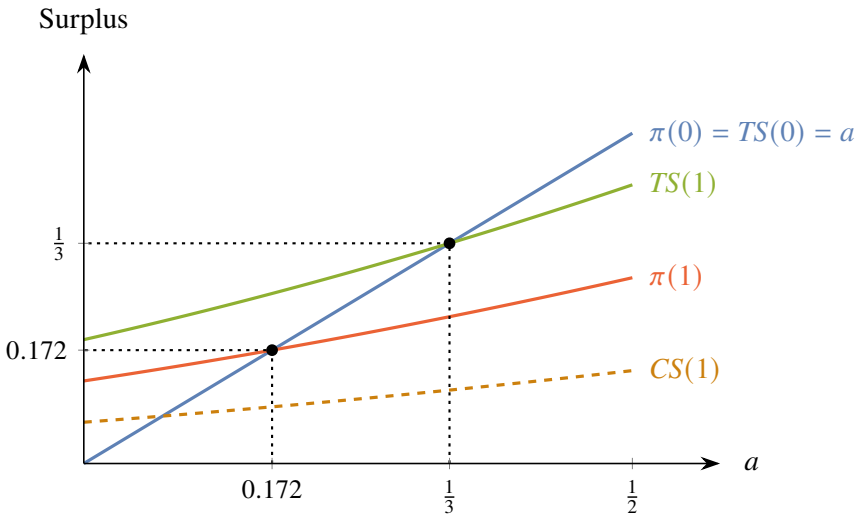


Figure 33.6: Profit π, consumer surplus CS, and total surplus TS with no information ($s = 0$) and full information ($s = 1$), as functions of the margin a.

What the toy model keeps.

- The taxonomy of advertising: hype shifts demand, while real information and design idiosyncrasy rotate it around the consumer with the average valuation.
- U-shaped profits in the spread of valuations, and the two marketing mixes that follow: a plain design with pure hype for the mass market, and an idiosyncratic design with full information for a niche.
- The all-or-nothing choice of real information of Lewis and Sappington (1994), and the prediction that higher costs favor the niche.
- U-shaped profits under Cournot competition, which the paper obtains for constant-elasticity demand.

What it drops.

- Moving rotation points. With risk-averse consumers, information raises the average valuation and design idiosyncrasy lowers it, so the two rotate demand around different points, and the paper shows that the firm may then pair detailed advertising with a plain design.
- Products as mixes of characteristics, as in Lancaster. In the paper, the firm either puts all weight on the characteristic whose valuations are most dispersed or chooses the mix that minimizes dispersion, as in a minimum-variance portfolio.
- Design-dependent costs. If marginal cost rises with the spread at an increasing rate, the paper explains that the extremes need no longer be optimal.
- Product lines of quality-differentiated goods, where the paper shows that more dispersion lengthens the product line, and perfect competition.

33.11 Connections to Other Chapters

- Chapter 32 studies informative advertising that tells consumers about a product, and asks which consumers firms should reach.
- Chapter 34 studies the other side of advertising: free media that sell their viewers' attention to advertisers.
- Chapter 24 studies a firm that discloses its quality or signals it through its price, another way to inform consumers about a product.

References

Johnson, J. P., and Myatt, D. P. (2006). On the simple economics of advertising, marketing, and product design. *American Economic Review*, 96(3), 756–784.

Lewis, T. R., and Sappington, D. E. M. (1994). Supplying information to facilitate price discrimination. *International Economic Review*, 35(2), 309–327.

Chapter 34

Too Many or Too Few Ads? The Market Provision of Broadcasting

Companion note to Anderson, S. P., and S. Coate (2005), "Market Provision of Broadcasting: A Welfare Analysis," Review of Economic Studies 72(4): 947–972.

34.1 Introduction

Free-to-air television does not charge its audience. Consider a station that shows a popular program for free. It earns money by selling commercial slots to advertisers, and viewers pay with the time they spend watching commercials. Many countries cap the number of commercial minutes per hour, which presumes that the market delivers too much advertising. Does the market show too many or too few commercials, and does it offer too many or too few programs?

Broadcasting is a market with two sides: viewers, who value programs and dislike advertising, and advertisers, who value contact with viewers. A station that shows one more commercial earns more from its current viewers, but some of them move to the rival channel. Like any seller with market power, the station also limits the number of commercials to keep their price high. Both forces push advertising down. In contrast, the station ignores the nuisance that commercials impose on the viewers who stay, which pushes advertising up.

Anderson and Coate (2005) show that advertising and programming can each be over- or under-provided. Advertising is excessive when viewers find commercials very annoying. It is insufficient when commercials are a small nuisance, programs are close substitutes, or advertisers value viewers highly. Entry can also be excessive, because a new station takes viewers and revenue from the incumbent. As a result, a cap on advertising can lower welfare, either because the market already shows too few commercials or because the cap drives a station out of the market.

The toy model places viewers on a line, with one channel at each end, and uses a linear demand for advertising. The first step studies a single channel and the first best, the second step solves the duopoly, and the third step compares the duopoly with the first best. The fourth step studies the

number of channels, the fifth step studies advertising caps, and the sixth step lists what the toy model keeps and drops.

34.2 The Toy Model

Viewers. A unit mass of viewers is uniformly distributed on the interval $[0, 1]$, as in Hotelling (1929). Channel A shows a program of type 0, located at 0, and channel B a program of type 1, located at 1. A viewer at λ obtains

$$\beta - \gamma a_A - \tau\lambda \quad \text{from channel } A, \qquad \beta - \gamma a_B - \tau(1-\lambda) \quad \text{from channel } B, \tag{34.1}$$

and zero from not watching, where a_J is the number of advertisements on channel J. The parameter $\gamma > 0$ is the nuisance cost of one advertisement, and $\tau > 0$ measures how different the two programs are for viewers: a lower τ means closer substitutes. Each viewer watches at most one channel. Advertising levels lie in $[0, 1]$, and we assume $\beta \geq \tau + \gamma$, so that every viewer prefers watching any channel to not watching, at any advertising level.

Advertisers. A unit mass of producers of new goods can advertise. A producer's willingness to pay for contacting one viewer is uniformly distributed on $[0, \omega]$ across producers. If a station with V viewers sells a advertisements, the price of one advertisement is $V\,p(a)$, where

$$p(a) = \omega(1 - a) \tag{34.2}$$

is the per-viewer inverse demand for advertising, and the station's revenue is $V\,R(a)$, with per-viewer revenue

$$R(a) = p(a)\,a = \omega a(1 - a).$$

In the paper, a producer of type σ makes a sale of value ω with probability σ, and types have distribution F on $[0, \bar{\sigma}]$; with a uniform F, this is (34.2), where ω stands for the paper's $\bar{\sigma}\omega$. Producers extract all the surplus from their sales, so advertisements give viewers no benefit.

Stations and timing. Programs are free to viewers, and stations are financed only by advertising. In stage 1, each of two stations decides whether to operate, which costs $K > 0$. In stage 2, operating stations simultaneously choose their advertising levels. Viewers then choose a channel, and advertisers buy advertisements. Aggregate surplus is the sum of viewers' net benefits and advertisers' gross benefits from contacts, minus program costs; advertising revenues are transfers from advertisers to stations.

34.3 First Step: A Single Channel

The monopoly station. With one channel, every viewer watches, since

$$\beta - \gamma a - \tau \geq \beta - \gamma - \tau \geq 0.$$

The station's revenue is $R(a) = \omega a(1-a)$, and it chooses the advertising level at which marginal revenue $\omega(1-2a)$ is zero:

$$a^m = \frac{1}{2}, \qquad p(a^m) = \frac{\omega}{2}, \qquad \pi^m = \frac{\omega}{4}. \tag{34.3}$$

The first best. The planner places the a advertisements with the producers who value them most. With one channel, aggregate surplus is

$$W_1(a) = \beta - \frac{\tau}{2} - \gamma a + \int_0^a p(x)\,dx = \beta - \frac{\tau}{2} - \gamma a + \omega\left(a - \frac{a^2}{2}\right), \tag{34.4}$$

where $\frac{\tau}{2}$ is the average distance cost. The marginal social benefit of an advertisement is the willingness to pay of the marginal advertiser, $p(a)$, and its marginal social cost is the nuisance γ it imposes on every viewer. Hence

$$a^o = \max\left\{0,\ 1 - \frac{\gamma}{\omega}\right\}. \tag{34.5}$$

One owner, two channels. If a single owner operates both channels, its revenue is $\lambda R(a_A) + (1-\lambda)R(a_B)$, where λ is channel A's audience. This is a weighted average of per-viewer revenues, so it is at most $R(\frac{1}{2}) = \frac{\omega}{4}$, which the owner reaches by setting $a_A = a_B = \frac{1}{2}$.

Result 1 (single channel, Proposition 1 of the paper). With one channel:

- the station sets $a^m = \frac{1}{2}$, charges $\frac{\omega}{2}$ per viewer, and earns $\frac{\omega}{4}$;
- the station advertises more than optimal, $a^m > a^o$, if and only if $\gamma > \frac{\omega}{2}$, and less than optimal if and only if $\gamma < \frac{\omega}{2}$;
- a single owner of both channels sets $\frac{1}{2}$ advertisements on each channel and earns $\frac{\omega}{4}$, the same revenue as with one channel.

Figure 34.1 depicts the per-viewer demand $p(a)$, the marginal revenue $\omega(1-2a)$, and the marginal social cost γ for $\omega = 1$ and two nuisance costs. With $\gamma = \frac{1}{4}$, the optimal level is $a^o = \frac{3}{4}$, above $a^m = \frac{1}{2}$; with $\gamma = \frac{3}{4}$, it is $a^o = \frac{1}{4}$, below a^m.

Intuitively, the station makes two mistakes that work in opposite directions. Since it sells access to its viewers as a monopolist, it stops where marginal revenue is zero, while the planner continues as long as the marginal advertiser's willingness to pay exceeds the nuisance cost; this restricts advertising. Since no viewer leaves, the station also ignores the nuisance cost, which expands advertising. When advertising is a small nuisance, the first mistake dominates, and the station shows too few advertisements; when it is a large nuisance, the second mistake dominates. A second channel does not raise the owner's revenue, because every viewer already watches and the owner already earns the highest revenue per viewer.

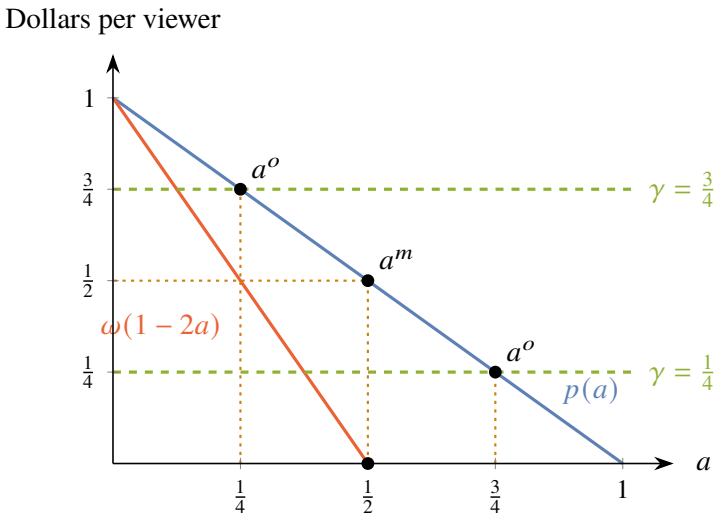


Figure 34.1: Single channel: per-viewer demand for advertising, marginal revenue, and marginal social cost ($\omega = 1$).

34.4 Second Step: Two Competing Channels

Audiences. Consider now that two stations compete. Station A shows a_A advertisements and station B shows a_B. By (34.1), the viewer indifferent between the two channels is located at

$$\hat{\lambda} = \frac{1}{2} + \frac{\gamma\,(a_B - a_A)}{2\tau}, \tag{34.6}$$

so that station A loses $\frac{\gamma}{2\tau}$ viewers for each additional advertisement.

Equilibrium. Station A maximizes $\hat{\lambda}\,R(a_A)$. The first-order condition is

$$\hat{\lambda}\,R'(a_A) = \frac{\gamma}{2\tau}\,R(a_A),$$

and in a symmetric equilibrium, where $\hat{\lambda} = \frac{1}{2}$, it becomes

$$R'(a) = \frac{\gamma}{\tau}\,R(a) \quad \Longleftrightarrow \quad \omega(1-2a) = \frac{\gamma}{\tau}\,\omega a(1-a). \tag{34.7}$$

The left-hand side is marginal revenue with a fixed audience, and the right-hand side is the revenue lost from viewers who move to the rival. The root of (34.7) in $[0, \frac{1}{2})$ is

$$a^* = \frac{2\tau}{\gamma + 2\tau + \sqrt{\gamma^2 + 4\tau^2}}, \tag{34.8}$$

the other root exceeds one, and a^* is a global best response to a^*. The per-viewer price and each station's revenue are

$$p^* = \omega(1-a^*), \qquad \pi^* = \frac{1}{2}R(a^*) = \frac{\omega\tau}{2\left(\sqrt{\gamma^2+4\tau^2}+2\tau\right)}. \tag{34.9}$$

Every viewer watches, since the farthest viewer of each channel obtains $\beta - \gamma a^* - \frac{\tau}{2} > 0$.

Result 2 (duopoly advertising). With two competing stations:

- each station shows a^* in (34.8) advertisements, with $0 < a^* < \frac{1}{2} = a^m$, and charges $p^* > \frac{\omega}{2}$ per viewer;
- a^* depends only on $\frac{\gamma}{\tau}$: it decreases in γ, increases in τ, does not depend on ω, tends to $\frac{1}{2}$ as $\gamma \to 0$, and tends to 0 as $\tau \to 0$.

Figure 34.2 depicts the two sides of (34.7) for $\omega = 1$, $\gamma = \frac{1}{2}$, and $\tau = \frac{1}{3}$, so that $\frac{\gamma}{\tau} = \frac{3}{2}$. Marginal revenue $1 - 2a$ meets the revenue lost to the rival, $\frac{3}{2}a(1-a)$, at $a^* = \frac{1}{3}$, where the per-viewer price is $\frac{2}{3}$. A monopoly station would continue to $a^m = \frac{1}{2}$.

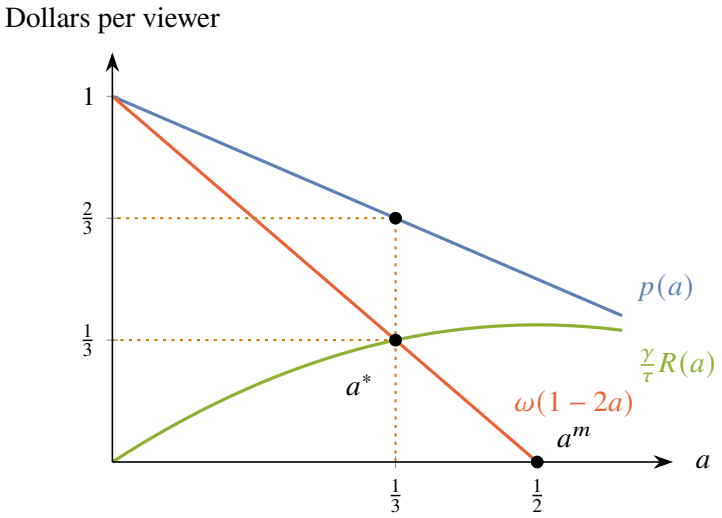


Figure 34.2: Duopoly: marginal revenue and revenue lost to the rival ($\omega = 1$, $\gamma = \frac{1}{2}$, $\tau = \frac{1}{3}$).

Intuitively, a station can attract viewers only by showing fewer advertisements, since programs are free. When programs are close substitutes, many viewers move to the rival after one more advertisement, and competition for viewers keeps advertising low. Each station also keeps a monopoly on access to its own viewers, so it restricts advertising to keep its price high, as in a competitive bottleneck (Armstrong, 2006); both forces push advertising down. The scale ω of advertisers' willingness to pay multiplies both sides of (34.7), so it changes prices and revenues, but not the number of advertisements.

34.5 Third Step: Too Much or Too Little Advertising?

The first best with two channels. With both channels and a advertisements on each, aggregate surplus is

$$W_2(a) = \beta - \frac{\tau}{4} - \gamma a + \omega\left(a - \frac{a^2}{2}\right), \tag{34.10}$$

since every viewer watches the closer program and the average distance cost falls to $\frac{\tau}{4}$. The planner gives both channels the same level, and the same producers advertise on both. The optimal level is again a^o in (34.5): with full coverage, the marginal advertiser's willingness to pay must equal the nuisance cost, $p(a^o) = \gamma$. Since a^o does not depend on τ, while a^* increases in τ, the comparison depends on the cutoff

$$\hat{\tau}(\gamma) \equiv \frac{\gamma^2(\omega - \gamma)}{\omega(2\gamma - \omega)}, \qquad \frac{\omega}{2} < \gamma < \omega, \tag{34.11}$$

which solves $a^* = a^o$. The cutoff decreases in γ, from $+\infty$ as $\gamma \to \frac{\omega}{2}$ to 0 as $\gamma \to \omega$.

Result 3 (Proposition 1 of the paper). With two competing stations:

- if $\gamma \leq \frac{\omega}{2}$, advertising is below the optimal level, $a^* < a^o$;
- if $\frac{\omega}{2} < \gamma < \omega$, advertising is below the optimal level if $\tau < \hat{\tau}(\gamma)$, optimal if $\tau = \hat{\tau}(\gamma)$, and above it if $\tau > \hat{\tau}(\gamma)$;
- if $\gamma \geq \omega$, advertising is above the optimal level, $a^* > a^o = 0$;
- equivalently, for every $\tau > 0$ there is a unique critical nuisance cost $\gamma_2(\tau) \in (\frac{\omega}{2}, \omega)$ such that $a^* < a^o$ if and only if $\gamma < \gamma_2(\tau)$; γ_2 decreases in τ, from ω as $\tau \to 0$ to $\frac{\omega}{2}$ as $\tau \to \infty$;
- the per-viewer price p^* exceeds the Pigouvian price γ if and only if $a^* < a^o$.

Figure 34.3 depicts a^* and a^o as functions of γ for $\omega = 1$ and $\tau = \frac{4}{9}$. The two curves cross at $\gamma_2 = \frac{2}{3}$, where $a^* = a^o = \frac{1}{3}$, since $\hat{\tau}(\frac{2}{3}) = \frac{4}{9}$. The dashed line is the monopoly level $\frac{1}{2}$, which equals a^o at $\gamma = \frac{1}{2}$, as in Result 1.

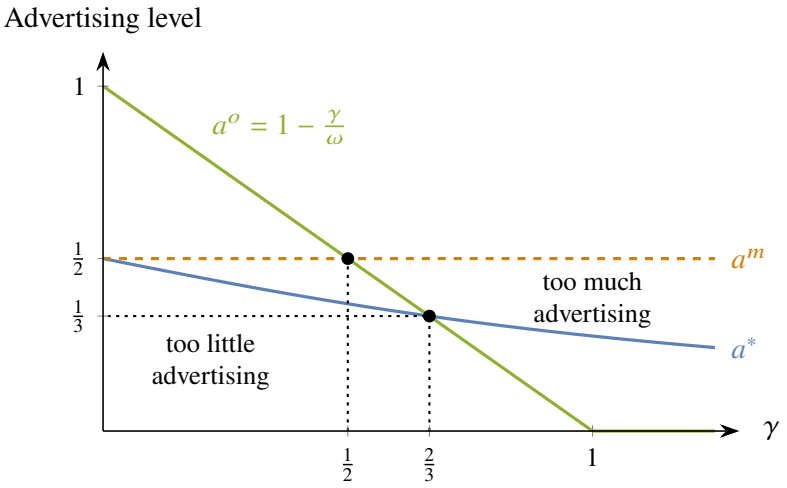


Figure 34.3: Duopoly and optimal advertising levels as functions of the nuisance cost ($\omega = 1$, $\tau = \frac{4}{9}$).

Figure 34.4 depicts the two regions in the (γ, τ) plane for $\omega = 1$. Advertising is too low in the shaded region, to the left of the curve $\tau = \hat{\tau}(\gamma)$, and too high to its right. The dot marks the example of Figure 34.3. The paper shows that γ_2 lies in $(0, \bar{\sigma}\omega)$; with linear demand, it lies in $(\frac{\omega}{2}, \omega)$.

Intuitively, revenue maximization accounts for the nuisance cost only through the viewers who move to the rival, while the planner counts the nuisance imposed on every viewer. When programs are

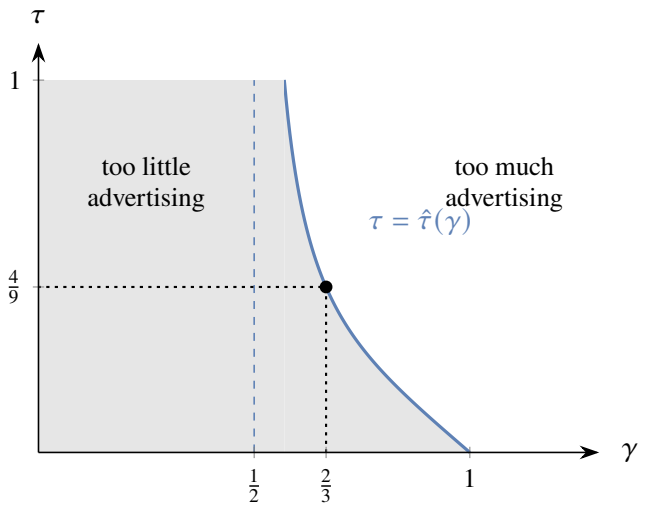


Figure 34.4: Duopoly advertising compared with the first best in the (γ, τ) plane ($\omega = 1$). The shaded region has too little advertising.

close substitutes, viewers move easily, and stations cut advertising, below the optimal level unless the nuisance cost exceeds ω; when programs are very different, stations face little competition and behave like the monopoly of the first step. Advertisers' benefits matter through a^o: a higher ω raises the optimal level but leaves a^* unchanged, so under-advertising is more likely when advertisers value contact with viewers more. The market price of advertising plays the role of a Pigouvian tax on advertisers, but nothing makes it equal to the nuisance cost.

34.6 Fourth Step: The Number of Channels

Social value of a second channel. With full coverage, the optimal advertising level is a^o with one or two channels, so advertisers gain nothing from a second channel. By (34.4) and (34.10), the social gain from the second program is the fall in viewers' distance cost,

$$\Delta B^o = W_2(a^o) - W_1(a^o) = \frac{\tau}{4}, \tag{34.12}$$

so the planner provides two programs if and only if $K \leq \frac{\tau}{4}$. One program is always worth providing in this range, since, under the assumption $\beta \geq \tau + \gamma$,

$$W_1(a^o) \geq \beta - \frac{\tau}{2} \geq \frac{\tau}{2} + \gamma.$$

Market entry. A first station operates if $K \leq \pi^m = \frac{\omega}{4}$, and both operate if $K \leq \pi^*$, where π^* in (34.9) satisfies $\pi^* < \frac{\omega}{8} < \frac{\omega}{4}$ and tends to $\frac{\omega}{8}$ as $\gamma \to 0$. The market over-provides programs when $\pi^* > \frac{\tau}{4}$ and under-provides them when $\pi^* < \frac{\tau}{4}$. Substituting (34.9), $\pi^* > \frac{\tau}{4}$ if and only if

$$2\omega > \sqrt{\gamma^2 + 4\tau^2} + 2\tau,$$

that is,

$$\omega > 2\tau \qquad \text{and} \qquad \gamma < 2\sqrt{\omega(\omega - 2\tau)}. \tag{34.13}$$

Result 4 (Propositions 2 and 5 of the paper). With competing stations:

- if $\omega \leq 2\tau$, the market never provides two programs when one is optimal, and it provides fewer programs than optimal for $K \in (\pi^*, \frac{\tau}{4}]$;
- if $\omega > 2\tau$ and $\gamma < 2\sqrt{\omega(\omega - 2\tau)}$, the market provides two programs when one is optimal, for $K \in (\frac{\tau}{4}, \pi^*]$;
- if $\omega > 2\tau$ and $\gamma > 2\sqrt{\omega(\omega - 2\tau)}$, the market provides fewer programs than optimal for $K \in (\pi^*, \frac{\tau}{4}]$;
- a single owner of both channels operates only one channel for every $K > 0$, while competing stations operate two for $K \leq \pi^*$.

Figure 34.5 depicts π^* and $\frac{\tau}{4}$ as functions of γ for $\omega = 1$ and $\tau = \frac{3}{8}$, so that $\omega > 2\tau$ and the cutoff in (34.13) is $\gamma = 1$. For $\gamma < 1$, every fixed cost between the two curves leads to two programs where one is optimal; for $\gamma > 1$, every fixed cost between them leads to one program where two are optimal. The condition $\omega > 2\tau$ is the paper's comparison of $P(\hat{a}, \frac{N}{2})\hat{a}$, the limit of π^* as $\gamma \to 0$, which equals $\frac{\omega}{8}$ in the toy model, with $\frac{N\tau}{4} = \frac{\tau}{4}$.

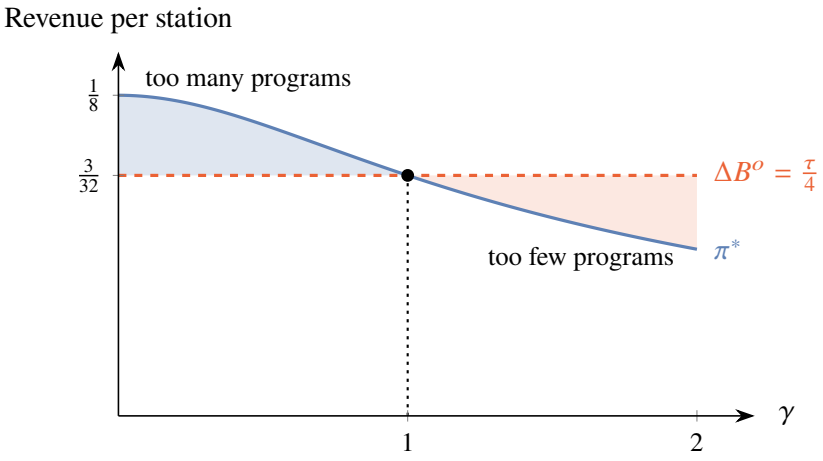


Figure 34.5: Revenue of each competing station and social value of the second program ($\omega = 1$, $\tau = \frac{3}{8}$).

Intuitively, the entrant's revenue misses the viewers' gain from a closer program, which pushes toward too few programs, but it includes revenue taken from the incumbent, which pushes toward too many, as in Mankiw and Whinston (1986). With full coverage, the second channel creates no new contacts for advertisers, so all of its viewers come from the incumbent's audience. When advertising is a small nuisance and advertisers value viewers highly relative to the differentiation of programs, this business stealing is large and the market provides too many programs. A single owner internalizes business stealing completely: since it already reaches every viewer, a second channel adds nothing to its revenue.

34.7 Fifth Step: Advertising Caps

Consider now a regulator that limits advertising to at most $\bar{a}$ per channel, as several European countries do. This step is an application of the toy model; the paper uses such ceilings as motivation. If $\bar{a} < a^*$, each station's profit increases in its own advertising up to $\bar{a}$ when the rival sets $\bar{a}$, so both stations set $\bar{a}$; if $\bar{a} \geq a^*$, the cap has no effect. If $\gamma \geq \omega$, W_2 in (34.10) decreases in a, so every binding cap raises aggregate surplus. If $\gamma < \omega$, W_2 is quadratic with maximum at a^o, so

$$W_2(\bar{a}) - W_2(a^*) = \frac{\omega}{2}\left[(a^* - a^o)^2 - (\bar{a} - a^o)^2\right]. \tag{34.14}$$

Result 5 (advertising caps). With two channels and a cap $\bar{a} < a^*$:

- both stations set $\bar{a}$, and each earns $\frac{\omega}{2}\bar{a}(1-\bar{a}) < \pi^*$;
- if both stations still operate, the cap raises aggregate surplus if and only if $a^o < a^*$ and $2a^o - a^* < \bar{a} < a^*$, with the largest gain at $\bar{a} = a^o$; if $a^* \leq a^o$, every binding cap lowers aggregate surplus;
- with $\omega = 1$, $\gamma = \frac{3}{4}$, and $\tau = \frac{1}{2}$, a cap $\bar{a} = a^o = \frac{1}{4}$ raises aggregate surplus by $\frac{1}{288}$ if $K \leq \frac{3}{32}$, and lowers it by
$$\frac{1}{8} - \frac{1}{288} - K \geq \frac{1}{96}$$
if $\frac{3}{32} < K \leq \frac{1}{9}$, since one station then exits.

Table 34.1 reports the outcomes of the example in the third bullet, with $\beta \geq \frac{5}{4}$ so that every viewer watches. In this example, the duopoly over-advertises, $a^* = \frac{1}{3} > a^o = \frac{1}{4}$, while $\omega = 2\tau$, so the market under-provides programs for $K \in (\frac{1}{9}, \frac{1}{8}]$.

Table 34.1: Example with $\omega = 1$, $\gamma = \frac{3}{4}$, and $\tau = \frac{1}{2}$.

Market structure	Advertising level	Price per viewer	Revenue per station
One channel, the first step	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{4}$
One owner, two channels, the first step	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{8}$
Two competing stations, the second step	$\frac{1}{3}$	$\frac{2}{3}$	$\frac{1}{9}$
Two competing stations, cap $\bar{a} = \frac{1}{4}$, the fifth step	$\frac{1}{4}$	$\frac{3}{4}$	$\frac{3}{32}$
First best, the third step	$\frac{1}{4}$	$\frac{3}{4}$ (Pigouvian)	n.a.

Intuitively, a cap works as a quantity ceiling in the market for viewers' attention. It corrects over-advertising, but it also lowers each station's revenue, since stations already advertise below the level that maximizes revenue per viewer. When entry is driven by advertising revenue, the cap can then cause a station to exit, and the loss of a program can outweigh the gain from fewer advertisements. When the market under-advertises, as in the shaded region of Figure 34.4, a cap only makes matters worse.

34.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- Free programs financed by advertisers, with advertising as a nuisance for viewers and a benefit for advertisers, and stations that compete for viewers only through advertising levels.
- Too little advertising when the nuisance cost is low, programs are close substitutes, or advertisers' benefits are high, and too much otherwise, with a critical nuisance cost that decreases in program differentiation (Proposition 1).
- Too many programs through business stealing when the nuisance cost is small, and too few otherwise (Proposition 2); higher advertising, lower prices, and less programming under single ownership (Proposition 5).

What it drops.

- Viewers who switch off. The assumption $\beta \geq \tau + \gamma$ keeps every viewer watching, which removes the kink in a single station's revenue and the case, studied in the paper, in which a monopoly station shows fewer advertisements than competing stations.
- General advertiser demand with a log-concave distribution, and the choice of program types and duplication.
- Subscription prices (pay-TV), which the paper shows may reduce aggregate surplus (Propositions 3 and 4), the welfare effect of single ownership (Proposition 6), and viewers who switch channels over time (Section 7 and Proposition 7 of the paper).

34.9 Connections to Other Chapters

- Chapter 35 studies platforms that sell targeted ads through auctions and invest in the quality of their free services to attract attention.
- Chapter 33 studies the content of advertising, and how real product information and hype change demand.
- Chapter 32 studies firms that choose which consumers their ads reach.

References

Anderson, S. P., and S. Coate (2005). Market provision of broadcasting: A welfare analysis. *Review of Economic Studies*, 72(4), 947–972.

Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics*, 37(3), 668–691.

Hotelling, H. (1929). Stability in competition. *Economic Journal*, 39(153), 41–57.

Mankiw, N. G., and M. D. Whinston (1986). Free entry and social inefficiency. *RAND Journal of Economics*, 17(1), 48–58.

Part IX

Platforms and Network Effects

The two chapters in this part study firms that use data on consumers to decide how to sell, in markets where what a consumer obtains depends on a platform or on other consumers. In Chapter 35, free platforms attract attention through the quality of their services and sell targeted ads through auctions, and the ads determine which products consumers discover. In Chapter 36, a good is more valuable when the people a consumer interacts with also use it, and firms buy data on how much each consumer influences others and how much each consumer is influenced by them. A monopolist then charges a premium for susceptibility and gives a discount for influence.

Both chapters ask what happens when firms have better data, and in both the answer is that better data need not help the party that uses them. In Chapter 35, more informative data raise product consumption but tend to lower platforms' ad revenue and investment. In Chapter 36, competing firms race to acquire data that lower their profits.

Chapter 35

The Attention Economy

Companion note to Chen, D. (2026), "A Model of the Attention Economy: Data, Platforms, and Discovery," American Economic Review 116(10): 3802–3836.

35.1 Introduction

Social media, streaming services, and podcasts are free to use. They earn their revenue by selling ads, and the ads decide which products consumers discover. How do the platforms' data and their competition for attention affect what consumers buy and how much platforms invest? Chen (2026) answers this question with a model of two linked markets.

In the product market, each consumer buys only from the firms she knows, her *consideration set*. In the attention market, platforms invest in the quality of their free services to attract attention, and sell each ad in an auction in which firms bid using the platforms' data about the consumer. Better data help the best-matched firm win each ad, so consumers discover better products. But better data also widens the gap between the highest and the second-highest bid, and the winner pays only the second, so ad revenue falls and platforms invest less.

The paper shows that the two markets cannot be studied separately. More informative data raises product consumption but tends to lower ad revenue and platform investment. Interoperability mandates, which make attention more sensitive to platform quality, lower product consumption but raise investment. The effects of both policies can change sign over time, and platform investment can be too high or too low relative to the planner's choice. A data or interoperability rule that helps one market can therefore hurt the other.

The toy model fixes the parameters of the paper at simple values and lets the pool of products be large, so that every steady state has a closed form. The first to third steps solve the product market, the attention market, and the ad auctions together with platform investment. The fourth and fifth steps study data policy and interoperability, in steady state and along the path after a policy change. The sixth step compares the equilibrium with the planner's allocation, and the seventh step lists what the toy model keeps and drops.

35.2 The Toy Model

Consumers and products. Time is continuous, $t \geq 0$. A unit mass of consumers has income $I = 1$ per unit of time. There is a large pool of products, each sold by a different firm at marginal cost 1. Consumer i values product j at v_{ij}, drawn independently from the uniform distribution on $[0, 1]$. Consumer i is aware of only a subset of products, her consideration set Ω_{it}, and buys only from it. Over the products in this set she has CES preferences with elasticity of substitution $\sigma = 2$,

$$C_{it} = \left[\int_{\Omega_{it}} v_{ij}^{\frac{1}{2}} \, c_{ijt}^{\frac{1}{2}} \, dj \right]^2. \tag{35.1}$$

Only two features of Ω_{it} matter: its mass M_t and the average value μ_t of the products in it. Both are common across consumers.

Platforms and attention. A unit mass of platforms, $K = 1$, offers free services. Each consumer has one unit of attention per unit of time and divides it across platforms to maximize

$$X_{it} = \left[\int_0^1 \left(v(a_k) \, q_k \, x_{ik} \right)^{\frac{\epsilon-1}{\epsilon}} dk \right]^{\frac{\epsilon}{\epsilon-1}}, \tag{35.2}$$

where x_{ik} is the attention she gives to platform k, q_k is the platform's quality, a_k is its ad rate, and $v(a) = 1 - a$ captures that ads are a nuisance. The elasticity $\epsilon \in (1, 3)$ measures how sensitive attention is to effective quality $v(a_k)q_k$; interoperability mandates raise ϵ. Quality depreciates at rate $\delta = 1$ and grows with investment ℓ_k,

$$\dot{q}_k = \sqrt{\ell_k} - q_k, \tag{35.3}$$

so investment has decreasing returns, with the paper's parameter $\varphi = \frac{1}{2}$.

Ads and discovery. A consumer who gives attention x_{ik} to platform k sees ads at rate $a_k x_{ik}$. Each ad is sold the instant it arises, in a second-price auction among $N = 5$ firms drawn at random from the products outside her consideration set. The winner enters the set, and every product in the set is forgotten at rate $\lambda_f = 1$. All firms and platforms discount the future at rate $\rho = 1$.

Data. Firms do not observe v_{ij}. The platforms' data give each bidder an expected value $\hat{v}_{ij}$, drawn from the uniform distribution G on $\left[\frac{1}{2} - s, \frac{1}{2} + s\right]$ with $s \in \left[0, \frac{1}{2}\right]$. With $s = 0$, data are uninformative and every firm expects the mean value $\frac{1}{2}$; with $s = \frac{1}{2}$, data are perfectly informative and $\hat{v}_{ij} = v_{ij}$. A higher s is a mean-preserving spread of G, that is, more informative data.

A large pool of products. The paper has a finite mass J of products. We let $J \to \infty$, which has two implications. First, the products outside a consideration set are a representative sample, so the five bidders in each auction have expected values drawn independently from G. Second, a product that a consumer forgets is invited to bid for her again only at rate $\frac{NA}{J} \approx 0$, where A is the rate at which she sees ads, so the firm treats being forgotten as permanent. The seventh step discusses what this simplification drops.

35.3 First Step: The Product Market

Maximizing (35.1) subject to the budget constraint gives consumer i's demand for product $j \in \Omega_{it}$,

$$c_{ij} = \frac{v_{ij}\, p_j^{-2}}{\int_{\Omega_{it}} v_{iz}\, p_z^{-1}\, dz}. \tag{35.4}$$

The value v_{ij} and the consideration set only scale demand, so firm j maximizes $(p_j - 1)p_j^{-2}$ regardless of what it knows about the consumer. The first-order condition,

$$p_j^{-2} - 2(p_j - 1)p_j^{-3} = 0,$$

gives $p_j = 2 = \frac{\sigma}{\sigma - 1}$. Substituting $p = 2$ in (35.4), $c_{ij} = \frac{v_{ij}}{2M_t\mu_t}$, so the firm's flow profit from consumer i and the consumer's aggregate consumption are

$$(p - 1)\, c_{ij} = \frac{v_{ij}}{2M_t\mu_t}, \qquad C_{it} = \frac{1}{2}\, M_t\mu_t. \tag{35.5}$$

Result 1 (product market, Proposition 1 of the paper). Every firm charges a price of 2, regardless of what it knows about consumers. A firm in a consideration set earns $\frac{v}{2M_t\mu_t}$ per unit of time from that consumer, and the firms in the set earn a total of $\frac{1}{2} = \frac{I}{\sigma}$, regardless of the size and composition of the set. Aggregate consumption $\frac{1}{2}M_t\mu_t$ increases in both the mass and the average value of the products the consumer knows.

Intuitively, a consumer who knows more products, or better ones, consumes more, but firms' total profit stays the same: each firm faces more, or better matched, rivals for the consumer's budget. This fixed profit pie is what the attention market competes for.

35.4 Second Step: Attention, Ad Load, and Consideration Sets

Attention. Maximizing (35.2) subject to $\int_0^1 x_{ik}\, dk \le 1$ gives the attention share

$$x_k = \frac{[v(a_k)q_k]^{\epsilon-1}}{\int_0^1 [v(a_z)q_z]^{\epsilon-1}\, dz}. \tag{35.6}$$

Ad rate. Let π_K denote the expected price of an ad. Platform k earns $\pi_K a_k x_k$ per unit of time, and since it is small, it takes the denominator of (35.6) as given and chooses a_k to maximize $a(1 - a)^{\epsilon-1}$. The first-order condition, $(1 - a) - (\epsilon - 1)a = 0$, gives the ad rate

$$A = \frac{1}{\epsilon}. \tag{35.7}$$

All platforms choose the same ad rate and, in a symmetric equilibrium, the same quality, so each receives attention $x_k = 1$, and consumers see ads at rate A.

Consideration sets. Products enter a consideration set at rate A and leave at rate $\lambda_f M_t = M_t$, so $\dot{M}_t = A - M_t$ and M_t converges to $M = A$. The winner of each auction is the bidder with the highest expected value, the maximum of five draws from G, and products leave at random, so in steady state the average value in the set is the expected maximum of five draws. For the uniform distribution on $\left[\frac{1}{2} - s, \frac{1}{2} + s\right]$, the expected highest and second-highest of five draws are

$$\mu_H = E[\hat{v}_{(1)}] = \frac{1}{2} + \frac{2s}{3}, \qquad E[\hat{v}_{(2)}] = \frac{1}{2} + \frac{s}{3}. \tag{35.8}$$

Result 2 (ad load and discovery, Theorem 1 of the paper). In steady state:

- every platform shows ads at rate $A = \frac{1}{\epsilon}$, which decreases in the elasticity of attention;
- consideration sets have mass $M = \frac{1}{\epsilon}$ and average value $\mu_H = \frac{1}{2} + \frac{2s}{3}$, which increases in data informativeness.

Interoperability therefore changes how many products consumers discover, and data change which ones. By (35.5), steady-state consumption is $C = \frac{\mu_H}{2\epsilon}$.

35.5 Third Step: Ad Auctions and Platform Investment

Bids. In a second-price auction, a firm bids the value of joining the consideration set. A member with expected value $\hat{v}$ earns $\frac{\hat{v}}{2M\mu_H}$ per unit of time until it is forgotten, and a firm that loses earns nothing from this consumer, so, discounting at $\rho + \lambda_f = 2$,

$$B(\hat{v}) = \frac{\hat{v}}{4M\mu_H}. \tag{35.9}$$

Ad revenue. The winner pays the second-highest bid, so the expected ad price is

$$\pi_K = \frac{E[\hat{v}_{(2)}]}{4M\mu_H}.$$

Total ad revenue is $R = \pi_K A$ and, since $M = A$,

$$R = \underbrace{\frac{I}{\sigma}}_{=\frac{1}{2}} \cdot \underbrace{\frac{\lambda_f}{\rho + \lambda_f}}_{=\frac{1}{2}} \cdot \frac{E[\hat{v}_{(2)}]}{E[\hat{v}_{(1)}]} = \frac{1}{4} \cdot \frac{3 + 2s}{3 + 4s}. \tag{35.10}$$

This is the toy-model version of equation (17) of the paper. Ad revenue is a share of the fixed profit pie $\frac{I}{\sigma}$, and the share falls with the ratio of the second-highest to the highest expected value in an auction, which measures the winner's information rent.

Investment. Platform k chooses ℓ_k to maximize

$$\int_0^\infty e^{-t} \left(\pi_K A x_k - \ell_k\right) dt$$

subject to (35.3). By (35.6), a higher q_k raises the platform's revenue at rate $\frac{(\epsilon-1)R}{q_k}$ in a symmetric equilibrium, so in steady state the shadow value of quality is

$$\lambda = \frac{(\epsilon-1)R}{(\rho+\delta)q} = \frac{(\epsilon-1)R}{2q}.$$

The first-order condition for investment, $1 = \frac{\lambda}{2\sqrt{\ell}}$, together with $q = \sqrt{\ell}$ from (35.3), gives

$$\ell = \frac{(\epsilon-1)R}{4}, \qquad q = \sqrt{\ell}. \tag{35.11}$$

Result 3 (steady state, Theorem 1 of the paper). In steady state:

- ad revenue is
$$R = \frac{1}{4} \cdot \frac{3+2s}{3+4s},$$
which decreases in data informativeness and does not depend on the elasticity of attention;
- platform investment is $\ell = \frac{(\epsilon-1)R}{4}$, and quality is $q = \sqrt{\ell}$;
- product consumption is
$$C = \frac{1}{2\epsilon}\left(\frac{1}{2} + \frac{2s}{3}\right).$$

The two markets are linked in both directions. Platform investment depends on ad revenue, which depends on how firms compete in the product market, and product consumption depends on the ad rate and the data that platforms supply. The paper requires $\epsilon - 1 < \frac{1}{\varphi} = 2$ for the platform's problem to be concave, which explains the restriction $\epsilon < 3$.

35.6 Fourth Step: Data Policy

Steady state. Set $\epsilon = 2$, so $A = M = \frac{1}{2}$, and compare three levels of data informativeness. From (35.8)–(35.11):

- uninformative data, $s = 0$: $\mu_H = E[\hat{v}_{(2)}] = \frac{1}{2}$, $R = \frac{1}{4}$, $C = \frac{1}{8}$, and $\ell = \frac{1}{16}$;
- intermediate data, $s = \frac{3}{10}$: $\mu_H = \frac{7}{10}$, $E[\hat{v}_{(2)}] = \frac{3}{5}$, $R = \frac{3}{14}$, $C = \frac{7}{40}$, and $\ell = \frac{3}{56}$;
- perfectly informative data, $s = \frac{1}{2}$: $\mu_H = \frac{5}{6}$, $E[\hat{v}_{(2)}] = \frac{2}{3}$, $R = \frac{1}{5}$, $C = \frac{5}{24}$, and $\ell = \frac{1}{20}$.

Result 4 (data, Proposition 2 of the paper). In steady state, more informative data:

- raise product consumption, $C = \frac{1}{8} + \frac{s}{6}$;
- lower ad revenue, R in (35.10), which reaches its maximum,
$$\frac{I}{\sigma} \cdot \frac{\lambda_f}{\rho+\lambda_f} = \frac{1}{4},$$
with uninformative data;

- lower platform investment and quality.

Figure 35.1 depicts steady-state ad revenue and product consumption as functions of data informativeness s, with $\epsilon = 2$. Consumption increases from $\frac{1}{8}$ to $\frac{5}{24}$, while ad revenue decreases from $\frac{1}{4}$ to $\frac{1}{5}$. The dots mark the policy experiment below, from $s = \frac{3}{10}$ to $s = \frac{1}{2}$.

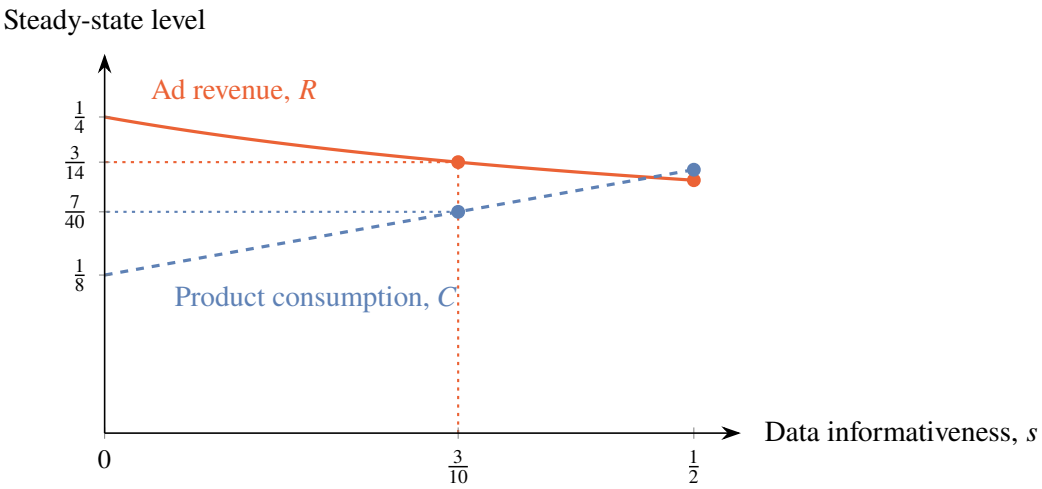


Figure 35.1: Steady-state ad revenue and product consumption as functions of data informativeness ($\epsilon = 2$).

Intuitively, three forces drive the decline in ad revenue, as in the paper. First, by Result 1, firms' total profits are fixed at $\frac{1}{2}$ regardless of the data. Second, with uninformative data all five bidders expect the same value and bid the same, so competition in the auction extracts the full value of joining a consideration set. Third, informative data create information rents: the winner pays the second-highest bid, and the gap between the highest and second-highest expected values, $\frac{s}{3}$, increases in s. Better targeting does not raise the winner's profit either: the winner's expected value rises, but so does the average value μ_H of its rivals in the consideration set, and in steady state the two coincide.

Transition. Consider an unanticipated improvement in data at $t = 0$, from $s = \frac{3}{10}$ to $s = \frac{1}{2}$, as in Figure 2 of the paper. The auctions use the new data at once, so the second-highest expected value of the bidders jumps from $\frac{3}{5}$ to $\frac{2}{3}$. The composition of consideration sets, in contrast, changes only as products are forgotten and replaced. The products already in a set receive new expected values, but by the law of iterated expectations their average stays at $\frac{7}{10}$, and from then on

$$\mu_t = \frac{5}{6} - \frac{2}{15}\, e^{-t}.$$

A firm that joins at time t earns $\frac{\hat{v}}{2M\mu_\tau}$ at every later date τ, so its bid is

$$\hat{v} \int_0^\infty \frac{e^{-2u}}{\mu_{t+u}}\, du,$$

and ad revenue on impact is

$$R_0 = \frac{1}{2} \cdot \frac{2}{3} \int_0^\infty \frac{e^{-2u}\, du}{\frac{5}{6} - \frac{2}{15} e^{-u}} = \frac{75}{4} \left[\frac{5}{6} \ln \frac{25}{21} - \frac{2}{15} \right] \approx 0.224, \tag{35.12}$$

which exceeds the initial level $\frac{3}{14} \approx 0.214$ by about 5%. Ad revenue then declines toward $\frac{1}{5}$, about 7% below its initial level, and product consumption rises gradually from $\frac{7}{40}$ to $\frac{5}{24}$.

Figure 35.3 depicts this path (red curve), relative to the initial level. Ad revenue rises on impact, crosses its initial level at about $t = 0.48$, and settles below it. Firms bid more at first because they value the better targeting before their rivals in consideration sets become better matched. As in the paper, the rise on impact is larger when firms are less patient: it is about 2% with $\rho = 0.1$ and about 9% with $\rho = 10$. Steady-state investment falls from $\frac{3}{56}$ to $\frac{1}{20}$. With $\rho = 1$, platforms are impatient enough that investment first rises slightly, to about 0.055, before falling; in the paper, with more patient platforms, investment falls from the start.

35.7 Fifth Step: Interoperability

Steady state. Set $s = \frac{1}{2}$ and raise the elasticity of attention from $\epsilon = \frac{3}{2}$ to $\epsilon = 2$. By Results 2 and 3:

- the ad rate and the mass of consideration sets fall from $\frac{2}{3}$ to $\frac{1}{2}$;
- product consumption, $C = \frac{5}{12\epsilon}$, falls from $\frac{5}{18}$ to $\frac{5}{24}$;
- ad revenue stays at $R = \frac{1}{5}$;
- investment, $\ell = \frac{\epsilon - 1}{20}$, doubles from $\frac{1}{40}$ to $\frac{1}{20}$.

Result 5 (interoperability, Proposition 3 of the paper). In steady state, a higher elasticity of attention:

- lowers the ad rate and the mass of consideration sets, and hence product consumption;
- leaves ad revenue unchanged in the toy model (in the paper it raises ad revenue, under a sufficient condition on G that the uniform distribution on $[0, 1]$, the case $s = \frac{1}{2}$, satisfies);
- raises platform investment and quality.

Figure 35.2 depicts steady-state product consumption, ad revenue, and investment as functions of ϵ, with $s = \frac{1}{2}$. Consumption decreases and investment increases with ϵ, while ad revenue is flat.

Intuitively, when attention is more elastic, ads drive users away faster, so platforms show fewer ads and compete for attention through quality instead. Fewer ads mean that consumers discover fewer products. Each ad, however, becomes more valuable: by (35.9), a smaller consideration set means fewer rivals for each consumer, which raises bids. With $M = A$, the two effects cancel out. In the paper, fewer ads also reduce the rate at which forgotten firms are rediscovered, which raises bids further, so ad revenue increases; the large pool of products in the toy model shuts down this channel.

Transition. Consider an unanticipated increase in ϵ from $\frac{3}{2}$ to 2 at $t = 0$. Platforms cut the ad rate to $\frac{1}{2}$ at once, but consideration sets shrink only gradually,

$$M_t = \frac{1}{2} + \frac{1}{6} e^{-t},$$

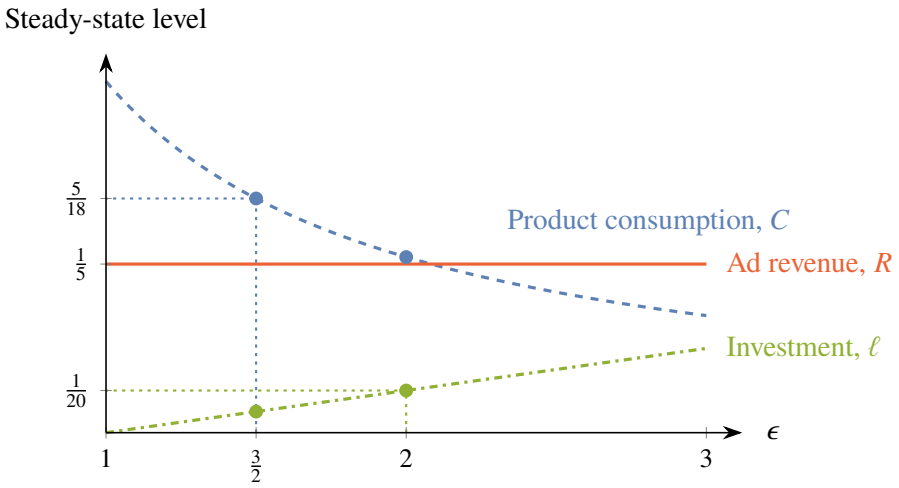


Figure 35.2: Steady-state product consumption, ad revenue, and platform investment as functions of the elasticity of attention ($s = \frac{1}{2}$).

so bids rise only gradually. A firm that joins at time t bids

$$\hat{v} \int_0^\infty \frac{e^{-2u}}{2M_{t+u}\mu_H}\, du,$$

and since

$$\frac{A\, E[\hat{v}_{(2)}]}{2\mu_H} = \frac{\frac{1}{2}\cdot\frac{2}{3}}{\frac{5}{3}} = \frac{1}{5},$$

ad revenue at time t is

$$R_t = \frac{1}{5}\int_0^\infty \frac{e^{-2u}}{M_{t+u}}\, du.$$

On impact,

$$R_0 = \frac{1}{5}\int_0^\infty \frac{e^{-2u}\, du}{\frac{1}{2}+\frac{1}{6}e^{-u}} = \frac{36}{5}\left[\frac{1}{6} - \frac{1}{2}\ln\frac{4}{3}\right] \approx 0.164, \tag{35.13}$$

about 18% below the steady-state level $\frac{1}{5}$. Figure 35.3 depicts this path (blue curve): ad revenue falls on impact and recovers as consideration sets shrink. Product consumption falls gradually from $\frac{5}{18}$ to $\frac{5}{24}$. In the toy model, investment rises on impact, to about 0.058, and then declines toward $\frac{1}{20}$: the elasticity $\epsilon - 1$ doubles, which outweighs the dip in ad revenue. In the paper, with steeper nuisance costs, the ad rate falls more sharply on impact, and investment first falls before rising in the long run.

35.8 Sixth Step: Welfare

General equilibrium. To evaluate welfare, the paper closes the model. Consumers supply $L = 1$ unit of labor, the only input: platforms hire ℓ units for investment, and firms use the rest for production, one unit of labor per unit of output. Consumers own all firms and platforms, and their

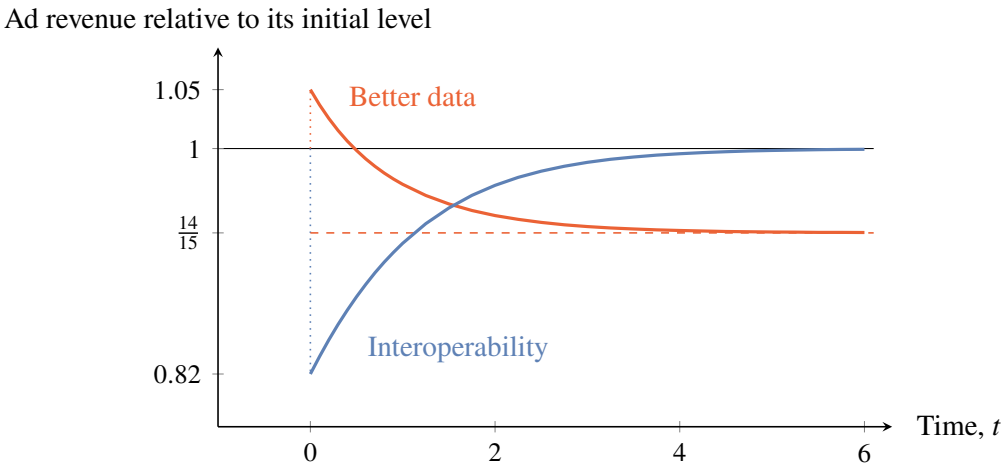


Figure 35.3: Ad revenue after an unanticipated policy change at $t = 0$, relative to its initial level: better data, from $s = \frac{3}{10}$ to $s = \frac{1}{2}$ with $\epsilon = 2$ (red), and interoperability, from $\epsilon = \frac{3}{2}$ to $\epsilon = 2$ with $s = \frac{1}{2}$ (blue). The dashed line is the long-run level after better data.

income equals firms' revenue, $I = 2(1-\ell)$, the price times production. Ad revenue per unit of income,

$$\hat{R} = \frac{R}{I} = \frac{1}{4} \cdot \frac{3+2s}{3+4s},$$

does not depend on I, so (35.11) becomes

$$\ell = \frac{(\epsilon-1)\hat{R}}{4} 2(1-\ell),$$

and

$$\ell = \frac{x}{4+x}, \qquad \textit{where} \quad x \equiv \frac{\sigma}{\sigma-1}\hat{R}(\epsilon-1) = \frac{(3+2s)(\epsilon-1)}{2(3+4s)}. \tag{35.14}$$

Planner. Consumers' flow utility is $C^{1-\tau}X^{\tau}$, with $\tau \in (0,1)$, where $C = (1-\ell)M\mu_H$ and $X = \nu(a)q$. The planner chooses investment and the ad rate, while consumers and firms behave as in the equilibrium. In steady state, the shadow value of quality is $\frac{\tau u}{2q}$, and investing one more unit of labor costs $\frac{(1-\tau)u}{1-\ell}$ in utility and raises quality at rate $\frac{1}{2\sqrt{\ell}}$, so the planner's investment satisfies

$$\frac{(1-\tau)u}{1-\ell} = \frac{\tau u}{4\ell},$$

that is,

$$\ell^* = \frac{\frac{\tau}{1-\tau}}{4+\frac{\tau}{1-\tau}} = \frac{\tau}{4-3\tau}. \tag{35.15}$$

For the ad rate, the paper compares steady states (its limit $\rho \to 0$). Since $M = a$ and μ_H does not depend on a, the planner maximizes $a^{1-\tau}(1-a)^{\tau}$, which gives $a^* = 1-\tau$.

Result 6 (welfare, Proposition 5 and Theorem 2 of the paper). Comparing the equilibrium with the planner's allocation:

- platforms invest too much if $x > \frac{\tau}{1-\tau}$, and too little if $x < \frac{\tau}{1-\tau}$, where x is the sufficient statistic in (35.14);
- platforms show too many ads if $\frac{1}{\epsilon} > 1 - \tau$, that is, if $\tau > 1 - \frac{1}{\epsilon}$, and too few ads otherwise;
- a tax on ad revenue at rate $1 - \frac{\tau}{(1-\tau)x}$, rebated to consumers, restores efficient investment (a negative rate is a subsidy).

Figure 35.4 depicts these regions in the (ϵ, τ) plane with perfectly informative data, $s = \frac{1}{2}$, where $x = \frac{2(\epsilon-1)}{5}$. Investment is too high below the blue curve, $\tau = \frac{2(\epsilon-1)}{2\epsilon+3}$, and the ad rate is too high above the red curve, $\tau = 1 - \frac{1}{\epsilon}$. The dot marks $\epsilon = 2$ and $\tau = \frac{1}{4}$. There, $x = \frac{2}{5}$, so equilibrium investment is $\frac{1}{11}$, above the planner's $\frac{1}{13}$, and a tax of $\frac{1}{6}$ on ad revenue restores efficiency. The ad rate, $\frac{1}{2}$, is below the planner's $\frac{3}{4}$.

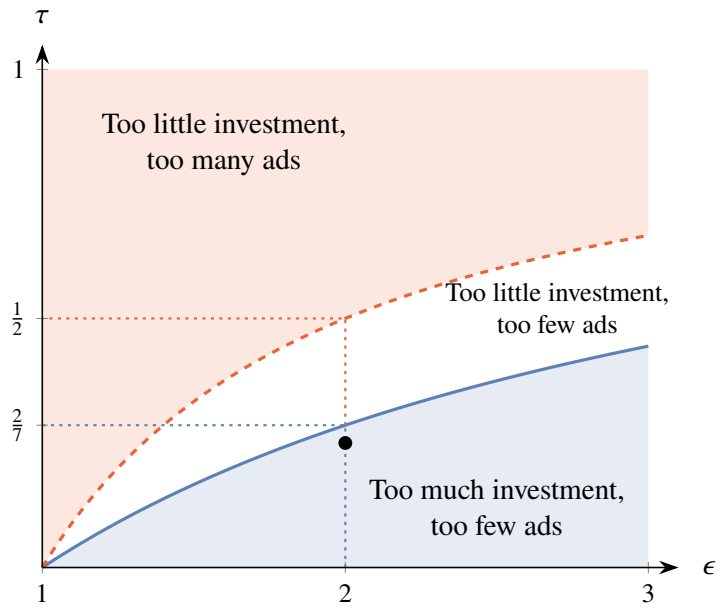


Figure 35.4: Investment and ad distortions in the (ϵ, τ) plane with $s = \frac{1}{2}$. Investment is too high in the blue region and the ad rate is too high in the red region.

Intuitively, the two sides of the investment condition reflect different distortions. The statistic x is the product of the markup $\frac{\sigma}{\sigma-1} = 2$, ad revenue per unit of income, and the elasticity $\epsilon - 1$. Ad revenue and $\epsilon - 1$ measure business stealing: platforms invest to take attention, and ad revenue, from their rivals, as in Anderson and Coate (2005). The markup reflects product-market power, which shifts labor toward platform investment. The term $\frac{\tau}{1-\tau}$ measures the value of quality to consumers, which platforms do not collect because their services are free. The ad rate is too high when attention is inelastic, since platforms then care little about the nuisance cost of ads, and too low when consumers value products far more than platform services, since platforms do not collect the value of the products that ads help consumers discover.

The investment condition also tells us how to evaluate the policies of the fourth and fifth steps. With uninformative data, $x = \frac{1}{2}$ at $\epsilon = 2$, so equilibrium investment rises to $\frac{1}{9}$, further above the planner's $\frac{1}{13}$, and the corrective tax rises to $\frac{1}{3}$. Hence, when investment is already too high, a ban on data

use moves investment further from the efficient level, as the paper notes. Similarly, interoperability raises x but not the planner's investment ℓ^*, so it can move investment toward the efficient level only when platforms initially invest too little.

35.9 Seventh Step: What the Toy Model Keeps and Drops

Table 35.1 reports the effects of the two policy experiments of the fourth and fifth steps.

Table 35.1: Effects of better data (s from $\frac{3}{10}$ to $\frac{1}{2}$, $\epsilon = 2$) and interoperability (ϵ from $\frac{3}{2}$ to 2, $s = \frac{1}{2}$).

Policy	Ad revenue, impact	Ad revenue, long run	Consumption	Investment, long run
Better data	+5%	−7%	$\frac{7}{40}$ to $\frac{5}{24}$ (+19%)	$\frac{3}{56}$ to $\frac{1}{20}$ (−7%)
Interoperability	−18%	0%	$\frac{5}{18}$ to $\frac{5}{24}$ (−25%)	$\frac{1}{40}$ to $\frac{1}{20}$ (+100%)

What the toy model keeps.

- The link between the two markets: ads determine which products consumers discover, and product-market competition determines how much firms pay for ads.
- The fixed profit pie under CES demand, which explains why better targeting does not raise firms' total willingness to pay for ads.
- Information rents in second-price auctions, which make ad revenue highest with uninformative data.
- The trade-off between ad load and quality when attention becomes more elastic.
- Short-run effects that differ from long-run effects, because consideration sets adjust gradually while auctions and ad rates respond at once.
- The two welfare distortions, business stealing and unpriced consumer surplus, and the sufficient statistic that signs the investment distortion.

What it drops.

- A finite pool of products. In the paper, forgotten firms can be rediscovered, and fewer ads slow down rediscovery, which raises bids; this channel makes interoperability raise long-run ad revenue.
- Strategic investment by a duopoly of large platforms, which the paper solves numerically and finds to preserve the effects of both policies.
- Data asymmetry between data-rich digital platforms and traditional media. The paper shows that data-rich platforms can gain a larger share of ad revenue while total ad revenue falls, consistent with the decline in advertising's share of GDP during 2000–2018 documented by Silk and Berndt (2021).
- Network effects, platforms with different productivity, prices for platform services, and entry of firms and platforms, which the paper studies in its online appendix.

35.10 Connections to Other Chapters

- Chapter 32 studies firms that target their ads to loyal consumers and comparison shoppers, and shows that targeting raises profits.
- Chapter 36 studies firms that buy data on consumers. With competition, firms race to acquire data that lower their profits.
- Chapter 4 studies firms that collect data on their customers' preferences and can sell them to a rival.
- Chapter 34 studies free media financed by advertising, where viewers dislike ads and the market can deliver too many or too few of them.

References

Anderson, S. P., and Coate, S. (2005). Market provision of broadcasting: A welfare analysis. *Review of Economic Studies*, 72(4), 947–972.

Chen, D. (2026). A model of the attention economy: Data, platforms, and discovery. *American Economic Review*, 116(10), 3802–3836.

Silk, A. J., and Berndt, E. R. (2021). Aggregate advertising expenditure in the US economy: Measurement and growth issues in the digital era. *Foundations and Trends in Marketing*, 15(1), 1–85.

Chapter 36

Pricing Network Effects: Influencers and Followers

Companion note to Fainmesser, I. P., and A. Galeotti (2016), "Pricing Network Effects," Review of Economic Studies 83(1): 165–198, and to its competition extension, Fainmesser and Galeotti (2020), American Economic Journal: Microeconomics 12(3): 1–32.

36.1 Introduction

Many goods are worth more when the people we interact with also use them, such as a messaging app, an online game, or a fashion item. Firms can now buy data on who listens to whom, and they can use the data to set prices. Which consumers should a firm charge more, and who gains from these data? Fainmesser and Galeotti (2016) answer for a monopolist, and Fainmesser and Galeotti (2020) for two competing firms.

Two numbers describe a consumer's place in the network. Her *susceptibility* is how many others she listens to, and her *influence* is how many others listen to her. A susceptible consumer values the good more, because she enjoys the consumption of the people she follows, so the monopolist charges her a premium. An influential consumer does not value the good more herself, but when she buys more, the people who follow her value it more, so the monopolist gives her a discount.

Pricing on influence raises consumer surplus, while pricing on susceptibility can lower it. With competing firms, data on influence lead to targeted discounts for influencers. A consumer whom only one firm targets may then buy the product she likes less, a loss called *misallocation*. Firms also race to acquire data that lower their profits. The same network data can therefore help or hurt consumers, depending on what they reveal and on whether one firm or two compete for consumers.

The toy model has four equally likely types of consumers (isolated consumers, influencers, followers, and hubs), so every price has a closed form. The first step solves the consumption stage and the second step the uniform price. The third and fourth steps solve pricing on susceptibility and on influence, and the fifth step compares consumer surplus and the value of network data. The sixth

to eighth steps study competing firms, misallocation, and the race for data, and the ninth step lists what the toy model keeps and drops.

36.2 The Toy Model

Network. A monopolist produces a divisible good at zero cost and sells it to a unit mass of consumers. Each consumer interacts with others, and two numbers describe her position in the network. Her *susceptibility* k is the number of other consumers whose consumption affects her (her out-degree), and her *influence* l is the number of other consumers whose consumption she affects (her in-degree). In the toy model, k and l are independent, and each equals 0 or 2 with probability $\frac{1}{2}$. There are thus four equally likely types:

- isolated consumers, $(k, l) = (0, 0)$;
- influencers, $(k, l) = (0, 2)$, whom others follow but who follow no one;
- followers, $(k, l) = (2, 0)$, who follow two others but whom no one follows;
- hubs, $(k, l) = (2, 2)$, who follow two others and are followed by two others.

Average susceptibility and average influence both equal $\bar{k} = 1$, since every link has two ends.

Preferences. A consumer of susceptibility k who buys x units at a unit price p obtains utility

$$U = (1 - p)\, x - \frac{1}{2}x^2 + \gamma\, k\, x\, A, \tag{36.1}$$

where A is the expected consumption of a consumer with whom she interacts, and $\gamma \in \left(0, \frac{1}{4}\right)$ measures the strength of network effects. The upper bound on γ is the paper's Assumption 1, $\gamma k_{\max} < \frac{1}{2}$, which guarantees that the monopolist's problem has a unique interior solution.

Who are the neighbors? A consumer samples the people she interacts with in proportion to their influence: a consumer with influence l is l times as likely to be sampled as a consumer with influence 1. In the toy model, consumers with $l = 0$ are never sampled, so A is the average consumption of the consumers with influence 2, half of whom are influencers and half hubs.

Information. The monopolist knows the distribution of types. We compare three cases: it charges a uniform price; it observes susceptibility and charges $p(k)$; or it observes influence and charges $p(l)$. Consumers observe the prices and then choose their consumption simultaneously.

36.3 First Step: The Consumption Stage

Utility (36.1) is concave in x, and the first-order condition gives the demand of a consumer with susceptibility k who pays p,

$$x = 1 - p + \gamma\, k\, A. \tag{36.2}$$

Demand increases in susceptibility, since a consumer who follows more people enjoys more of their consumption. Her influence l does not affect her own demand. Influence matters only through A, which must equal the average consumption of the consumers with $l = 2$,

$$A = \frac{1}{2}\big[x(0,2) + x(2,2)\big]. \tag{36.3}$$

Substituting (36.2) into (36.3) gives a linear equation in A with slope $\gamma < 1$, so it has a unique solution for any prices. This is the paper's Proposition 1. Finally, at the optimal x, utility (36.1) equals $U = x \cdot x - \frac{1}{2}x^2$, that is,

$$U = \frac{1}{2}x^2. \tag{36.4}$$

A consumer is better off whenever she consumes more, which we use in the fifth step.

36.4 Second Step: The Uniform Price

With a uniform price p, consumers with $k = 0$ buy $1 - p$ and consumers with $k = 2$ buy $1 - p + 2\gamma A$. By (36.3),

$$A = \frac{1}{2}\big[(1 - p) + (1 - p + 2\gamma A)\big] = (1 - p) + \gamma A,$$

so

$$A = \frac{1 - p}{1 - \gamma}. \tag{36.5}$$

Average consumption is

$$(1 - p) + \gamma A = \frac{1 - p}{1 - \gamma},$$

so the monopolist maximizes $\Pi = \frac{p(1-p)}{1-\gamma}$.

Result 1 (uniform price, Proposition 2 of the paper). The monopolist charges $p = \frac{1}{2}$, regardless of the strength of network effects. Then:

- consumers with $k = 0$ buy $\frac{1}{2}$, and consumers with $k = 2$ buy $\frac{1+\gamma}{2(1-\gamma)}$;
- expected neighbor consumption is $A = \frac{1}{2(1-\gamma)}$;
- profit is $\Pi^U = \frac{1}{4(1-\gamma)}$, which increases in γ.

Intuitively, network effects scale up total demand by the factor $\frac{1}{1-\gamma}$ at every price, so the monopoly price does not change. Profit rises with γ, but a uniform price cannot exploit the fact that some consumers matter more for others' demand.

36.5 Third Step: Pricing Susceptibility

Consider that the monopolist observes susceptibility and charges p_0 to consumers with $k = 0$ and p_2 to consumers with $k = 2$. Since half of the consumers with influence 2 have $k = 0$ and half have

$k = 2$, condition (36.3) gives $A = \frac{1-\bar{p}}{1-\gamma}$, where $\bar{p} = \frac{p_0+p_2}{2}$, and profit is

$$\Pi = \frac{1}{2}\, p_0(1 - p_0) + \frac{1}{2}\, p_2(1 - p_2 + 2\gamma A). \tag{36.6}$$

Raising either price lowers A at rate $\frac{1}{2(1-\gamma)}$, which lowers the demand of the consumers with $k = 2$. The first-order conditions are therefore

$$1 - 2p_0 - \frac{\gamma\, p_2}{1-\gamma} = 0, \qquad 1 - 2p_2 + 2\gamma A - \frac{\gamma\, p_2}{1-\gamma} = 0, \tag{36.7}$$

which we can write as

$$p(k) = \underbrace{\frac{1}{2}}_{\text{uniform price}} + \underbrace{\frac{\gamma A}{2}\, k}_{\text{susceptibility premium}} - \underbrace{\frac{\gamma\, p_2}{2(1-\gamma)}}_{\text{common discount}}. \tag{36.8}$$

Solving (36.7) and letting

$$D \equiv 4 - 4\gamma - \gamma^2$$

denote a common denominator gives the following result.

Result 2 (pricing susceptibility, Propositions 3–4 of the paper). When the monopolist prices on susceptibility:

- it charges

  $$p_0 = \frac{2 - 3\gamma - \gamma^2}{D} \qquad \text{and} \qquad p_2 = \frac{(1-\gamma)(2+\gamma)}{D},$$

  below and above the uniform price, respectively, so the premium per link is $\frac{\gamma A}{2} = \frac{\gamma}{D}$ and the average price is below $\frac{1}{2}$;
- consumers with $k = 0$ buy $\frac{2-\gamma}{D}$ and consumers with $k = 2$ buy $\frac{2+\gamma}{D}$, so more susceptible consumers still consume more;
- expected neighbor consumption is $A = \frac{2}{D}$, and profit is $\Pi^S = \frac{1}{D}$.

The last two terms of (36.8) reflect two forces. The premium is standard price discrimination: susceptible consumers have higher demand because they enjoy more network effects, so they pay more. The discount reflects influence: a higher price, to anyone, lowers the consumption that followers observe, which lowers their demand. The discount is the same for every consumer because, without data on influence, the monopolist cannot tell who is influential.

36.6 Fourth Step: Pricing Influence

Consider instead that the monopolist observes influence and charges q_0 to consumers with $l = 0$ and q_2 to consumers with $l = 2$. Now (36.3) gives $A = \frac{1-q_2}{1-\gamma}$: only the price of influential consumers affects A. Average susceptibility within each influence group is 1, so profit is

$$\Pi = \frac{1}{2}\, q_0(1 - q_0 + \gamma A) + \frac{1}{2}\, q_2(1 - q_2 + \gamma A), \tag{36.9}$$

and the first-order conditions are

$$1 - 2q_0 + \gamma A = 0, \qquad 1 - 2q_2 + \gamma A - \frac{\gamma\,(q_0 + q_2)}{1-\gamma} = 0. \tag{36.10}$$

With $\bar{q} = \frac{q_0+q_2}{2}$, both conditions are summarized by

$$q(l) = \underbrace{\frac{1}{2}}_{\text{uniform price}} + \underbrace{\frac{\gamma A}{2}\,\bar{k}}_{\text{common premium}} - \underbrace{\frac{\gamma\,\bar{q}}{2(1-\gamma)}\,l}_{\text{influence discount}}. \tag{36.11}$$

Result 3 (pricing influence, Propositions 3–4 of the paper). When the monopolist prices on influence:

- it charges
$$q_0 = \frac{2-\gamma}{D} \qquad \text{and} \qquad q_2 = \frac{2-3\gamma}{D},$$
above and below the uniform price, respectively, so the discount for influence is $q_0 - q_2 = \frac{2\gamma}{D}$ and the average price is above $\frac{1}{2}$;
- consumption is $\frac{2-3\gamma-\gamma^2}{D}$ for isolated consumers, $\frac{(1-\gamma)(2+\gamma)}{D}$ for influencers, $\frac{2+\gamma+\gamma^2}{D}$ for followers, and $\frac{(1+\gamma)(2+\gamma)}{D}$ for hubs;
- expected neighbor consumption is $A = \frac{2+\gamma}{D}$, and profit is $\Pi^I = \frac{1}{D}$.

Figure 36.1 depicts the three price schedules with $\gamma = \frac{1}{5}$, where $D = \frac{79}{25}$. Pricing susceptibility gives an increasing schedule, from $\frac{34}{79}$ to $\frac{44}{79}$, while pricing influence gives a decreasing one, from $\frac{45}{79}$ to $\frac{35}{79}$. Pricing susceptibility charges more to the consumers with higher demand, while pricing influence charges less to the consumers who raise others' demand.

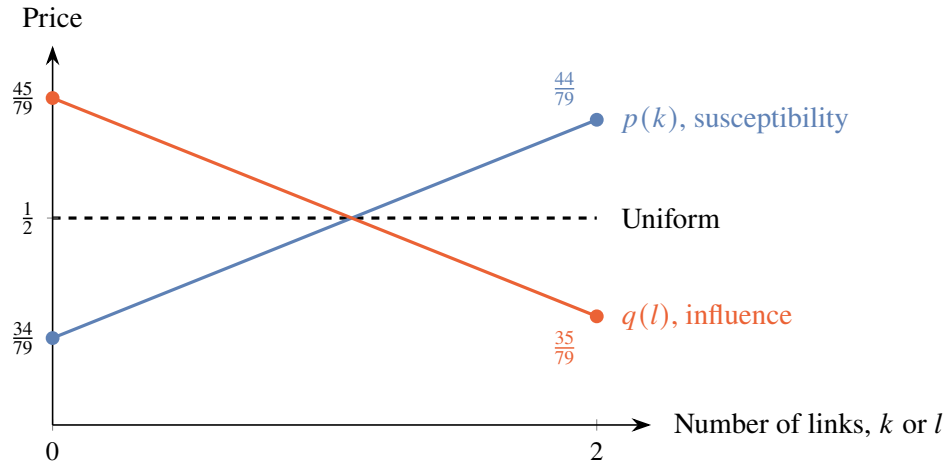


Figure 36.1: Optimal prices with $\gamma = \frac{1}{5}$: pricing on susceptibility k (blue), on influence l (red), and the uniform price (dashed).

Intuitively, an influencer's purchase raises the demand of the consumers who follow her, so the monopolist subsidizes her consumption and recovers the subsidy from everyone else through the common premium. The premium reflects average susceptibility, $\bar{k}$, because without data on susceptibility the monopolist charges every consumer for the network effects of an average consumer.

36.7 Fifth Step: Consumer Surplus and the Value of Network Data

Consumer surplus. By (36.4), consumer surplus is

$$CS = \frac{1}{8} \sum_{(k,l)} x(k,l)^2,$$

the sum running over the four types. Using Results 1–3,

$$CS^U = \frac{1+\gamma^2}{8(1-\gamma)^2}, \qquad CS^S = \frac{4+\gamma^2}{2D^2}, \qquad CS^I = \frac{\gamma^4+4\gamma^3+5\gamma^2+4}{2D^2}. \tag{36.12}$$

Value of network data. The value to the monopolist of learning susceptibility, or influence, is the increase in profit relative to the uniform price,

$$\Pi^S - \Pi^U = \Pi^I - \Pi^U = \frac{\gamma^2}{4(1-\gamma)D} > 0. \tag{36.13}$$

Result 4 (welfare and the value of data, Propositions 5–7 of the paper). Relative to the uniform price:

- pricing on influence raises consumer surplus: influencers and hubs gain, while isolated consumers and followers lose;
- pricing on susceptibility lowers consumer surplus in the toy model: consumers with $k = 0$ gain and consumers with $k = 2$ lose;
- both forms of price discrimination raise average consumption, profit, and total welfare;
- data on influence and data on susceptibility are equally valuable to the monopolist, since the two degrees have the same variance.

Figure 36.2 depicts consumer surplus under each form of price discrimination, relative to the uniform price, as a function of γ. Pricing on influence raises consumer surplus, by about 4% at $\gamma = \frac{1}{5}$, and the gain grows with network effects; pricing on susceptibility lowers it slightly, by about 0.4% at $\gamma = \frac{1}{5}$.

Intuitively, two facts drive these results. First, by (36.4), each consumer's surplus is convex in her consumption, so consumer surplus rises with both the average and the dispersion of consumption.



Figure 36.2: Consumer surplus under price discrimination on influence (red) and on susceptibility (blue, dashed), relative to the uniform price.

Second, both forms of price discrimination raise average consumption, but they have opposite effects on its dispersion. Pricing on susceptibility charges more to the consumers who consume more, which compresses consumption; pricing on influence lowers the price of influencers and raises it for everyone else, which spreads consumption out. In the paper, pricing on susceptibility raises consumer surplus only when average network effects $\bar{k}$ are large enough, a threshold that decreases with the dispersion of susceptibility; with $\bar{k} = 1$, the toy model is below the threshold for every admissible γ. Table 36.1 reports all outcomes at $\gamma = \frac{1}{5}$.

Table 36.1: Outcomes of the monopoly model with $\gamma = \frac{1}{5}$.

	Uniform price	Pricing susceptibility	Pricing influence
Prices	$\frac{1}{2}$	$p_0 = \frac{34}{79}$, $p_2 = \frac{44}{79}$	$q_0 = \frac{45}{79}$, $q_2 = \frac{35}{79}$
Consumption, $k = 0$	$\frac{1}{2}$	$\frac{45}{79}$	$\frac{34}{79}$ $(l = 0)$, $\frac{44}{79}$ $(l = 2)$
Consumption, $k = 2$	$\frac{3}{4}$	$\frac{55}{79}$	$\frac{56}{79}$ $(l = 0)$, $\frac{66}{79}$ $(l = 2)$
Neighbor consumption, A	$\frac{5}{8}$	$\frac{50}{79}$	$\frac{55}{79}$
Profit	0.3125	0.3165	0.3165
Consumer surplus	0.2031	0.2023	0.2120
Total welfare	0.5156	0.5187	0.5284

36.8 Sixth Step: Extension to Competing Firms

Model. Fainmesser and Galeotti (2020) let two firms compete. Firm Y is located at 0 and firm X at 1 of a unit line, both produce at zero cost, and a unit mass of consumers with unit demand is

uniformly distributed on the line. A consumer at ξ who buys from X at price p^X obtains

$$-(1-\xi)\tau - p^X + \gamma k A^X,$$

where $\tau \geq 1$ is the transport cost and A^X is the probability that a consumer with whom she interacts buys from X; buying from Y is symmetric. Products are incompatible: network effects come only from neighbors who buy the same product. As in the paper, every consumer has the same susceptibility, here $k = 1$, and consumers differ only in influence, $l \in \{0, 2\}$ with probability $\frac{1}{2}$ each, so a neighbor is always a consumer with $l = 2$. We keep $\gamma \in \left(0, \frac{1}{4}\right)$, the paper's Assumption 1, $\gamma l_{\max} < \frac{1}{2}$.

Data. Each firm learns the influence of a random fraction $w \in [0, 1]$ of consumers, independently of its rival. It offers $p(l)$ to the consumers it targets and p to the rest. With $w = 0$, firms charge uniform prices; with $w = 1$, they price discriminate on every consumer.

Demand. Let $\wp$ denote the gap between the prices a consumer is offered by X and by Y, and $\bar{\wp}$ the expected gap for a consumer with whom she interacts. The paper's Proposition 1 shows that she buys from X with probability

$$\frac{1}{2}\left(1 - \frac{\wp}{\tau} - \frac{\gamma k}{\tau(\tau - \gamma k)}\,\bar{\wp}\right). \tag{36.14}$$

A lower price from X to her neighbors raises her demand for X, since she expects them to buy X.

Prices. Firm X chooses its prices taking Y's prices as given. Raising the price $p(l)$ to a targeted consumer has three effects: it raises the margin on her, lowers her demand, and, through $\bar{\wp}$, lowers the demand of everyone who follows her, an effect that is larger the more influential she is. In a symmetric equilibrium $\bar{\wp} = 0$, a consumer whom X targets is targeted by Y with probability w, and the first-order conditions for the untargeted and the targeted prices, with $k = 1$, are

$$1 - \frac{p}{\tau} - \frac{\gamma\, p}{\tau(\tau - \gamma)} = 0, \qquad 1 - \frac{(1-w)\,[p(l) - p]}{\tau} - \frac{p(l)}{\tau} - \frac{\gamma\, l\, p}{\tau(\tau - \gamma)} = 0, \tag{36.15}$$

where we used that the average price each firm charges equals p, as the solution confirms.

Result 5 (equilibrium prices, Proposition 2 of the 2020 paper). Both firms charge $p = \tau - \gamma$ to the consumers they do not target, and

$$p(l) = \tau - \gamma + \frac{\gamma}{2 - w}\,(1 - l) \tag{36.16}$$

to the consumers they target. Hence:

- targeted followers ($l = 0$) pay a premium of $\frac{\gamma}{2-w}$;
- targeted influencers ($l = 2$) receive a discount of $\frac{\gamma}{2-w}$;
- the premium and the discount increase with the data w, and with $w = 1$ the schedule is $p(l) = \tau - \gamma l$.

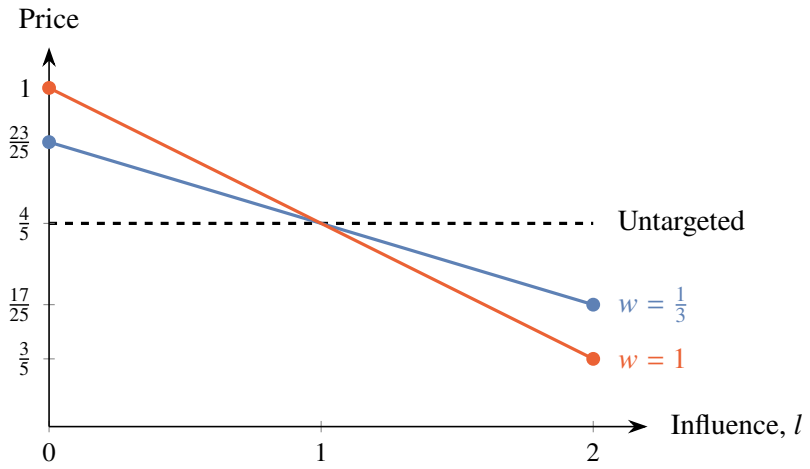


Figure 36.3: Equilibrium prices with competing firms ($\tau = 1$, $\gamma = \frac{1}{5}$): targeted prices with $w = \frac{1}{3}$ (blue) and $w = 1$ (red), and the untargeted price (dashed).

Figure 36.3 depicts the price schedules with $\tau = 1$ and $\gamma = \frac{1}{5}$. Untargeted consumers pay $\frac{4}{5}$. With $w = \frac{1}{3}$, targeted followers pay $\frac{23}{25}$ and targeted influencers $\frac{17}{25}$; with $w = 1$, the schedule widens to 1 and $\frac{3}{5}$.

Intuitively, each firm faces the monopolist's incentives of the fourth step, discounting influencers and charging followers more, but competition changes the role of network effects. With a monopolist that prices on influence, stronger network effects raise demand, the average price, and profit; with competing firms, stronger network effects make each firm's demand more elastic, since a price cut also attracts the consumers who follow the buyer, so the average price, $\tau - \gamma$, falls. More data intensify the competition for influencers: a firm that targets an influencer is more likely to face a rival that also targets her and offers her a discount, so it must discount more.

36.9 Seventh Step: Misallocation, Profits, and Consumer Surplus

Misallocation. Consumers targeted by both firms, or by neither, receive the same price from both, and by symmetry they expect the same network benefits from both products, so each buys the closer product. A mass $2w(1-w)$ of consumers is targeted by only one firm and faces a price gap of $\pm\frac{\gamma}{2-w}$. Since $\tau(2\xi - 1)$ is uniform on $[-\tau, \tau]$, a fraction $\frac{|\varphi|}{2\tau}$ of these consumers buys the farther product, at an average extra transport cost of $\frac{|\varphi|}{2}$, so the welfare loss is $\frac{\varphi^2}{4\tau}$ per consumer, and

$$M(w) = 2w(1-w) \cdot \frac{1}{4\tau}\left(\frac{\gamma}{2-w}\right)^2 = \frac{\gamma^2\, w(1-w)}{2\tau(2-w)^2}. \tag{36.17}$$

Profits and consumer surplus. Let $W^{\max}$ denote total surplus when every consumer buys the closer product. The paper's Corollary 1 shows that each firm's profit and consumer surplus are

$$\Pi(w) = \frac{\tau - \gamma}{2} - M(w), \qquad CS(w) = W^{\max} - (\tau - \gamma) + M(w), \tag{36.18}$$

so total welfare, $CS + 2\Pi = W^{\max} - M(w)$, falls by the misallocation.

Result 6 (misallocation and data, Propositions 3–4 of the 2020 paper). As firms obtain more data:

- for $w < \frac{2}{3}$, misallocation rises, each firm's profit falls, and consumer surplus rises;
- for $w > \frac{2}{3}$, the reverse holds, and with $w = 1$ misallocation disappears;
- total welfare is highest with no data or with full data, $w \in \{0, 1\}$.

Misallocation, and hence the loss in profits and the gain to consumers, increases with network effects γ and decreases with product differentiation τ.

Figure 36.4 depicts misallocation $M(w)$ with $\tau = 1$ and $\gamma = \frac{1}{5}$. It reaches its maximum, $\frac{1}{400}$, at $w = \frac{2}{3}$. At $w = \frac{1}{3}$, $M = \frac{1}{625}$, so each firm earns $\frac{2}{5} - \frac{1}{625}$ instead of $\frac{2}{5}$.

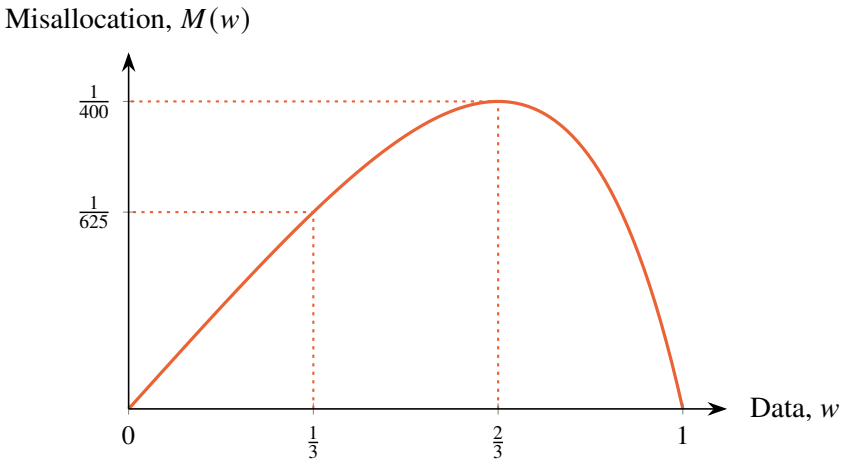


Figure 36.4: Misallocation as a function of the share of consumers on whom firms have data ($\tau = 1$, $\gamma = \frac{1}{5}$). Each firm's profit is $\frac{2}{5} - M(w)$.

Intuitively, misallocation requires a targeting mismatch: a consumer must be targeted by one firm but not the other. The mass of such consumers, $2w(1 - w)$, is largest at $w = \frac{1}{2}$, but the price gap they face, $\frac{\gamma}{2-w}$, keeps growing with w, so misallocation peaks at $w = \frac{2}{3}$. The two firms together lose $2M(w)$, consumers gain $M(w)$, and the difference, $M(w)$, is the welfare lost because some consumers buy the product they like less.

36.10 Eighth Step: The Race for Data

Setup. Consider that each firm chooses its data w_J, at cost $\frac{\alpha}{2}w_J^2$, at the same time as its prices. In a symmetric equilibrium, the marginal benefit of data to a firm, holding prices at their equilibrium levels, is

$$MB(w) = \frac{\gamma^2}{2\tau(2 - w)^2}, \tag{36.19}$$

which increases in w: the more data the rival has, the more valuable data are to a firm. Setting $MB(w) = \alpha w$ gives the equilibrium condition

$$w(2-w)^2 = r, \qquad \text{where} \quad r \equiv \frac{\gamma^2}{2\tau\alpha}. \tag{36.20}$$

The left side increases on $\left[0, \frac{2}{3}\right]$, reaches $\frac{32}{27}$ at $w = \frac{2}{3}$, and falls to 1 at $w = 1$.

Result 7 (the race for data, Propositions 6–7 of the 2020 paper). Data acquisition decisions are strategic complements, and:

- if $r < 1$, the unique stable equilibrium has $w^* < \frac{2}{3}$, the smallest root of (36.20);
- if $1 < r < \frac{32}{27}$, there are two stable equilibria, w^* and full data, $w = 1$;
- if $r > \frac{32}{27}$, firms acquire full data, $w = 1$.

Figure 36.5 depicts the left side of (36.20) and three values of r. With $\tau = 1$ and $\gamma = \frac{1}{5}$, a cost $\alpha = \frac{27}{1250}$ gives $r = \frac{25}{27}$ and $w^* = \frac{1}{3}$: misallocation is $\frac{1}{625}$, and each firm spends $\frac{3}{2500}$ on data, so its profit falls from $\frac{2}{5}$ to

$$\frac{2}{5} - \frac{1}{625} - \frac{3}{2500} = \frac{993}{2500}.$$

With $r = \frac{9}{8}$, the curve crosses the line at $w = \frac{1}{2}$ and at $w = \frac{7-\sqrt{13}}{4} \approx 0.85$, an unstable equilibrium, and full data is also an equilibrium.

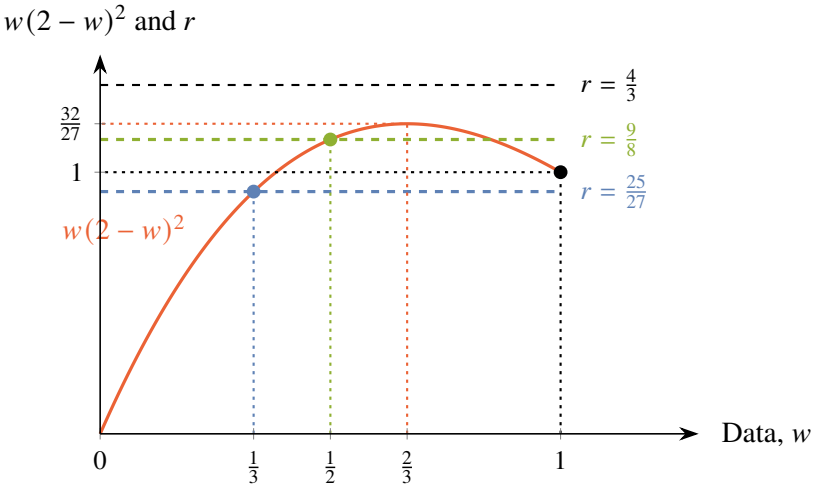


Figure 36.5: The equilibrium condition (36.20). Horizontal lines show $r = \frac{25}{27}$ (blue), $r = \frac{9}{8}$ (green), and $r = \frac{4}{3}$ (black). Colored dots mark the stable interior equilibria; the black dot marks the corner $w = 1$, an equilibrium whenever $r \geq 1$.

Intuitively, data on a consumer's influence are valuable to a firm because they let it set the premium and the discount of Result 5, $\frac{\gamma}{2-w}$, which widen as the rival acquires more data. Each firm therefore acquires data because the other does, and when the cost α of data falls, w^* rises. Starting from $w^* < \frac{2}{3}$, this raises misallocation, lowers profits, and raises consumer surplus, which is the paper's race for data: firms spend resources on information that leaves them worse off.

36.11 Ninth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The difference between influence and susceptibility, and the premium-discount structure of optimal prices: a premium for susceptibility and a discount for influence.
- The welfare comparison: pricing on influence raises consumer surplus, while pricing on susceptibility can lower it, both through the convexity of consumer surplus in consumption.
- The closed-form profits of the paper with a uniform price and with price discrimination,

$$\frac{1}{4(1-\gamma\bar{k})} \quad \text{and} \quad \frac{1}{4-4\gamma\bar{k}-\gamma^2\sigma^2},$$

  where $\sigma^2 = 1$ is the variance of the degree observed.
- With competition, the equilibrium price schedule, the misallocation that targeting mismatches create, its split between firms and consumers, and the race for data.

What it drops.

- General degree distributions and the comparative statics in average network effects and in their dispersion. In particular, the toy model cannot show that pricing on susceptibility raises consumer surplus when average network effects are large enough.
- Correlation between influence and susceptibility, and price discrimination on both, where the paper shows that data on influence and data on susceptibility are complements.
- Monopoly data on only part of the population, and negative network effects (such as smoking).
- In the 2020 paper, restrictions that let every consumer buy at the untargeted price, and how data affect entry: with data on influence, dispersed influence and strong network effects lower duopoly profits but raise monopoly profits, so they can act as a barrier to entry.

36.12 Connections to Other Chapters

- Chapter 1 studies the welfare effects of third-degree price discrimination without network effects.
- In Chapter 4, firms can also sell their customer data to a rival.
- In Chapter 35, more informative data on consumer preferences raise product consumption but tend to lower platforms' ad revenue.

References

Fainmesser, I. P., and Galeotti, A. (2016). Pricing network effects. *Review of Economic Studies*, 83(1), 165–198.

Fainmesser, I. P., and Galeotti, A. (2020). Pricing network effects: Competition. *American Economic Journal: Microeconomics*, 12(3), 1–32.

Part X

Innovation and Patents

The four chapters in this part study innovation, and each adds a force that the simplest analysis leaves out. Chapter 37 studies how rivals organize cost-reducing research when part of each firm's research leaks to its rival. Firms that coordinate their R&D spending invest more than independent firms only when this spillover is large, and firms that share all their results without coordinating free ride on each other and achieve the smallest cost reduction. Sharing results and coordinating spending together yield the largest cost reduction, the lowest prices, and the highest total surplus. Chapter 38 starts after the innovation has taken place. A cost reduction usually lowers prices and raises welfare, but when the firm that lowers its cost has a small output share, the reduction moves production away from more efficient rivals and welfare can fall, more easily under Bertrand than under Cournot competition. Chapters 39 and 40 study the incentive to innovate, and both build on Arrow's replacement effect: a firm that already earns a profit on a product gains less from replacing it. In Chapter 39, this effect weakens as a patent approaches its expiration date, and patent length and forward protection, which makes rivals pay damages to the leader, change who innovates and when. In Chapter 40, the replacement effect leads an incumbent to buy an entrant and shut down its project when the new product would compete with the incumbent's own.

Chapter 40 also links this part to the Mergers part. In Chapter 15, a merger between rival innovators lowers the merging firms' incentive to innovate through a force similar to the replacement effect; in Chapter 40, an acquisition can end a project altogether. Chapter 37 is also the base model of Chapters 42 and 44, which add pollution and emission taxes to its four ways of organizing research.

Chapter 37

Research Joint Ventures and R&D Cartels: Sharing versus Coordinating Research

Companion note to Kamien, M. I., E. Muller, and I. Zang (1992), "Research Joint Ventures and R&D Cartels," American Economic Review 82(5): 1293–1306.

37.1 Introduction

Rival firms sometimes do research together. Sematech, a consortium of 14 firms formed to develop technologies for producing computer chips, is an example (Kamien et al., 1992). Joint research avoids duplicated effort, but antitrust authorities worry that partners free ride on each other, so that less research is done and prices rise.

Cooperation in research has two parts that the analysis keeps apart. Firms can share their results, so that each firm's research lowers every partner's cost as much as its own; this is a *research joint venture* (RJV). Firms can also coordinate their research spending to maximize joint profit; this is an *R&D cartel*. Without sharing, part of a firm's research still leaks to its rival, at a rate called the *spillover*. Sharing removes duplication but invites free riding, and coordination makes each firm count its partner's profit.

Kamien et al. (1992) compare R&D competition, an R&D cartel, RJV competition, and an RJV cartel, with firms competing in the product market in all four. RJV competition yields the smallest cost reduction and the highest prices, while an RJV cartel yields the largest cost reduction and the highest consumer plus producer surplus, always under quantity competition and in most cases under price competition. Antitrust policy should therefore let research partners coordinate their spending, as long as they keep competing in the product market.

The toy model has two firms, linear demand for differentiated goods, and a quadratic R&D cost, so every equilibrium has a closed form. It is the base model behind the four arrangements used in the chapters on Poyago-Theotoky (2007) and Iida (2020). The first step solves the product market, the second and third steps compare R&D across the four arrangements, the fourth step compares

prices, profits, and surplus, the fifth step turns to price competition, and the sixth step lists what the toy model keeps and drops.

37.2 The Toy Model

Demand. Two firms, 1 and 2, produce differentiated goods. The inverse demand for good i is

$$P_i = 1 - Q_i - \gamma Q_j, \qquad \gamma = \frac{4}{5}, \tag{37.1}$$

where $\gamma \in [0, 1]$ measures how close substitutes the goods are: $\gamma = 0$ makes each firm a monopolist, and $\gamma = 1$ makes the goods identical. Demand (37.1) comes from a representative consumer with quasi-linear utility $Q_1 + Q_2 - \frac{1}{2}\left(Q_1^2 + Q_2^2 + 2\gamma Q_1 Q_2\right)$, so consumer surplus is

$$CS = \frac{1}{2}\left(Q_1^2 + Q_2^2 + 2\gamma Q_1 Q_2\right). \tag{37.2}$$

Costs and R&D. Both firms start with unit cost $c = \frac{1}{2}$, and we write $A \equiv 1 - c = \frac{1}{2}$. Firm i spends $x_i \geq 0$ on R&D. As in the paper, what lowers cost is the firm's *effective* R&D,

$$X_i = x_i + \beta x_j, \tag{37.3}$$

where $\beta \in [0, 1]$ is the spillover: a dollar spent by the rival counts as β dollars of firm i's own research. Firm i's unit cost in production is $c - y_i$, where $y_i = f(X_i)$ is its cost reduction and

$$f(X) = \sqrt{\frac{2X}{\eta}}, \qquad \eta = \frac{25}{4}. \tag{37.4}$$

Equivalently, lowering unit cost by y requires effective R&D $X = \frac{\eta}{2}y^2$, the quadratic R&D cost of d'Aspremont and Jacquemin (1988) applied to effective R&D. The function f is increasing and concave, with $f'(X) = \frac{1}{\eta f(X)}$.

Four arrangements. The paper's labels are used throughout:

- *R&D competition* (N): each firm chooses x_i to maximize its own profit, with X_i given by (37.3);
- *R&D cartel* (C): each firm chooses x_i to maximize the joint profit of the two firms, and the spillover remains β;
- *RJV competition* (NJ): firms share all results, so $X_1 = X_2 = x_1 + x_2$, the case $\beta = 1$, but each firm maximizes its own profit;
- *RJV cartel* (CJ): firms share all results and maximize joint profit.

Timing. In the first stage, firms choose their R&D spending. In the second stage, they observe their unit costs and compete in quantities (in prices in the fifth step). In no arrangement do firms coordinate their outputs or prices.

Why these values. With $\gamma = \frac{4}{5}$ and $\eta = \frac{25}{4}$ we have $(2+\gamma)^2\eta = 49$, which keeps the closed forms simple. The R&D cost η is large enough for four properties to hold in all four arrangements, at every β, and under both quantity and price competition. First, unit costs stay positive. Second, each firm's profit is concave in its own R&D at the equilibrium, and the equilibrium R&D is a best response against every alternative level. Third, each firm's best response has a slope between -1 and 0, so the R&D equilibrium is stable (Henriques, 1990); the slope is -1 only with a full spillover, where only total spending matters and, as in the paper, each firm spends half. Fourth, a cartel that could assign unequal R&D to the two firms would still choose equal R&D. The paper assumes a bounded f, which the square root is not, so the toy model checks these properties directly.

37.3 First Step: The Product Market

Given cost reductions y_1 and y_2, firm i chooses Q_i to maximize $\left(1 - Q_i - \gamma Q_j - c + y_i\right) Q_i$. The first-order condition is

$$A + y_i - 2Q_i - \gamma Q_j = 0,$$

and solving the two conditions together gives

$$Q_i = \frac{1}{2+\gamma}\left[A + \frac{2y_i - \gamma y_j}{2-\gamma}\right] = \frac{5}{14}\left[\frac{1}{2} + \frac{5}{3}\,y_i - \frac{2}{3}\,y_j\right]. \tag{37.5}$$

By the first-order condition, the price-cost margin equals output, $P_i - (c - y_i) = Q_i$, so firm i's profit net of R&D spending is

$$\pi_i = Q_i^2 - x_i. \tag{37.6}$$

Result 1 (the product market). Given the cost reductions y_1 and y_2:

- each firm produces (37.5) and earns (37.6);
- firm i's own cost reduction raises its output, $\frac{\partial Q_i}{\partial y_i} = \frac{2}{4-\gamma^2} = \frac{25}{42}$, and its rival's cost reduction lowers it, $\frac{\partial Q_i}{\partial y_j} = -\frac{\gamma}{4-\gamma^2} = -\frac{5}{21}$;
- when both firms have the same cost reduction y, each produces

$$Q = \frac{A+y}{2+\gamma} = \frac{5}{14}\left(\frac{1}{2} + y\right), \tag{37.7}$$

  and the price is $P = 1 - (1+\gamma)Q = 1 - \frac{9}{5}\,Q$.[1]

Intuitively, a lower own cost makes firm i more aggressive in the product market, and a lower rival cost makes the rival more aggressive, which reduces firm i's output. The ratio of the two effects, $\frac{\gamma}{2} = \frac{2}{5}$, measures how much of the gain from a cost reduction is lost when the rival obtains the same reduction, and it reappears in the second step as the critical spillover.

[1] On p. 1300, the paper writes the symmetric output as $\frac{a-c-f(X)}{2+\gamma(n-1)}$. By its equation (6), and as in its equation (18), the output is $\frac{a-c+f(X)}{2+\gamma(n-1)}$, which is (37.7) with $n = 2$. The argument there uses only the sign of this expression.

37.4 Second Step: R&D Competition and the R&D Cartel

R&D competition. Firm i chooses x_i to maximize (37.6). Since $\frac{\partial X_i}{\partial x_i} = 1$ and $\frac{\partial X_j}{\partial x_i} = \beta$, by (37.5) the first-order condition is

$$\frac{\partial \pi_i}{\partial x_i} = \frac{2Q_i}{4-\gamma^2}\Big[\underbrace{2f'(X_i)}_{\text{own cost}} \quad \underbrace{-\gamma\beta f'(X_j)}_{\text{spillover to the rival}}\Big] - 1 = 0. \tag{37.8}$$

The first term is the gain from a lower own cost. The second term is the paper's competitive-advantage externality: part of firm i's research lowers the rival's cost, which makes the rival a tougher competitor and reduces firm i's output.

R&D cartel. Firm i chooses x_i to maximize $\pi_1 + \pi_2$, taking x_j as given, as in the paper's problem (10). Its first-order condition adds to (37.8) the effect of x_i on the rival's profit,

$$\frac{\partial \pi_j}{\partial x_i} = \frac{2Q_j}{4-\gamma^2}\Big[2\beta f'(X_j) - \gamma f'(X_i)\Big], \tag{37.9}$$

the paper's combined-profits externality: the spillover lowers the rival's cost, which raises its profit, while firm i's lower cost makes firm i tougher, which reduces it.

Symmetric equilibrium. At a symmetric profile, $Q_1 = Q_2 = Q$ and $f'(X_1) = f'(X_2) = f'(X)$, so every term is a multiple of $\frac{2Qf'(X)}{4-\gamma^2}$. Table 37.1 reports the multiples for the four arrangements; the two RJV rows are used in the third step. In each arrangement the first-order condition becomes

$$\frac{2Q\,f'(X)}{2+\gamma}\,k^{\ell} = 1, \tag{37.10}$$

where k^{ℓ}, the sum of the row divided by $2-\gamma$, is the *return factor* of arrangement $\ell = N, C, NJ, CJ$:

$$k^{N} = \frac{2-\gamma\beta}{2-\gamma} = \frac{5-2\beta}{3}, \qquad k^{C} = 1+\beta, \qquad k^{NJ} = 1, \qquad k^{CJ} = 2. \tag{37.11}$$

Arrangement	Own cost	Spillover to rival	Rival's profit	Sum	Sum at $\gamma = \frac{4}{5}$
R&D competition (N)	2	$-\gamma\beta$	0	$2-\gamma\beta$	$2-\frac{4}{5}\beta$
R&D cartel (C)	2	$-\gamma\beta$	$2\beta-\gamma$	$(2-\gamma)(1+\beta)$	$\frac{6}{5}(1+\beta)$
RJV competition (NJ)	2	$-\gamma$	0	$2-\gamma$	$\frac{6}{5}$
RJV cartel (CJ)	2	$-\gamma$	$2-\gamma$	$2(2-\gamma)$	$\frac{12}{5}$

Table 37.1: Terms of the first-order condition for R&D at a symmetric profile, in units of $\frac{2Qf'(X)}{4-\gamma^2}$. The rival's profit enters only when firms coordinate their R&D, and in an RJV the spillover is $\beta = 1$.

Substituting $Q = \frac{A+y}{2+\gamma}$ from (37.7) and $f'(X) = \frac{1}{\eta y}$ into (37.10) gives $2(A+y)\,k^{\ell} = (2+\gamma)^2\eta\,y$, so the equilibrium cost reduction is

$$y^{\ell} = \frac{2A\,k^{\ell}}{(2+\gamma)^2\eta - 2k^{\ell}} = \frac{k^{\ell}}{49-2k^{\ell}}, \tag{37.12}$$

which increases in k^ℓ. Effective R&D is $X^\ell = \frac{\eta}{2}\left(y^\ell\right)^2$, and spending per firm is $x^\ell = \frac{X^\ell}{1+\beta}$ in the two arrangements without an RJV.

Result 2 (R&D competition versus the R&D cartel; Proposition 1 of the paper, cases N and C). The equilibrium cost reductions are

$$y^N = \frac{5-2\beta}{137+4\beta}, \qquad y^C = \frac{1+\beta}{47-2\beta}, \tag{37.13}$$

and

$$y^C - y^N = \frac{49\,(5\beta-2)}{(47-2\beta)(137+4\beta)}.$$

Therefore:

- the R&D cartel achieves the larger cost reduction if and only if $\beta > \frac{\gamma}{2} = \frac{2}{5}$, and the two coincide at $\beta = \frac{2}{5}$, where $y^C = y^N = \frac{1}{33}$;
- y^N decreases in β, from $\frac{5}{137}$ to $\frac{1}{47}$, and y^C increases in β, from $\frac{1}{47}$ to $\frac{2}{45}$;
- effective R&D $X = \frac{\eta}{2}y^2$ is ranked in the same way.

Intuitively, the cartel adds the combined-profits externality of (37.9), whose multiple in Table 37.1 is $2\beta - \gamma$. One dollar of firm i's research lowers the rival's cost through the spillover, which raises the rival's profit by a term proportional to 2β, and lowers firm i's own cost, which costs the rival a term proportional to γ. When the spillover is small, the second effect dominates, so the cartel invests less than independent firms to avoid hurting the partner; when $\beta > \frac{\gamma}{2}$, the first dominates and the cartel invests more. The paper extends this comparison from d'Aspremont and Jacquemin (1988), and with identical goods, $\gamma = 1$, the critical spillover is $\frac{1}{2}$.

37.5 Third Step: Research Joint Ventures and Free Riding

In an RJV, both firms' effective R&D is $X = x_1 + x_2$, so each dollar either firm spends lowers both costs equally. The first-order condition of RJV competition is (37.8) with $\beta = 1$, so $k^{NJ} = 1$, and that of the RJV cartel adds (37.9) with $\beta = 1$, so $k^{CJ} = 2$. By (37.12),

$$y^{NJ} = \frac{1}{47}, \qquad y^{CJ} = \frac{2}{45}, \tag{37.14}$$

and spending per firm is $x^\ell = \frac{X^\ell}{2}$, so $x^{NJ} = \frac{25}{35344}$ and $x^{CJ} = \frac{1}{324}$.

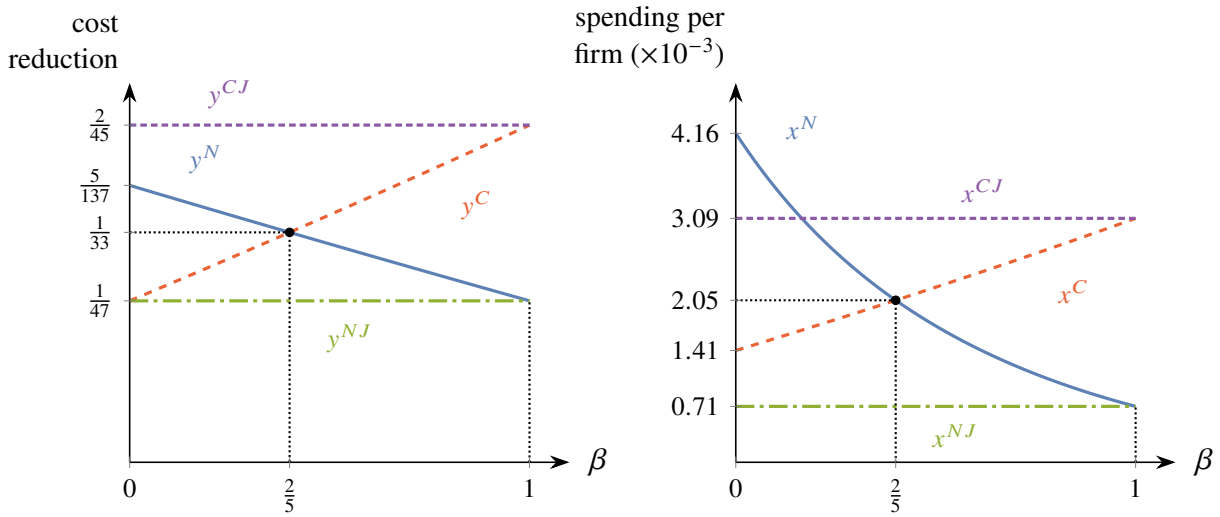
Free riding, term by term. Compare the rows N and NJ of Table 37.1. The own-cost term is 2 in both. The spillover to the rival is $-\gamma\beta$ under R&D competition and $-\gamma$ in RJV competition: in an RJV, every dollar firm i spends gives the rival the same cost reduction as firm i, so none of it buys a competitive advantage, and the private return falls to its lowest level, $k^{NJ} = 1$. The receiving side works in the same direction. By (37.3), a dollar of the rival's spending counts as β dollars of firm i's own, and in an RJV as a full dollar, so the more firm i receives from its rival, the less it

spends itself. The slope of firm i's best response at the R&D-competition equilibrium is -0.051 at $\beta = 0$, -0.155 at $\beta = \frac{1}{5}$, and -1 at $\beta = 1$, where only total spending matters, as in RJV competition.

Result 3 (Proposition 1 and Corollary 1 of the paper). For every spillover $\beta \in [0, 1]$:

- the cost reductions are ranked $y^{CJ} \geq y^{C} \geq y^{NJ}$ and $y^{CJ} \geq y^{N} \geq y^{NJ}$, and effective R&D is ranked in the same way;
- RJV competition coincides with R&D competition only at $\beta = 1$, the R&D cartel coincides with the RJV cartel only at $\beta = 1$, and the R&D cartel and RJV competition reach the same cost reduction only at $\beta = 0$;
- R&D spending per firm under R&D competition falls with the spillover, from $x^{N} = \frac{625}{150152} \approx 4.16 \times 10^{-3}$ at $\beta = 0$ to $x^{NJ} \approx 0.71 \times 10^{-3}$ at $\beta = 1$.

Figure 37.1 depicts the four arrangements against the spillover β: the cost reduction in panel (a) and R&D spending per firm in panel (b). The RJV arrangements do not depend on β, since sharing sets the spillover to 1. The R&D cartel overtakes R&D competition at $\beta = \frac{2}{5}$, where both reach $y = \frac{1}{33}$ with spending $\frac{125}{60984} \approx 2.05 \times 10^{-3}$ per firm.



(a) Cost reduction y^{ℓ} (b) R&D spending per firm x^{ℓ}

Figure 37.1: Cost reduction and R&D spending per firm against the spillover β under R&D competition (N, solid), the R&D cartel (C, dashed), RJV competition (NJ, dash-dotted), and the RJV cartel (CJ, short dashes), with $\gamma = \frac{4}{5}$, $\eta = \frac{25}{4}$, and $c = \frac{1}{2}$.

Intuitively, RJV competition combines the largest spillover with no coordination, so free riding is at its worst: each firm counts on its partner's spending and gives away the full benefit of its own. The RJV cartel removes both problems. Sharing makes each dollar lower both firms' costs, and coordination makes each firm value the reduction in its partner's cost as its own, so $k^{CJ} = 2$. Panel (b) shows the gain from eliminating duplication: at $\beta = 0$, the RJV cartel spends $\frac{1}{324} \approx 3.09 \times 10^{-3}$ per firm, less than the 4.16×10^{-3} of R&D competition, and yet obtains the larger cost reduction, $\frac{2}{45} > \frac{5}{137}$.

37.6 Fourth Step: Prices, Profits, and Surplus

With the same cost reduction y^ℓ for both firms, (37.7) gives output $Q^\ell = \frac{5}{14}\left(\frac{1}{2} + y^\ell\right)$ and the price

$$P^\ell = 1 - \frac{9}{5}Q^\ell = 1 - \frac{9}{14}\left(\frac{1}{2} + y^\ell\right). \tag{37.15}$$

Profit per firm is $\pi^\ell = \left(Q^\ell\right)^2 - x^\ell$, consumer surplus (37.2) is $CS^\ell = (1+\gamma)\left(Q^\ell\right)^2 = \frac{9}{5}\left(Q^\ell\right)^2$, and total surplus, the sum of consumer surplus and both firms' profits, is

$$TS^\ell = CS^\ell + 2\pi^\ell = \frac{19}{5}\left(Q^\ell\right)^2 - 2x^\ell. \tag{37.16}$$

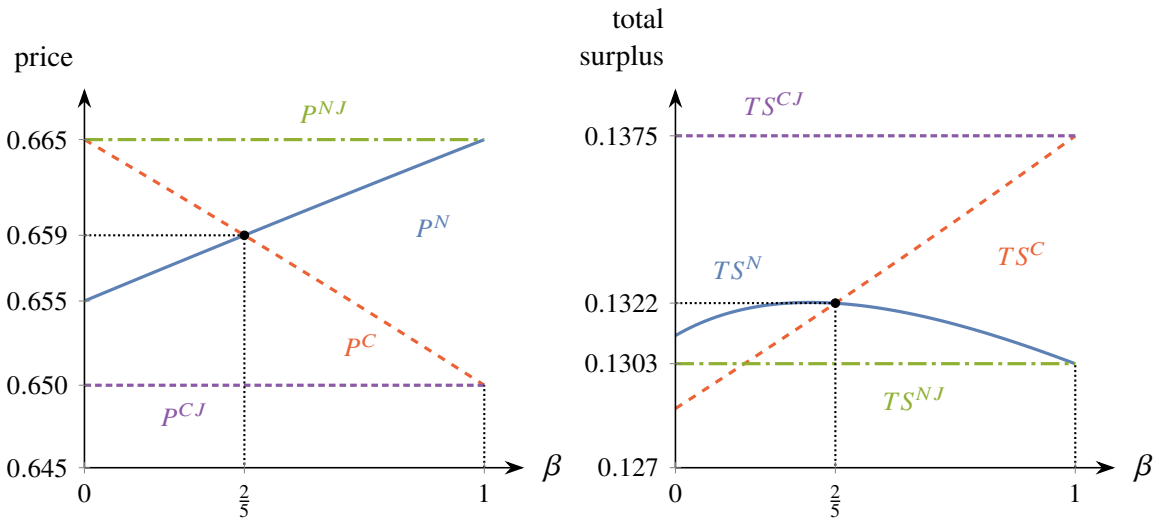
Result 4 (Corollary 2, Proposition 2, and Theorem 1 of the paper). For every $\beta \in [0,1]$:

- prices are ranked $P^{CJ} \leq P^C \leq P^{NJ}$ and $P^{CJ} \leq P^N \leq P^{NJ}$, with $P^{CJ} = \frac{13}{20}$ and $P^{NJ} = \frac{125}{188} \approx 0.665$, and $P^C < P^N$ if and only if $\beta > \frac{2}{5}$;
- profit per firm is highest in the RJV cartel, $\pi^{CJ} \geq \pi^\ell$, and $\pi^C \geq \pi^N$;
- total surplus is highest in the RJV cartel, $TS^{CJ} = \frac{11}{80} = 0.1375$;
- in the toy model, $TS^C > TS^N$ if and only if $\beta > \frac{2}{5}$, and $TS^{NJ} < TS^N$ for every $\beta < 1$.

Figure 37.2 depicts the price in panel (a) and total surplus in panel (b) against the spillover β. The price mirrors the cost reduction of Figure 37.1(a): it is highest in RJV competition and lowest in the RJV cartel, and the R&D cartel and R&D competition cross at $\beta = \frac{2}{5}$, where the price is $\frac{29}{44} \approx 0.659$. Total surplus is $\frac{4605}{35344} \approx 0.1303$ in RJV competition, below its level under R&D competition, which peaks at 0.1322 near $\beta = 0.34$, and both lie below 0.1375 in the RJV cartel.

Table 37.2 reports every step of the calculation at $\beta = \frac{1}{5}$, below the critical spillover $\frac{2}{5}$. R&D competition yields more cost reduction than the R&D cartel, a lower price, and higher total surplus, while the R&D cartel earns more profit. RJV competition yields the least cost reduction and the highest price, and the RJV cartel the most cost reduction, the lowest price, the highest profit, and the highest total surplus.

Intuitively, the price follows the cost reduction, by (37.15). Profits favor the RJV cartel because it can always reproduce the effective R&D of any other arrangement at no higher spending: with full sharing, effective R&D X costs each firm $\frac{X}{2}$, while without sharing it costs $\frac{X}{1+\beta}$. The same argument, with the cartel choosing the spending of R&D competition, gives $\pi^C \geq \pi^N$. Since the RJV cartel has both the lowest price and the highest profit, it has the highest total surplus. The toy model adds two observations. First, the R&D cartel lowers total surplus below the level of R&D competition when $\beta < \frac{2}{5}$, because consumers lose more from the higher price than firms gain. Second, firms earn more in RJV competition than in R&D competition when $\beta < 0.515$, as in Table 37.2, where $\pi^{NJ} = 0.0340 > \pi^N = 0.0334$: firms then want to form an RJV even without coordinating their spending, which is the arrangement with the least cost reduction and the highest price.



(a) Price P^ℓ (b) Total surplus TS^ℓ

Figure 37.2: Price and total surplus against the spillover β in the four arrangements, with the line styles of Figure 37.1; $\gamma = \frac{4}{5}$, $\eta = \frac{25}{4}$, and $c = \frac{1}{2}$. The vertical axes start at 0.645 and 0.127.

Variable	N	C	NJ	CJ
Return factor k^ℓ, from (37.11)	$\frac{5-2\beta}{3} = \frac{23}{15}$	$1+\beta = \frac{6}{5}$	1	2
Cost reduction $y = \frac{k}{49-2k}$	$\frac{23}{689} = 0.0334$	$\frac{6}{233} = 0.0258$	$\frac{1}{47} = 0.0213$	$\frac{2}{45} = 0.0444$
Effective R&D $X = \frac{25}{8}y^2$ ($\times 10^{-3}$)	3.482	2.072	1.415	6.173
Spending per firm x ($\times 10^{-3}$)	$\frac{3.482}{1.2} = 2.902$	$\frac{2.072}{1.2} = 1.727$	$\frac{1.415}{2} = 0.707$	$\frac{6.173}{2} = 3.086$
Output $Q = \frac{5}{14}\left(\frac{1}{2}+y\right)$	0.1905	0.1878	0.1862	0.1944
Price $P = 1 - \frac{9}{5}Q$	0.6571	0.6620	0.6649	0.6500
Profit $\pi = Q^2 - x$	0.0334	0.0335	0.0340	0.0347
Consumer surplus $CS = \frac{9}{5}Q^2$	0.0653	0.0635	0.0624	0.0681
Total surplus $TS = CS + 2\pi$	0.1321	0.1305	0.1303	0.1375

Table 37.2: Equilibrium of the four arrangements at $\beta = \frac{1}{5}$, with $\gamma = \frac{4}{5}$, $\eta = \frac{25}{4}$, and $c = \frac{1}{2}$. Spending per firm is $\frac{X}{1+\beta}$ without an RJV and $\frac{X}{2}$ in an RJV. Each entry is computed from the unrounded entries above it.

37.7 Fifth Step: Price Competition

Consider now that firms set prices in the second stage. Inverting (37.1), the demand for good i is

$$Q_i = \frac{(1-\gamma) - P_i + \gamma P_j}{1-\gamma^2},$$

the paper's system with $n = 2$. Solving the two first-order conditions, firm i's price-cost margin is

$$m_i = \frac{(1-\gamma)(2+\gamma)A + (2-\gamma^2)\, y_i - \gamma\, y_j}{4-\gamma^2}, \tag{37.17}$$

its output is $Q_i = \frac{m_i}{1-\gamma^2}$, and its profit is $\frac{m_i^2}{1-\gamma^2} - x_i$. With the same cost reduction y for both firms, the price is $P = \frac{(1-\gamma)+c-y}{2-\gamma}$. Repeating the second and third steps, the first-order condition of each arrangement again differs from the others only through a return factor,

$$k_B^N = \frac{2-\gamma^2-\gamma\beta}{(1-\gamma)(2+\gamma)} = \frac{17-10\beta}{7}, \qquad k_B^C = 1+\beta, \qquad k_B^{NJ} = 1, \qquad k_B^{CJ} = 2, \tag{37.18}$$

and the cost reductions are

$$y_B^\ell = \frac{2A\, k_B^\ell}{\phi_B - 2k_B^\ell}, \qquad \phi_B = \frac{(2-\gamma)^2(1+\gamma)\,\eta}{1-\gamma} = 81, \tag{37.19}$$

so that

$$y_B^N = \frac{17-10\beta}{533+20\beta}, \qquad y_B^C = \frac{1+\beta}{79-2\beta}, \qquad y_B^{NJ} = \frac{1}{79}, \qquad y_B^{CJ} = \frac{2}{77}.$$

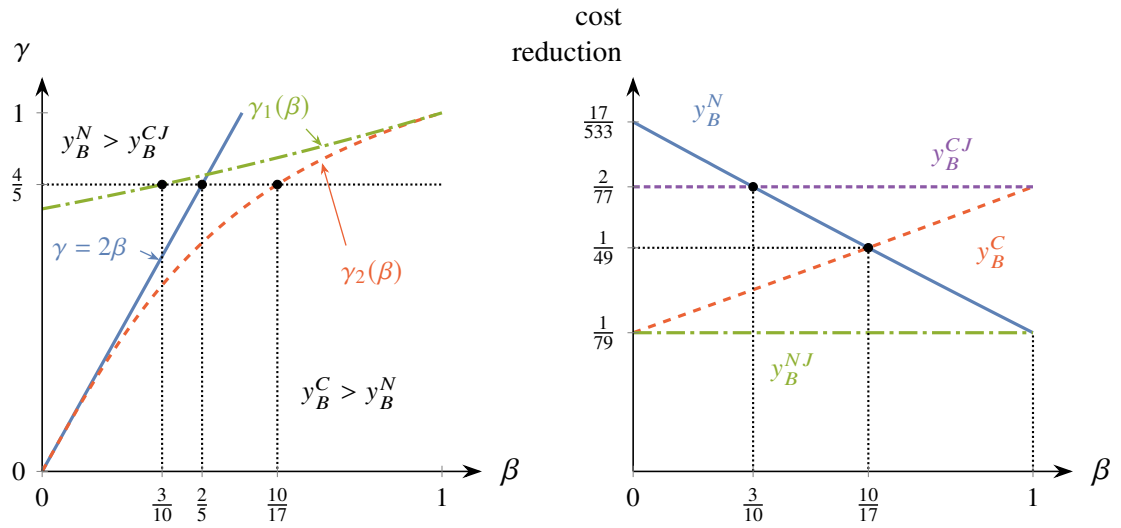
In the R&D cartel, the combined-profits externality now has the sign of $(2-\gamma^2)\beta - \gamma$, positive if and only if $\beta > \frac{\gamma}{2-\gamma^2}$, which equals $\frac{10}{17}$ at $\gamma = \frac{4}{5}$. Define

$$\gamma_1(\beta) = \frac{\sqrt{(2-\beta)^2+8} - (2-\beta)}{2} \qquad \text{and} \qquad \gamma_2(\beta) = \frac{4\beta}{1+\sqrt{1+8\beta^2}},$$

the roots in $(0,1]$ of $k_B^{CJ} = k_B^N$ and $k_B^C = k_B^N$, that is, of $\gamma^2 + (2-\beta)\gamma - 2 = 0$ and $\beta\gamma^2 + \gamma - 2\beta = 0$.

Result 5 (Proposition 1 of the paper).* Under price competition, for every $\beta \in [0,1]$ and $\gamma \in (0,1)$:

- the cost reductions are ranked $y_B^{CJ} \geq y_B^C \geq y_B^{NJ}$ and $y_B^N \geq y_B^{NJ}$;
- $y_B^{CJ} \geq y_B^N$ if and only if $\gamma \leq \gamma_1(\beta)$, and $y_B^C \geq y_B^N$ if and only if $\gamma \leq \gamma_2(\beta)$;
- γ_1 and γ_2 increase in β, $\gamma_2(\beta) < \gamma_1(\beta)$ for $\beta < 1$, $\gamma_1(1) = \gamma_2(1) = 1$, and $\gamma_1(\beta) \geq \gamma_1(0) = \sqrt{3} - 1 \approx 0.732$, above the bound of $\frac{2}{3}$ that the paper reports;
- at $\gamma = \frac{4}{5}$, the RJV cartel achieves a smaller cost reduction than R&D competition if and only if $\beta < \frac{3}{10}$, and the R&D cartel a smaller one if and only if $\beta < \frac{10}{17}$.



(a) Cutoffs in the (β, γ) plane

(b) Cost reduction at $\gamma = \frac{4}{5}$

Figure 37.3: Price competition. Panel (a): R&D competition yields a larger cost reduction than the RJV cartel above $\gamma_1(\beta)$ (dash-dotted) and than the R&D cartel above $\gamma_2(\beta)$ (dashed); the solid line is the cutoff $\gamma = 2\beta$ of the R&D cartel under quantity competition. Panel (b): cost reductions at $\gamma = \frac{4}{5}$, with the line styles of Figure 37.1, $\eta = \frac{25}{4}$, and $c = \frac{1}{2}$.

Figure 37.3 depicts the cutoffs in panel (a) and the cost reductions at $\gamma = \frac{4}{5}$ in panel (b). In panel (a), R&D competition achieves a larger cost reduction than the RJV cartel above $\gamma_1(\beta)$, and a larger one than the R&D cartel above $\gamma_2(\beta)$; under quantity competition the second cutoff is the line $\gamma = 2\beta$ of the second step, and the first does not exist. Along the dotted line $\gamma = \frac{4}{5}$, the cutoffs are $\beta = \frac{3}{10}$, $\frac{2}{5}$, and $\frac{10}{17}$. In panel (b), y_B^N crosses $y_B^{CJ} = \frac{2}{77}$ at $\beta = \frac{3}{10}$ and y_B^C at $\beta = \frac{10}{17}$, where both equal $\frac{1}{49}$.

Intuitively, the RJV cartel lowers both firms' costs together, while R&D competition gives each firm a cost advantage over its rival. By (37.17), a cost reduction common to both firms raises each margin by $\frac{1-\gamma}{2-\gamma} = \frac{1}{6}$ per unit, since price competition passes most of it on to consumers, while a cost reduction for firm i alone raises its margin by $\frac{2-\gamma^2}{4-\gamma^2} = \frac{17}{42}$. A common reduction is therefore worth $\frac{7}{17} \approx 41\%$ of an own reduction under price competition, against 60% under quantity competition, where by (37.5) a common reduction raises each output by $\frac{5}{14}$ per unit and an own reduction raises it by $\frac{25}{42}$. With close substitutes and small spillovers, the pursuit of a cost advantage leads firms that compete in R&D to a larger cost reduction than the RJV cartel. The welfare comparison depends on how close substitutes the goods are. At $\gamma = \frac{4}{5}$, the RJV cartel still has the highest total surplus for every β, but for $\beta < \frac{3}{10}$ consumers pay more than under R&D competition: at $\beta = 0$ the prices are $P^{CJ} = 0.562$ and $P^N = 0.557$. With closer substitutes the RJV cartel can lose its lead in total surplus: at $\beta = 0$ this happens for every $\gamma > 0.821$, and at $\gamma = \frac{9}{10}$ total surplus is 0.1410 under R&D competition and 0.1374 in the RJV cartel. This is the sense in which the RJV cartel is best only "in most cases" under price competition.

37.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The four arrangements and the two externalities behind them: the competitive-advantage externality, which every firm takes into account, and the combined-profits externality, which only coordinating firms take into account (the second step).
- The critical spillover $\beta = \frac{\gamma}{2}$ above which an R&D cartel invests more than independent firms, with lower prices, as in Proposition 1 and Corollaries 1 and 2 (the second and fourth steps).
- Free riding in RJV competition, which yields the smallest cost reduction and the highest price, and the RJV cartel, which yields the largest cost reduction, the lowest price, the highest profit, and the highest total surplus under quantity competition (the third and fourth steps, Proposition 2 and Theorem 1).
- Under price competition, the cutoffs $\gamma_1(\beta)$ and $\gamma_2(\beta)$ of Proposition 1*, so that with close substitutes and small spillovers the RJV cartel achieves a smaller cost reduction than independent firms (the fifth step).

What it drops.

- More than two firms and a general concave R&D production function. The paper assumes that f is bounded; the square root is not, and the toy model instead checks concavity, stability, and positive costs directly for its parameter values.
- Asymmetric equilibria. The paper considers only symmetric equilibria and notes that an R&D cartel may prefer unequal R&D (Salant and Shaffer, 1998); the R&D cost of the toy model is high enough that it does not.
- An RJV that raises the spillover only partly, toward but below 1, and a complete cartel that also coordinates output (Section III of the paper).
- Environmental policy, which the chapters on Poyago-Theotoky (2007) and Iida (2020) add to these four arrangements.

37.9 Connections to Other Chapters

- Chapter 42 adds pollution and an emission tax set after firms invest to the comparison between R&D competition and an R&D cartel, and treats an RJV as full spillovers.
- Chapter 44 compares the same four arrangements when two governments set emission taxes strategically, and finds that an RJV cartel can lower welfare, unlike in this chapter.
- Chapter 16 also compares independent and coordinated cost-reducing investment, with coordination through a merger instead of an R&D agreement.

References

d'Aspremont, C., and A. Jacquemin (1988). "Cooperative and Noncooperative R&D in Duopoly with Spillovers." *American Economic Review*, 78(5), pp. 1133–1137.

Henriques, I. (1990). "Cooperative and Noncooperative R&D in Duopoly with Spillovers: Comment." *American Economic Review*, 80(3), pp. 638–640.

Iida, T. (2020). "International R&D Formations and Strategic Environmental Policy." *Environment and Development Economics*, 25(5), pp. 509–528.

Kamien, M. I., E. Muller, and I. Zang (1992). "Research Joint Ventures and R&D Cartels." *American Economic Review*, 82(5), pp. 1293–1306.

Poyago-Theotoky, J. A. (2007). "The Organization of R&D and Environmental Policy." *Journal of Economic Behavior & Organization*, 62(1), pp. 63–75.

Salant, S. W., and G. Shaffer (1998). "Optimal Asymmetric Strategies in Research Joint Ventures." *International Journal of Industrial Organization*, 16(2), pp. 195–208.

Chapter 38

When Cost Reductions Lower Welfare: Bertrand versus Cournot

Companion note to Wang, X. H., and J. Zhao (2007), "Welfare Reductions from Small Cost Reductions in Differentiated Oligopoly," International Journal of Industrial Organization 25(1): 173–185.

38.1 Introduction

Consider a small, high-cost coffee shop next to a large, efficient chain. The small shop finds a cheaper supplier and lowers its cost. Prices fall and consumers gain. Can total welfare still fall? Lahiri and Ono (1988) showed that it can in Cournot competition with a homogeneous good. Wang and Zhao (2007) show that the same happens with differentiated products, under both Bertrand and Cournot competition.

The reason is that a cost reduction moves production between firms. The inefficient firm saves on the few units it already produces, and it expands at the expense of its efficient rival. Each unit that the efficient firm gives up was worth its large markup, while each unit that the inefficient firm adds is worth only its small markup. When the inefficient firm produces little, its cost saving is small, and this shift of production toward the higher-cost firm dominates, so total welfare falls even though consumer surplus rises.

Wang and Zhao (2007) show that a small cost reduction lowers welfare whenever the firm's output share is below a critical level. They also show that a welfare loss is more likely when firms compete in prices: if a cost reduction lowers welfare under Cournot competition, it also lowers welfare under Bertrand competition, but not the reverse. Intuitively, price competition is tougher on the inefficient firm, so it leaves that firm with a smaller share of the market. A lower cost is therefore not always good news for welfare, and the form of competition affects when it is.

The toy model is a duopoly with linear demand in which one firm produces at zero cost. The first step solves the Bertrand equilibrium. The second and third steps find when a cost reduction lowers

profits and welfare under Bertrand and under Cournot competition, and the fourth step compares the two. The fifth step extends the comparison to many firms, and the sixth step lists what the toy model keeps and drops.

38.2 The Toy Model

Demand. Two firms sell differentiated products. A representative consumer with utility

$$u(q_1, q_2) = q_1 + q_2 - \frac{q_1^2 + q_1 q_2 + q_2^2}{3} \tag{38.1}$$

has inverse demands $p_i = 1 - \frac{2q_i + q_j}{3}$ and direct demands

$$q_i = 1 - 2p_i + p_j. \tag{38.2}$$

This is the Shubik demand of the paper,

$$q_i = V - p_i - \gamma(p_i - \bar{p}),$$

with two firms, $V = 1$, and substitutability $\gamma = 2$, where $\bar{p}$ is the average price.

Costs and welfare. Firm 1 produces at zero cost and firm 2 at constant marginal cost $c \geq 0$, so firm 2 is the inefficient firm. Consumer surplus, total profit, and total welfare are

$$CS = u - p_1 q_1 - p_2 q_2, \qquad \Pi = p_1 q_1 + (p_2 - c) q_2, \qquad W = CS + \Pi = u - c q_2.$$

Question. How does a small reduction in c affect CS, Π, and W when firms compete in prices (Bertrand) and in quantities (Cournot)?

38.3 First Step: The Bertrand Equilibrium

Firm 1 maximizes $p_1(1 - 2p_1 + p_2)$ and firm 2 maximizes $(p_2 - c)(1 - 2p_2 + p_1)$. The first-order conditions

$$1 - 4p_1 + p_2 = 0 \qquad \text{and} \qquad 1 - 4p_2 + p_1 + 2c = 0$$

give

$$p_1^B = \frac{5 + 2c}{15}, \quad p_2^B = \frac{5 + 8c}{15}, \qquad q_1^B = \frac{2(5 + 2c)}{15}, \quad q_2^B = \frac{2(5 - 7c)}{15}. \tag{38.3}$$

Firm 2 produces if and only if $c < \frac{5}{7}$, and its output share is

$$\frac{q_2^B}{Q^B} = \frac{5 - 7c}{5(2 - c)},$$

where $Q^B = \frac{2(2-c)}{3}$ is total output.

Result 1 (comparative statics, Section 2 of the paper). A small reduction in firm 2's cost:

- lowers both prices and raises total output;
- raises firm 2's output and profit;
- lowers firm 1's output and profit.

38.4 Second Step: Welfare under Bertrand Competition

Substituting (38.3) into the definitions gives

$$CS^B = \frac{4(13c^2 - 25c + 25)}{225}, \qquad \Pi^B = \frac{2(53c^2 - 50c + 50)}{225}, \qquad W^B = \frac{2(79c^2 - 100c + 100)}{225}. \tag{38.4}$$

Result 2 (Bertrand competition, Proposition 1 and Corollary 1 of the paper). A small reduction in firm 2's cost:

- always raises consumer surplus;
- lowers total profit if and only if $c > \frac{25}{53}$, that is, if firm 2's output share is below $\frac{2}{9}$;
- lowers total welfare if and only if $c > \frac{50}{79}$, that is, if firm 2's output share is below $\frac{1}{12}$.

Figure 38.1 depicts the gains from a small cost reduction,

$$-\frac{\partial CS^B}{\partial c}, \qquad -\frac{\partial \Pi^B}{\partial c}, \qquad \text{and} \qquad -\frac{\partial W^B}{\partial c}.$$

Consumers always gain, but by less as c rises. Total profit falls once $c > \frac{25}{53}$, and the loss exceeds the consumers' gain once $c > \frac{50}{79}$.

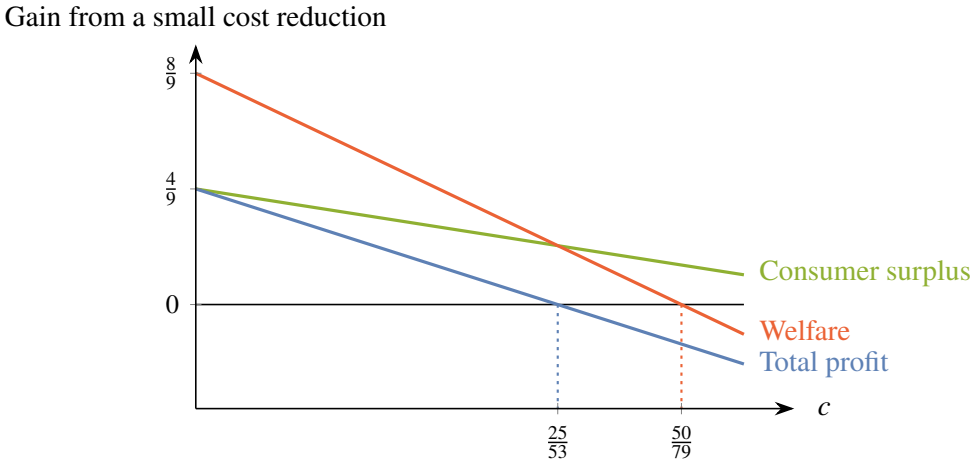


Figure 38.1: Effect of a small reduction in firm 2's cost under Bertrand competition on consumer surplus (green), total profit (blue), and total welfare (red).

Intuitively, the welfare effect of a small cost reduction has three parts:

$$-\frac{\partial W}{\partial c} = \underbrace{q_2}_{\text{cost saving}} \underbrace{-(p_1-0)\frac{\partial q_1}{\partial c}}_{\text{firm 1's lost margin}} \underbrace{+(p_2-c)\left(-\frac{\partial q_2}{\partial c}\right)}_{\text{firm 2's added margin}}. \tag{38.5}$$

Firm 2 saves on each unit it already produces, and it expands at the expense of firm 1. Each unit firm 1 gives up was worth its large markup p_1, while each unit firm 2 adds is worth only its small markup $p_2 - c$. When firm 2 is very inefficient, it produces little, so the cost saving is small, and the shift of output toward the high-cost firm lowers welfare.

38.5 Third Step: Welfare under Cournot Competition

With quantity competition, firm i maximizes $(p_i - c_i)q_i$ using the inverse demands. The first-order conditions

$$4q_1 + q_2 = 3 \qquad \text{and} \qquad q_1 + 4q_2 = 3(1-c)$$

give

$$q_1^C = \frac{3+c}{5}, \quad q_2^C = \frac{3-4c}{5}, \qquad p_1^C = \frac{2(3+c)}{15}, \quad p_2^C = \frac{6+7c}{15}. \tag{38.6}$$

Firm 2 produces if and only if $c < \frac{3}{4}$, its output share is

$$\frac{q_2^C}{Q^C} = \frac{3-4c}{3(2-c)},$$

and

$$CS^C = \frac{13c^2 - 27c + 27}{75}, \qquad \Pi^C = \frac{2(17c^2 - 18c + 18)}{75}, \qquad W^C = \frac{47c^2 - 63c + 63}{75}. \tag{38.7}$$

Result 3 (Cournot competition, Proposition 2 and Corollary 2 of the paper). A small reduction in firm 2's cost:

- always raises consumer surplus;
- lowers total profit if and only if $c > \frac{9}{17}$, that is, if firm 2's output share is below $\frac{1}{5}$;
- lowers total welfare if and only if $c > \frac{63}{94}$, that is, if firm 2's output share is below $\frac{2}{25}$.

Figure 38.2 depicts the three parts of (38.5) under Cournot competition, the cost saving $\frac{3-4c}{5}$, firm 1's lost margin $-\frac{2(3+c)}{75}$, and firm 2's added margin $\frac{8(3-4c)}{75}$, together with their sum, the welfare gain

$$-\frac{\partial W^C}{\partial c} = \frac{63 - 94c}{75}.$$

Intuitively, the three parts of (38.5) work as under Bertrand competition, but markups are larger and the inefficient firm keeps a larger share of the market, so its cost saving matters more and the critical cost for a welfare loss is higher.

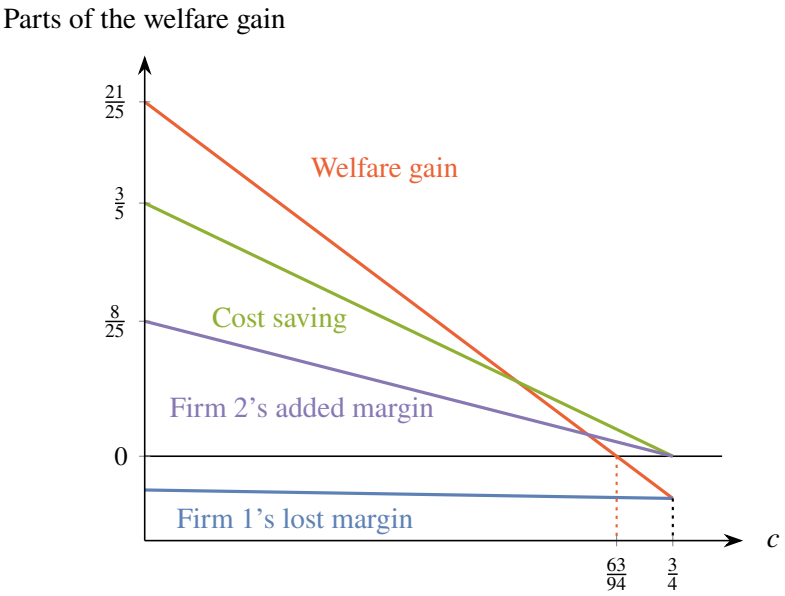


Figure 38.2: The three parts of the welfare gain from a small reduction in firm 2's cost, and their sum, under Cournot competition.

38.6 Fourth Step: Bertrand versus Cournot

Result 4 (comparison, Section 4 and Proposition 3 of the paper). Comparing the two forms of competition:

- total output is lower under Cournot competition, $Q^C - Q^B = -\frac{2-c}{15}$;
- firm 2's output share is lower under Bertrand competition, by $\frac{c}{15(2-c)}$;
- the critical output share for a welfare loss is higher under Bertrand competition, $\frac{1}{12} > \frac{2}{25}$, and the critical cost is lower, $\frac{50}{79} < \frac{63}{94}$;
- hence, if a cost reduction lowers welfare under Cournot competition, it also lowers welfare under Bertrand competition, but not the reverse.

Figure 38.3 depicts the welfare gain from a small cost reduction under both forms of competition. For $\frac{50}{79} < c < \frac{63}{94}$, a cost reduction lowers welfare under Bertrand competition but raises it under Cournot competition. At $c = \frac{13}{20}$, for instance, firm 2's share is $\frac{1}{15}$ under Bertrand competition, below $\frac{1}{12}$, and $\frac{8}{81}$ under Cournot competition, above $\frac{2}{25}$. By (38.5), the welfare gain under Bertrand competition is

$$\frac{3}{50} - \frac{4}{15} \cdot \frac{21}{50} + \frac{14}{15} \cdot \frac{3}{100} = -\frac{3}{125},$$

and the welfare gain under Cournot competition is

$$\frac{2}{25} - \frac{1}{5} \cdot \frac{73}{150} + \frac{4}{5} \cdot \frac{4}{75} = \frac{19}{750}.$$

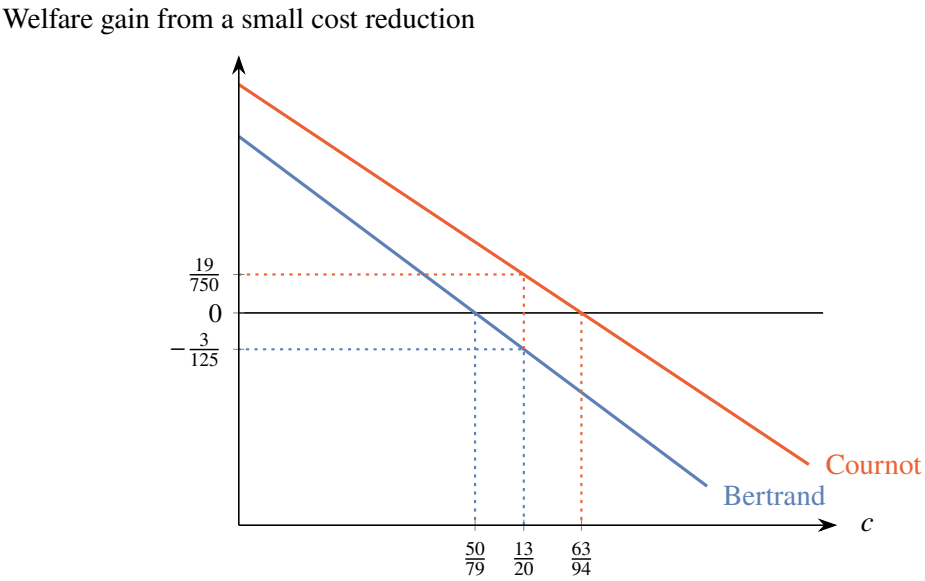


Figure 38.3: Welfare gain from a small reduction in firm 2's cost under Bertrand (blue) and Cournot (red) competition, for c between 0.55 and the exit thresholds.

Intuitively, price competition is tougher on the inefficient firm. Its share is smaller, so the cost saving on its own units is smaller, and the output it takes from the efficient firm weighs more in the welfare comparison.

38.7 Fifth Step: Extension to Many Firms

With n firms, Shubik demand $q_i = V - p_i - \gamma(p_i - \bar{p})$, and costs $c_1, \dots, c_n$, the paper shows that a small reduction in c_i lowers welfare if and only if firm i's output share is below

$$t_W^B = \frac{\gamma\,(n + (n-1)\gamma)}{(2n + (n-1)\gamma)\,(3n + (3n-1)\gamma)} \quad \text{(Bertrand)},$$
$$t_W^C = \frac{\gamma(n+\gamma)}{(2n + (n+1)\gamma)\,(3n + 2\gamma)} \quad \text{(Cournot)}. \tag{38.8}$$

With $n = 2$ and $\gamma = 2$, these are the cutoffs $\frac{1}{12}$ and $\frac{2}{25}$ of Results 2 and 3.

Result 5 (many firms). For any number of firms and any substitutability:

- a small cost reduction always raises consumer surplus and lowers welfare if and only if the firm's share is below the cutoff in (38.8);
- every inefficient firm has a smaller share under Bertrand competition than under Cournot competition;
- if a cost reduction lowers welfare under Cournot competition, it also lowers welfare under Bertrand competition (Proposition 3 of the paper).

Figure 38.4 depicts the cutoffs (38.8) as functions of γ. With two firms (solid), the Bertrand cutoff is above the Cournot cutoff for every γ, and as γ grows the two approach $\frac{1}{5}$ and $\frac{1}{6}$; the dotted lines mark the Bertrand cutoff of the toy model, $\frac{1}{12}$ at $\gamma = 2$. With six firms (dashed), the Bertrand cutoff is above the Cournot cutoff only for $\gamma < 2$.[1]

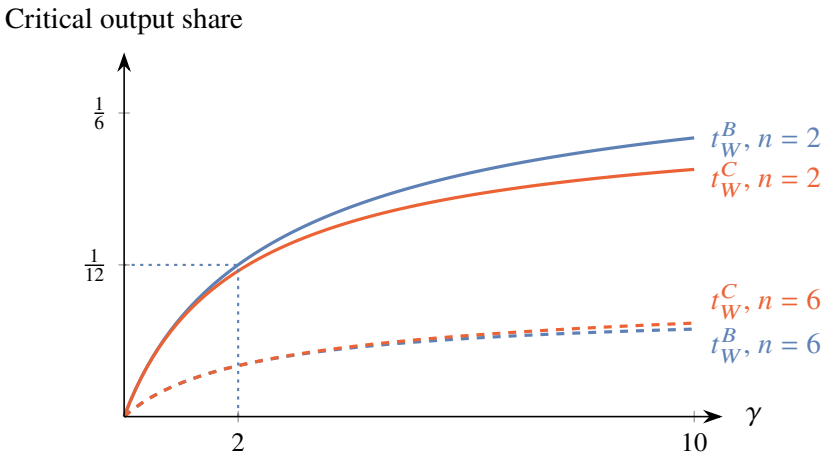


Figure 38.4: Critical output shares below which a small cost reduction lowers welfare, under Bertrand (blue) and Cournot (red) competition, with two firms (solid) and six firms (dashed).

Intuitively, competition in prices is fiercer than competition in quantities, so it leaves an inefficient firm with a smaller share of the market (second bullet). Its share then falls below the welfare cutoff at a lower cost, so a cost reduction that lowers welfare under Cournot competition also lowers it under Bertrand competition.

38.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The welfare-reducing cost reduction, with critical costs and critical output shares under both forms of competition.
- The split of the welfare effect into a cost saving and a reallocation of output toward the inefficient firm.
- The comparison: a cost reduction lowers welfare under Bertrand competition whenever it does so under Cournot competition.

What it drops.

- More than two firms and general substitutability γ, except in the fifth step.

[1] The paper states that $t_W^B > t_W^C$ always. The difference equals $\frac{\gamma^3(n-1)[n-(n-3)\gamma]}{(2n+(n-1)\gamma)(2n+(n+1)\gamma)(3n+(3n-1)\gamma)(3n+2\gamma)}$, which is negative when $n \geq 4$ and $\gamma > \frac{n}{n-3}$. Proposition 3 still holds for every n and γ, because the critical cost of firm i is lower under Bertrand competition in all cases.

- The logit and Hotelling versions of price competition, which the paper discusses briefly.
- Homogeneous Bertrand competition, where only the most efficient firm produces and the perverse effect cannot arise.

38.9 Connections to Other Chapters

- Chapter 11 studies when merger control should accept cost savings as a defense. This chapter shows that a cost reduction by a small firm can lower welfare.
- Chapter 45 studies firms that invest both in cost-reducing R&D and in abatement under an emission fee.
- In Chapter 5, a more efficient entrant can always outbid an off-list-price rebate offered by a less efficient incumbent.

References

Lahiri, S., and Ono, Y. (1988). Helping minor firms reduces welfare. *Economic Journal*, 98(393), 1199–1202.

Wang, X. H., and Zhao, J. (2007). Welfare reductions from small cost reductions in differentiated oligopoly. *International Journal of Industrial Organization*, 25(1), 173–185.

Chapter 39

Patent Length, Forward Protection, and the Timing of R&D

Companion note to Parra, Á. (2019), "Sequential Innovation, Patent Policy, and the Dynamics of the Replacement Effect," RAND Journal of Economics 50(3): 568–590.

39.1 Introduction

A drug company that holds a patent on the best medicine on the market has less to gain from inventing a better one, since the new drug would replace its own profitable product. This is Arrow's (1962) *replacement effect*. How should patent policy reward innovation when the patent holder, the leader, is also a potential source of the next innovation? Parra (2019) answers this question by following the replacement effect over the life of a patent.

The key observation is that the replacement effect changes over time. When the patent is new, it is worth as much as a new patent, so the leader gains nothing from innovating and waits. As the expiration date approaches, the current patent is worth less, and the leader's incentive to innovate grows. Patent policy has two tools. The first is the patent's length. The second is *forward protection*: the probability that a rival's improvement infringes the patent, so that the rival must pay the leader damages. Through these damages, rivals bear part of the leader's loss from being replaced.

Longer patents delay the leader's R&D, because the leader waits longer before its current patent loses value, and forward protection delays the rivals' R&D. The policy that maximizes the pace of innovation therefore has a finite length. It relies more on length when R&D is unproductive, as in pharmaceuticals, and more on forward protection when R&D is productive, as in software.

The toy model is the paper's simplified version in which each rival lives for only an instant, so that the leader's problem has a closed-form solution. The first step sets up the leader's problem and the second step solves the R&D paths over the life of a patent. The third and fourth steps study patent length and forward protection, the fifth step finds the policy that minimizes the expected waiting time between innovations, and the sixth step lists what the toy model keeps and drops.

39.2 The Toy Model

Innovation ladder. Time is continuous and the discount rate is $r = 1$. Firms climb an infinite ladder of innovations. The firm with the latest innovation, the leader, holds a patent and earns a profit flow of 1 while the patent is active; all other firms earn zero. Each new innovation makes the previous technology obsolete.

Patent policy. A patent lasts T years. Let $t \in [0, T]$ be the time since the last innovation. If a rival innovates at $t < T$, with probability $b \in [0, 1]$ (forward protection) its innovation infringes the patent and it pays the leader damages equal to the value the leader loses, v_t. When the patent expires, the leader's profit falls to zero.

R&D. A firm investing x pays $\frac{x^2}{2}$ per unit of time and innovates at rate x. As in Section 4 of the paper, the leader is a long-lived firm, while at each instant it faces a new rival that plays only once (a short-run follower). Let V be the value of a new patent, v_t the leader's value at time t, with $v_0 = V$ in equilibrium and $v_T = 0$.

39.3 First Step: The Leader's Problem

The rival at instant t gains V if it innovates but expects to pay damages bv_t, so it invests

$$x_{f,t} = V - b\, v_t. \tag{39.1}$$

The leader earns the profit flow, invests $x_{l,t}$, gains $V - v_t$ if it innovates first, and loses $(1 - b)v_t$ if the rival innovates first. Its Hamilton-Jacobi-Bellman equation is

$$r v_t = 1 + x_{l,t}\,(V - v_t) - x_{f,t}\,(1 - b)\, v_t - \frac{x_{l,t}^2}{2} + \dot{v}_t, \tag{39.2}$$

so the leader invests

$$x_{l,t} = V - v_t. \tag{39.3}$$

Substituting (39.1) and (39.3) into (39.2) gives a Riccati equation,

$$-\dot{v}_t = \left[\tfrac{1}{2} + b(1 - b)\right] v_t^2 - [1 + (2 - b)V]\, v_t + 1 + \frac{V^2}{2}, \qquad v_T = 0, \tag{39.4}$$

which has a closed-form solution for each V. The value of a new patent then solves the fixed point $v_0(V) = V$, which has a unique solution (Proposition 4 of the paper).

39.4 Second Step: R&D over the Patent Life

Result 1 (R&D dynamics, Theorem 1 of the paper). Over the life of a patent:

- the leader invests nothing when the patent is new, since $v_0 = V$;

- the leader's value falls and its R&D rises as the expiration date approaches;
- the rival's R&D also rises if $b > 0$, and it is constant at V if $b = 0$;
- at expiration, both invest V.

Figure 39.1 depicts both investment paths with $T = 2$ and $b = \frac{1}{2}$, where a new patent is worth $V = 0.765$. The leader's R&D starts at zero and the rival's at $\frac{V}{2} = 0.382$, and both rise to 0.765 at expiration.

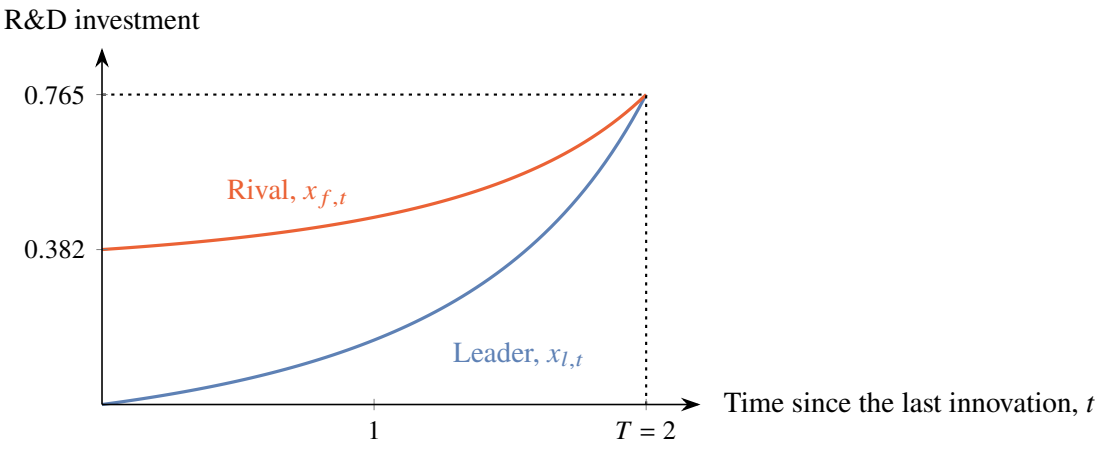


Figure 39.1: R&D of the leader (blue) and of the rival (red) over the life of a patent, with $T = 2$ and $b = \frac{1}{2}$.

Intuitively, the leader's gain from innovating is the new patent minus the patent it already holds, $V - v_t$. Early on, the current patent is worth as much as a new one, so innovating replaces a patent with an identical one and the leader waits. As the patent ages, it is worth less, and the leader's incentive grows. By (39.1) and (39.3), the rival invests $x_{f,t} - x_{l,t} = (1 - b)v_t$ more than the leader: with full forward protection, $b = 1$, the rival pays the leader's entire loss and invests as the leader does, so it internalizes the leader's replacement effect.[1]

39.5 Third Step: Patent Length

Result 2 (patent length, Theorems 5 and 6 and Lemma 3 of the paper). A longer patent:

- raises the value of a new patent and the leader's R&D at expiration, but lowers the leader's R&D during the patent's early life, so it delays the leader's R&D;
- raises the rival's R&D at every date if $b = 0$, since the rival invests V;
- delays the rival's R&D as well if $b = 1$;
- in the limit $T \to \infty$, leaves the leader with no R&D for most of the patent's life.

[1] In the paper's main model, rivals are long-lived and also value the option of innovating later, so with strong forward protection they do not invest early in the patent's life and then invest less than the leader, reversing Arrow's result (Theorem 2 of the paper). Short-run rivals lack this option, so the toy model reaches only the equality at $b = 1$.

Figure 39.2 depicts the leader's R&D with $b = \frac{1}{2}$ and patents of length 1 and 2. With $T = 2$, the leader invests less at every date up to $t = 1$ (0.159 instead of 0.603 at $t = 1$), but more at its own expiration date (0.765).

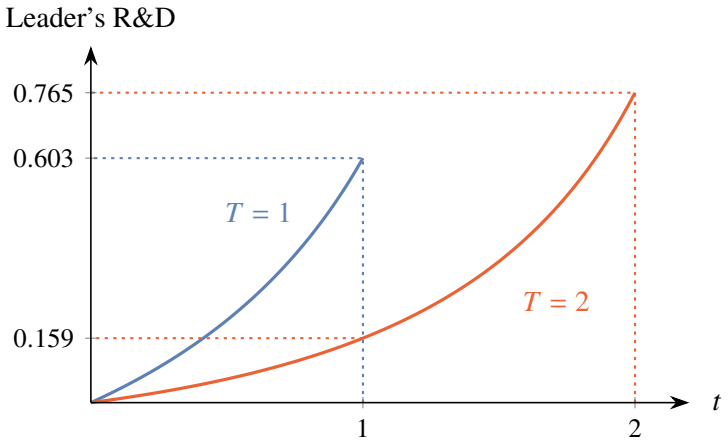


Figure 39.2: Leader's R&D over the patent life with $T = 1$ (blue) and $T = 2$ (red), with $b = \frac{1}{2}$.

Intuitively, a longer patent raises the value of the current patent more when it is close to the old expiration date, since an extension is most likely to be enjoyed by a patent that has already survived. The leader's replacement effect therefore grows at every date, and its R&D moves toward the end of the longer patent. For the rival, a longer patent raises the prize V but also the damages bv_t; with full forward protection the rival invests as the leader does, so it delays too.

39.6 Fourth Step: Forward Protection

Result 3 (forward protection, Theorem 7 of the paper). Raising forward protection from $b = 0$ to $b = \frac{1}{2}$:

- raises the value of a new patent, from 0.522 to 0.603 with $T = 1$ and from 0.602 to 0.765 with $T = 2$;
- lowers the rival's R&D early in the patent's life, from 0.602 to 0.382 at $t = 0$ with $T = 2$;
- raises both firms' R&D at expiration.

Table 39.1 reports the investment paths for both patent lengths and three levels of forward protection.

Intuitively, damages lower the rival's prize, and they matter most early in the patent's life, when the leader has the most to lose. At expiration damages are zero, so only the higher value of a new patent remains, and both firms invest more.

Table 39.1: Value of a new patent and R&D at three dates.

Outcome	Date	$T = 1$			$T = 2$		
		$b = 0$	$b = \frac{1}{2}$	$b = 1$	$b = 0$	$b = \frac{1}{2}$	$b = 1$
Value of a new patent, V		0.522	0.603	0.655	0.602	0.765	0.890
Leader's R&D	$t = \frac{T}{2}$	0.156	0.203	0.228	0.089	0.159	0.195
Rival's R&D	$t = 0$	0.522	0.302	0	0.602	0.382	0
Rival's R&D	$t = \frac{T}{2}$	0.522	0.403	0.228	0.602	0.462	0.195
Both firms' R&D	$t = T$	0.522	0.603	0.655	0.602	0.765	0.890

39.7 Fifth Step: The Pace of Innovation

The two firms innovate at the combined rate $x_{l,t} + x_{f,t}$, and after expiration, as in the paper's approximation, at the rate $2V$ that both firms reach at the expiration date. The policy goal in the paper is to minimize the expected waiting time between innovations, $E[t]$.

Result 4 (optimal patent, Theorem 8 and Result 9 of the paper). The expected waiting time:

- is minimized at a finite, positive patent length for every b: $T = 1.14$ with $b = 0$, $T = 1.04$ with $b = \frac{1}{2}$, and $T = 0.89$ with $b = 1$;
- is minimized overall at $T = 1.10$ and $b = 0.19$, where $E[t] = 1.205$;
- calls for shorter patents and more forward protection when R&D is more productive.

Figure 39.3 depicts the expected waiting time as a function of patent length, with $b = 0$, $b = \frac{1}{2}$, and $b = 1$.

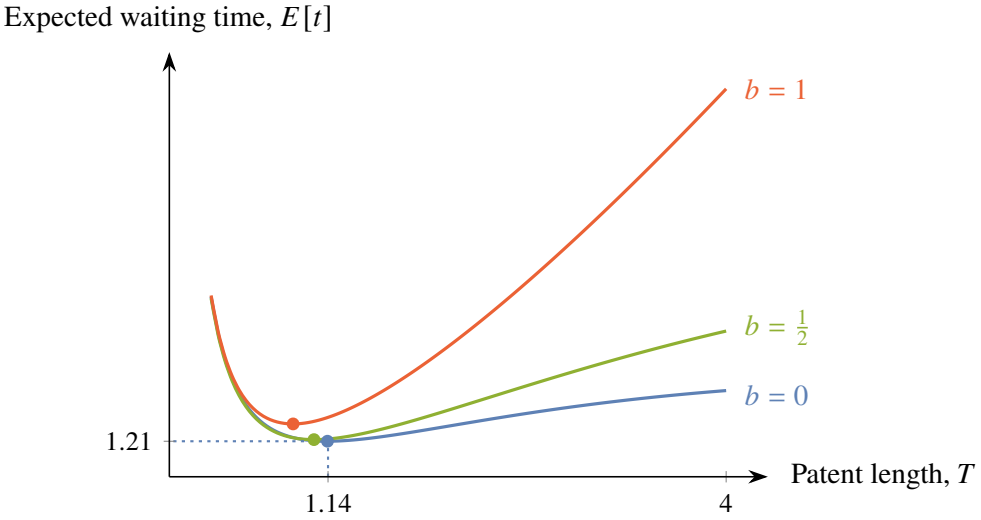


Figure 39.3: Expected waiting time between innovations as a function of patent length, with $b = 0$ (blue), $b = \frac{1}{2}$ (green), and $b = 1$ (red). Dots mark each curve's minimum.

Intuitively, very short patents give little reward and few innovations, and long patents delay R&D; with $b = 1$ the cost of long patents is largest, since both firms delay.

R&D productivity. Let each firm innovate at rate λx instead of x, where λ measures R&D productivity. Table 39.2 reports the policy that minimizes the waiting time.

Table 39.2: Policy that minimizes the expected waiting time, by R&D productivity λ.

Policy and outcome	$\lambda = \frac{1}{2}$	$\lambda = 1$	$\lambda = 2$	$\lambda = 4$
Patent length, T^*	1.92	1.10	0.61	0.33
Forward protection, b^*	0	0.19	0.39	0.52
Expected waiting time	3.20	1.21	0.51	0.23

Intuitively, when R&D is productive, innovations arrive quickly and a patent is unlikely to survive to its statutory end, so extra length adds little reward while it still delays the leader. Forward protection, which pays the leader whenever it is replaced, is then the better reward. When R&D is unproductive, patents are likely to reach their expiration date, so length is an effective reward, and weak forward protection keeps rivals investing while the leader waits. The paper's examples are pharmaceuticals, where innovations are costly and long patents with little forward protection work best, and software, where short, protective patents do.

39.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The nonstationary replacement effect: R&D rises as the patent ages.
- The asymmetric effects of the two tools: length delays the leader, forward protection delays the rival.
- A finite optimal patent length, and the trade-off between length and forward protection across markets.

What it drops.

- Long-lived rivals, whose option value of innovating later produces the reversal of Arrow's result (Theorem 2 of the paper).
- Entry: the paper determines the number of rivals by an entry cost and shows that long, protective patents can reduce entry (Result 10 of the paper).
- The value of competing after expiration, which the toy model sets to zero, and the extensions in Section 6 of the paper (leaders that extend their lead, bargaining over license fees, undiscounted damages).

39.9 Connections to Other Chapters

- In Chapter 15, a merged firm innovates less because an improved product takes sales from its other product, a force similar to the replacement effect of this chapter.
- Chapter 45 studies how environmental policy shapes firms' incentives to invest in R&D.
- In Chapter 40, the replacement effect leads an incumbent that acquires an innovator to develop the project for a smaller range of development costs than the innovator would.

References

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors* (pp. 609–626). Princeton University Press.

Parra, Á. (2019). Sequential innovation, patent policy, and the dynamics of the replacement effect. *RAND Journal of Economics*, 50(3), 568–590.

Chapter 40

Killer Acquisitions: Buying Entrants to Shut Down Their Projects

Companion note to Cunningham, C., F. Ederer, and S. Ma (2021), "Killer Acquisitions," Journal of Political Economy 129(3): 649–702.

40.1 Introduction

Some firms buy a start-up only to shut down its project. Consider a large firm that sells a drug and a start-up that is developing a new drug for the same patients. If the new drug reaches the market, the two drugs compete and prices fall. If the large firm buys the start-up, it decides whether to develop the project or terminate it. When does the buyer terminate the project, and what does this mean for consumers?

Cunningham et al. (2021) call these deals killer acquisitions. The key force is the replacement effect of Arrow (1962): an incumbent gains less from launching a new product than an independent entrant does, because the new product takes sales away from the incumbent's existing product. The entrant has no product to protect, so she counts the new product's full profit. Buying the entrant also protects the incumbent's profit from a new competitor. As a result, the incumbent may buy and terminate projects that the entrant would have developed.

Using data on drug projects, the paper finds that acquired projects that overlap with the buyer's existing drugs are less likely to be developed, and estimates that 5.3 to 7.4 percent of the acquisitions in its sample are killer acquisitions. These deals cluster just below the deal size that triggers antitrust review. In the toy model, killer acquisitions are most likely when the two products overlap at intermediate levels and the market has few incumbents. They hurt consumers, and their low takeover prices keep them below merger review thresholds.

The toy model has one incumbent and one entrant that sell differentiated products and compete in prices. The first step solves the product market, and the second step compares the development decisions of the entrant and of an incumbent that acquired the project. The third step adds existing

competition, the fourth step studies when the incumbent acquires the entrant, the fifth step studies welfare and merger review thresholds, and the sixth step lists what the toy model keeps and drops.

40.2 The Toy Model

Products and demand. A representative consumer buys quantities q_i of several differentiated products to maximize $U(q) - \sum_i p_i q_i$, where

$$U(q) = \sum_i q_i - \frac{1}{2}\Big(\sum_i q_i^2 + \gamma \sum_i \sum_{j \neq i} q_i q_j \Big), \tag{40.1}$$

which yields the linear inverse demand $p_i = 1 - q_i - \gamma \sum_{j \neq i} q_j$ for each product i. The parameter $\gamma \in [0, 1)$ measures product market overlap: products are independent when $\gamma = 0$ and become closer substitutes as γ increases. Consumer surplus is $CS = U(q) - \sum_i p_i q_i$, evaluated at the equilibrium prices and quantities. All products are produced at zero marginal cost. The paper uses the same demand system, following Häckner (2000), and also allows the new product to be of higher quality; we set every quality equal to one.

Firms and the project. An incumbent A sells one existing product. An entrant E owns a project that, if developed at cost $k \geq 0$, yields a new product with probability $\rho \in (0, 1]$ and nothing otherwise. In the fourth step we let the success probability depend on the owner, ρ_A or ρ_E. A terminated project is worth nothing; the paper gives it a liquidation value L, which lowers every development threshold below by L. In the third step the market has $n \geq 1$ incumbents, each with one existing product, and A is one of them.

Timing. In $t = 0$, the incumbent can acquire the entrant. In $t = 1$, the owner of the project, A if it acquired it and E otherwise, decides whether to develop it. In $t = 2$, the outcome of development is observed and the active firms compete in prices. As in the paper, the incumbent makes a take-it-or-leave-it offer, so the takeover price P equals the entrant's expected payoff if she remains independent.

40.3 First Step: Product Market Competition

Three market structures can arise in $t = 2$.

Incumbent alone. If the project is not developed or fails, the incumbent is a monopolist facing $p = 1 - q$. It sets $p = \frac{1}{2}$, sells $q = \frac{1}{2}$, and earns $\pi_1 = \frac{1}{4}$, while consumer surplus is $CS_1 = \frac{1}{8}$.

Duopoly. If the independent entrant develops the project successfully, each firm faces the direct demand

$$q_i = \frac{1 - \gamma - p_i + \gamma p_j}{1 - \gamma^2}.$$

Maximizing $p_i q_i$ gives the best response $p_i = \frac{1-\gamma+\gamma p_j}{2}$, so in the symmetric equilibrium

$$p_2 = \frac{1-\gamma}{2-\gamma}, \qquad q_2 = \frac{1}{(2-\gamma)(1+\gamma)}, \qquad \pi_2 = \frac{1-\gamma}{(2-\gamma)^2(1+\gamma)}, \qquad CS_2 = \frac{1}{(2-\gamma)^2(1+\gamma)}. \tag{40.2}$$

Two-product monopoly. If the incumbent acquired the project and developed it successfully, it sets both prices to maximize $p_1 q_1 + p_2 q_2$. The first-order condition for p_1,

$$1 - \gamma - 2p_1 + 2\gamma p_2 = 0,$$

and its counterpart for p_2 give $p_1 = p_2 = \frac{1}{2}$, so each product sells $\frac{1}{2(1+\gamma)}$ and

$$\Pi = \frac{1}{2(1+\gamma)}, \qquad CS_M = \frac{1}{4(1+\gamma)}. \tag{40.3}$$

Result 1 (product market profits, profit ranking (1) of the paper). For every $\gamma \in (0, 1)$:

- the incumbent's profit satisfies

$$\Pi = \frac{1}{2(1+\gamma)} > \pi_1 = \frac{1}{4} > \pi_2 > 0;$$

- the entrant earns π_2 if her product succeeds and zero otherwise.

Table 40.1 reports prices, profits, and surplus in the three market structures for $\gamma = \frac{1}{2}$. The duopoly price, $\frac{1}{3}$, is lower than the monopoly price, $\frac{1}{2}$, which the two-product monopolist charges for both products.

Table 40.1: Product market outcomes for $\gamma = \frac{1}{2}$.

Market structure	Price	Profit per product	Producer surplus	Consumer surplus	Total surplus
Incumbent alone	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{3}{8}$
Duopoly	$\frac{1}{3}$	$\frac{4}{27}$	$\frac{8}{27}$	$\frac{8}{27}$	$\frac{16}{27}$
Two-product monopoly	$\frac{1}{2}$	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{1}{6}$	$\frac{1}{2}$

Intuitively, the new product takes demand away from the existing one. An independent entrant ignores this loss, which falls on the incumbent, whose profit drops from $\frac{1}{4}$ to π_2. An incumbent that owns both products takes the loss into account and keeps both prices at the monopoly level, so it earns more than with one product, but less than twice its monopoly profit. Only with independent products, $\gamma = 0$, does the second product add $\frac{1}{4}$ to the incumbent's profit.

40.4 Second Step: Development Decisions and Product Overlap

In $t = 1$, the owner of the project compares the expected gain from developing it with the cost k. The entrant earns π_2 with a new product and nothing without it, while the incumbent earns Π with the new product and π_1 without it. Their gains from a successful development are

$$\Delta_E = \pi_2 - 0 = \frac{1-\gamma}{(2-\gamma)^2(1+\gamma)} \qquad \text{and} \qquad \Delta_A = \Pi - \pi_1 = \frac{1-\gamma}{4(1+\gamma)} = \frac{(2-\gamma)^2}{4}\,\Delta_E. \tag{40.4}$$

With a common success probability ρ, the entrant develops the project if $\rho\Delta_E \geq k$ and the incumbent develops it if $\rho\Delta_A \geq k$, which defines the development thresholds

$$k_E = \rho\,\Delta_E \qquad \text{and} \qquad k_A = \rho\,\Delta_A. \tag{40.5}$$

Result 2 (development and overlap, Proposition 1 of the paper). With a common success probability ρ:

- the entrant develops the project if and only if $k \leq k_E$, and an incumbent that acquired it develops it if and only if $k \leq k_A$;
- for every $\gamma \in (0,1)$, $k_A < k_E$, so an acquired project is terminated although the entrant would develop it whenever $k_A < k \leq k_E$, an interval of width

$$k_E - k_A = \rho\,\frac{\gamma(4-\gamma)(1-\gamma)}{4(2-\gamma)^2(1+\gamma)}; \tag{40.6}$$

- the width (40.6) is zero at $\gamma = 0$, tends to zero as $\gamma \to 1$, increases in γ up to $\gamma^* \approx 0.556$, where it reaches 0.0655ρ, and decreases afterwards.

Figure 40.1 depicts the gains Δ_E and Δ_A, the thresholds (40.5) per unit of success probability, as functions of γ. The shaded area between the two curves is the killer-acquisition interval. With $\gamma = \frac{1}{2}$, $\Delta_E = \frac{4}{27} \approx 0.148$ and $\Delta_A = \frac{1}{12} \approx 0.083$; if, in addition, $\rho = \frac{1}{2}$, the entrant develops the project for $k \leq \frac{2}{27}$ but the incumbent only for $k \leq \frac{1}{24}$. With different success probabilities, $k_A < k_E$ if and only if

$$\frac{\rho_A}{\rho_E} < \frac{\Delta_E}{\Delta_A} = \frac{4}{(2-\gamma)^2},$$

which is the form of the condition in the paper.

Intuitively, the incumbent's gain from the new product is smaller because part of the new product's sales come from its own existing product, which it would have sold anyway: this is the replacement effect. The entrant has no product to protect, so she counts the new product's full profit. With independent products there is nothing to replace and both gains equal $\frac{1}{4}$. With nearly identical products, price competition leaves the entrant almost no profit, so neither firm wants to develop the project. The gap between the two incentives, and with it the scope for killer acquisitions, is therefore largest at intermediate levels of overlap.

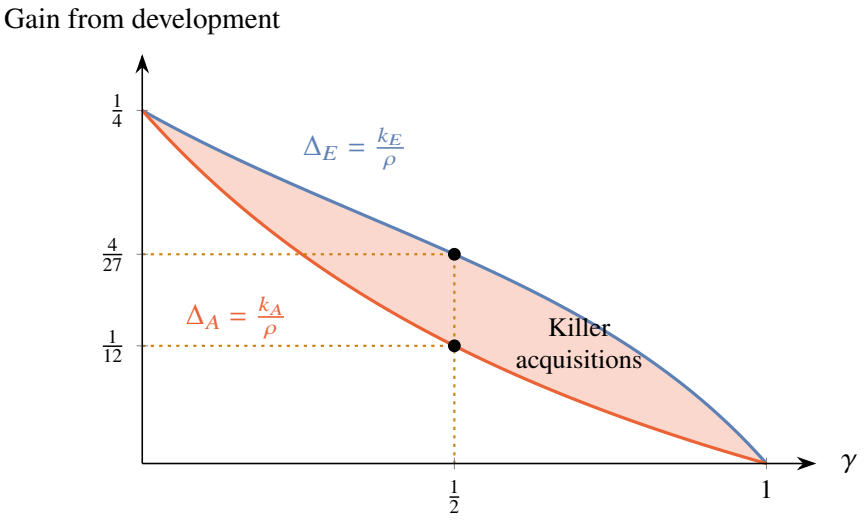


Figure 40.1: Development gains of the entrant and of the incumbent as functions of product overlap.

40.5 Third Step: Existing Competition

Consider now $n \geq 1$ incumbents, each selling one existing product, and let the incumbent A be one of them. With N single-product firms, the symmetric Bertrand equilibrium yields the profit per firm

$$\pi_N = \frac{(1-\gamma)\,[1+(N-2)\gamma]}{[2+(N-3)\gamma]^2\,[1+(N-1)\gamma]}, \tag{40.7}$$

which reduces to $\pi_1 = \frac{1}{4}$ and to π_2 in (40.2). If the incumbent sells two products against $n-1$ single-product rivals, its profit is

$$\Pi_n = \frac{(1-\gamma)\,[1+(n-2)\gamma]\,[2+(2n-1)\gamma]^2}{2(1+n\gamma)\,\{2+(3n-4)\gamma+[1+(n-3)n]\gamma^2\}^2}, \tag{40.8}$$

as in the appendix of the paper. The development gains are now $\Delta_E(n) = \pi_{n+1}$ and $\Delta_A(n) = \Pi_n - \pi_n$. For $\gamma = \frac{1}{2}$ these expressions simplify to

$$\pi_n = \frac{2n}{(n+1)^3}, \qquad \Delta_E(n) = \frac{2(n+1)}{(n+2)^3}, \qquad \Pi_n = \frac{n(2n+3)^2}{(n+2)(n^2+3n+1)^2}. \tag{40.9}$$

Result 3 (development and competition, Proposition 2 of the paper). With $\gamma = \frac{1}{2}$ and a common success probability ρ:

- the width of the killer-acquisition interval, $k_E - k_A = \rho[\Delta_E(n) - \Delta_A(n)]$, strictly decreases in n: it equals $\frac{7}{108}\rho \approx 0.065\rho$ with one incumbent, 0.039ρ with two, 0.023ρ with three, and 0.0035ρ with eight;
- the width tends to zero, and the ratio $\frac{\Delta_A(n)}{\Delta_E(n)}$ tends to one, as $n \to \infty$; this ratio equals $\frac{9}{16}$ with one incumbent.

Figure 40.2 depicts $\Delta_E(n)$ and $\Delta_A(n)$ for $n = 1, \ldots, 8$. The two gains approach each other as the number of incumbents grows. The same pattern holds beyond $\gamma = \frac{1}{2}$: on a grid of γ between 0.02 and 0.98, the width is positive and strictly decreasing in n for every n from 1 to 30.

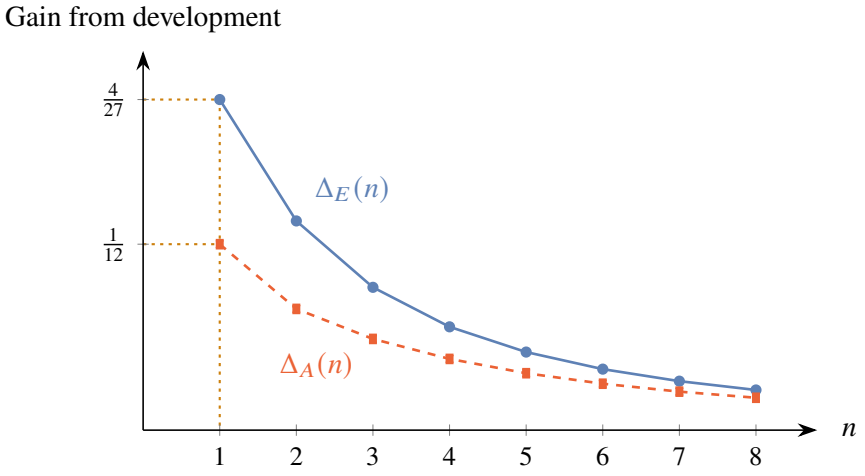


Figure 40.2: Development gains of the entrant, $\Delta_E(n) = \frac{2(n+1)}{(n+2)^3}$, and of the incumbent, $\Delta_A(n) = \Pi_n - \pi_n$, as functions of the number of incumbents, for $\gamma = \frac{1}{2}$.

Intuitively, a new product draws demand away from every existing product. With one incumbent, all of this loss falls on the acquirer's own product. With many incumbents, most of the loss falls on the acquirer's rivals, which the acquirer does not care about, so the replacement effect becomes small and the acquirer's development decision approaches the entrant's.

40.6 Fourth Step: The Acquisition Decision

Takeover price. Return to one incumbent. If the entrant remains independent, she develops the project when $k \leq k_E$, so the takeover price is $P = \max\{\rho_E \Delta_E - k,\, 0\}$. Since the price compensates the entrant for her stand-alone payoff, the incumbent acquires if and only if the joint profit of the two firms is higher with the acquisition.

Joint profits. Consider first a common success probability ρ and $k \leq k_E$. Without an acquisition, the joint expected profit is $2\rho\pi_2 + (1-\rho)\pi_1 - k$. With an acquisition, it is π_1 if the incumbent kills the project and $\rho\Pi + (1-\rho)\pi_1 - k$ if it develops it. Hence the gain from a killer acquisition is

$$\underbrace{\rho\,(\pi_1 - \pi_2)}_{\text{profit protected}} - \underbrace{(\rho\pi_2 - k)}_{\text{value of the project to the entrant}} = k - \rho\left(2\pi_2 - \frac{1}{4}\right), \tag{40.10}$$

where the first term is the efficiency effect of the paper and the second its replacement effect at the acquisition stage. The gain from acquiring and continuing is

$$\rho\,(\Pi - 2\pi_2) = \rho\,\frac{\gamma^2}{2(2-\gamma)^2(1+\gamma)}. \tag{40.11}$$

Since $\Pi > 2\pi_2$,

$$2\pi_2 - \frac{1}{4} < \Pi - \frac{1}{4} = \Delta_A,$$

so (40.10) is positive for every $k > k_A$.

Result 4 (acquisition, Proposition 4 of the paper). With one incumbent and a common success probability ρ:

- if $k > k_E$, no owner develops the project and the incumbent gains nothing from acquiring it;
- if $k_A < k \leq k_E$ and $\gamma > 0$, the incumbent acquires the entrant and terminates the project, and joint profit rises by (40.10);
- if $k \leq k_A$ and $\gamma > 0$, the incumbent acquires the entrant and develops the project, and joint profit rises by (40.11);
- if $\gamma = 0$, the incumbent is indifferent between acquiring and not acquiring.

With n incumbents the comparison is the same, with gain $k - \rho(2\pi_{n+1} - \pi_n)$ from a killer acquisition: it is positive on the whole interval $k_A < k \leq k_E$ because $\Pi_n > 2\pi_{n+1}$, which holds for every n from 1 to 30 on the grid of the third step.

Development synergies. The incumbent may be better or worse at developing the project, $\rho_A \neq \rho_E$. Consider $\rho_E = \frac{1}{2}$ and $k = \frac{2}{27}$, so that, since Δ_E decreases in γ, the entrant develops the project if and only if $\gamma \leq \frac{1}{2}$. In that case, the gain from killing is (40.10) with $\rho = \rho_E$, which is positive if and only if $\pi_2 < \frac{43}{216}$, that is, $\gamma > \gamma_0 \approx 0.234$. Acquiring to continue is strictly better than killing if $\rho_A \Delta_A > k$, that is,

$$\rho_A \;>\; \frac{8(1+\gamma)}{27(1-\gamma)}, \tag{40.12}$$

and strictly better than not acquiring if $\rho_A \Delta_A > \rho_E(2\pi_2 - \frac{1}{4})$, that is,

$$\rho_A \;>\; b(\gamma) \equiv \frac{4 - 8\gamma + 3\gamma^2 - \gamma^3}{2(2-\gamma)^2(1-\gamma)}. \tag{40.13}$$

For $\gamma > \frac{1}{2}$, the entrant would not develop the project, $P = 0$, and the incumbent acquires only to continue, which requires (40.12).

Result 5 (acquisition and development synergies, Figure 2 of the paper). With one incumbent, $\rho_E = \frac{1}{2}$, and $k = \frac{2}{27}$:

- the incumbent acquires to kill if and only if $\gamma_0 < \gamma \leq \frac{1}{2}$ and (40.12) does not hold;
- it acquires to continue if and only if either $\gamma \leq \gamma_0$ and (40.13) holds, or $\gamma_0 < \gamma < \frac{19}{35}$ and (40.12) holds;
- it does not acquire otherwise;
- killer acquisitions occur with $\rho_A > \rho_E$ for every $\gamma \in (\frac{11}{43}, \frac{1}{2}]$, where $\frac{11}{43} \approx 0.256$, and acquisitions to continue occur with $\rho_A < \rho_E$ for every $\gamma \in (0, \frac{11}{43})$.

Figure 40.3 depicts the incumbent's choice for each pair (γ, ρ_A). The dashed line marks $\rho_A = \rho_E = \frac{1}{2}$, and the dot marks the corner $(\frac{1}{2}, \frac{8}{9})$ of the killer-acquisition region.

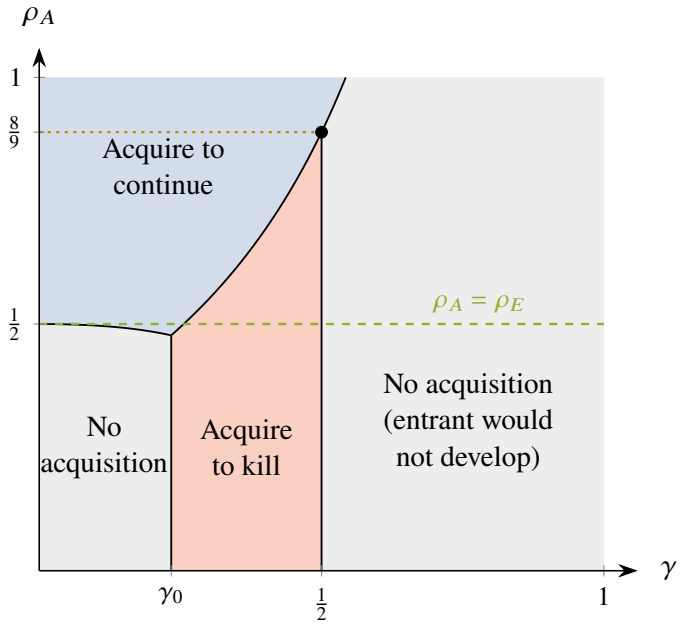


Figure 40.3: The incumbent's acquisition choice as a function of product overlap γ and its own success probability ρ_A, for $\rho_E = \frac{1}{2}$ and $k = \frac{2}{27}$.

Intuitively, two forces drive acquisitions. The first is the profit the incumbent protects by keeping the entrant out, which grows with overlap; the second is the incumbent's own ability to develop the project. With little overlap, the incumbent has little to protect, so it acquires only if it can develop the project well. With more overlap, killing becomes attractive, even for an incumbent that is better at development than the entrant, because developing the project would cannibalize its existing product. When overlap is so large that the entrant would not develop the project, the incumbent faces no threat and buys the project only to develop it.

40.7 Fifth Step: Welfare and Merger Review Thresholds

Welfare. Return to a common success probability ρ. In the killer-acquisition interval, without the acquisition, the entrant develops the project, and with probability ρ the market becomes a duopoly; with the acquisition, the incumbent remains alone. Expected consumer surplus therefore falls by $\rho(CS_2 - CS_1)$, and expected total surplus falls by $\rho(2\pi_2 + CS_2 - \frac{3}{8}) - k$. Since $k \le k_E = \rho\pi_2$ in this interval, the total surplus loss is at least

$$\rho\Big(\pi_2 + CS_2 - \frac{3}{8}\Big) = \rho\Big[\frac{1}{(2-\gamma)(1+\gamma)} - \frac{3}{8}\Big], \tag{40.14}$$

which is positive because $(2-\gamma)(1+\gamma) \le \frac{9}{4} < \frac{8}{3}$. With n incumbents, the total surplus loss is

$$\rho\,(TS_{n+1} - TS_n) - k,$$

where TS_N is total surplus with N single-product firms, and every killer acquisition lowers total surplus if and only if $TS_{n+1} - TS_n > \pi_{n+1}$.

Result 6 (welfare, appendix of the paper). With a common success probability ρ:

- with one incumbent, a killer acquisition lowers expected consumer surplus by $\rho(CS_2 - \frac{1}{8}) > 0$ and expected total surplus by at least (40.14);
- with one incumbent, an acquisition to continue lowers expected consumer surplus by

$$\rho \, \frac{\gamma(4-\gamma)}{4(2-\gamma)^2(1+\gamma)}$$

  and expected total surplus by

$$\rho \, \frac{\gamma(4-3\gamma)}{4(2-\gamma)^2(1+\gamma)};$$

- with two incumbents, every killer acquisition lowers total surplus if and only if $\gamma < \gamma_W \approx 0.716$; with three incumbents, the cutoff falls to 0.515.

With $\gamma = \rho = \frac{1}{2}$ and one incumbent, Table 40.1 gives a consumer surplus loss of

$$\frac{1}{2}\Big(\frac{8}{27} - \frac{1}{8}\Big) = \frac{37}{432} \approx 0.086$$

from a killer acquisition and of

$$\frac{1}{2}\Big(\frac{8}{27} - \frac{1}{6}\Big) = \frac{7}{108} \approx 0.065$$

from an acquisition to continue. Intuitively, a killer acquisition removes a product that consumers value and that would have lowered prices; an acquisition to continue keeps the product but lets the incumbent price both products as a monopolist. When the new product is a close substitute and the market already has several products, much of the entrant's profit consists of sales taken from the incumbents, a transfer rather than a gain for society; this is why, with two or more incumbents, a killer acquisition of a near-copy can raise total surplus. These are also the markets in which, by Result 3, killer acquisitions are least likely to occur.

Takeover prices. Under premerger notification rules, such as the Hart-Scott-Rodino Act in the United States, an acquisition must be notified to the antitrust authorities only if its value exceeds a threshold, which we denote by T. In the toy model, the takeover price is $P = \rho\Delta_E - k$ whenever the entrant would develop the project.

Result 7 (takeover prices and review thresholds). With one incumbent and a common success probability ρ:

- a killer acquisition takes place at a price $P \in [0, \rho(\Delta_E - \Delta_A))$, and an acquisition to continue at a price $P \in [\rho(\Delta_E - \Delta_A), \rho\Delta_E]$;
- every killer acquisition therefore has a lower price than every acquisition to continue, and a threshold $T \geq \rho(\Delta_E - \Delta_A)$ leaves every killer acquisition unreviewed;

- with $\gamma = \rho = \frac{1}{2}$, this bound is $\frac{7}{216} \approx 0.032$, which is $\frac{7}{16}$ of the highest takeover price, $\frac{2}{27}$.

Intuitively, killer acquisitions take place when development is costly, $k > k_A$, and the entrant's project is then worth little to her, so the incumbent buys it for a small price. A review rule based only on deal value thus misses the acquisitions that harm consumers the most: since $CS_M > CS_1$, a killer acquisition lowers consumer surplus by more than an acquisition to continue. The toy model points to two complements: reviewing small deals in which the target's project overlaps with the buyer's products, and paying more attention to markets with few incumbents, where Result 3 shows that the incentive to kill is strongest.

40.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- The replacement effect: an incumbent gains less from a new product that overlaps with its own, so it terminates projects that an independent entrant would develop.
- The role of overlap and of existing competition: the killer-acquisition interval vanishes with independent products and shrinks as the number of incumbents grows.
- The acquisition decision as a joint-profit comparison, in which the profit the incumbent protects is weighed against the value of the project to the entrant, and the roles of development synergies.
- The welfare losses from killer acquisitions and their low takeover prices, which keep them below review thresholds.

What it drops.

- Quality differences between the new and the existing products, and a positive liquidation value; the paper shows that its results hold without quality differences.
- Patent life and generic entry: the paper shows that the gap between the two development thresholds increases with the remaining patent life of the incumbent's product (Proposition 3 of the paper).
- Several incumbents bidding for the entrant, who free ride on each other's acquisitions, and incumbents with different degrees of overlap.
- Quantity competition, under which the paper's propositions hold but an acquisition may not be profitable, and the effect of acquisitions on the creation of new projects.

40.9 Connections to Other Chapters

- Chapter 39 studies Arrow's replacement effect, which also makes the acquirer in this chapter develop the project for a smaller range of development costs than the entrant would.

- Chapter 15 studies a merger between rival innovators, where the merged firm innovates less because an improved product takes sales from its other product.
- Chapter 14 studies how merger review should respond when the authority sees only the merger that the firms propose.

References

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In R. R. Nelson (Ed.), *The Rate and Direction of Inventive Activity: Economic and Social Factors* (pp. 609–626). Princeton, NJ: Princeton University Press.

Cunningham, C., Ederer, F., and Ma, S. (2021). Killer acquisitions. *Journal of Political Economy*, 129(3), 649–702.

Häckner, J. (2000). A note on price and quantity competition in differentiated oligopolies. *Journal of Economic Theory*, 93(2), 233–239.

Part XI

Environmental Regulation and Innovation

The five chapters in this part study how environmental policy and firms' investments in cleaner technology affect each other. Chapter 41 compares instruments: emission taxes, auctioned permits, free permits, and standards, ranked by the incentive they give firms to adopt an advanced abatement technology. A firm that does not adopt free rides on the fall in the price of permits caused by the other firms' adoption, so not all firms adopt. Chapters 42 through 45 fix the instrument, an emission tax, and study firms' investments in R&D. Chapter 42 compares independent research with an environmental R&D cartel when the regulator sets the tax after observing abatement, so firms invest partly to lower the tax. Chapter 43 keeps the emissions, spillovers, and R&D costs of Chapter 42, but the regulator commits to the tax before firms invest, and the number of firms varies. Without spillovers, total R&D rises with the number of firms; with spillovers, it first rises and then falls, an inverted U. Chapter 44 moves to two exporting countries. Governments again set their taxes before firms invest, but each chooses its tax to favor its own exporting firm, and a research joint venture can then lower each exporting country's welfare, and even the welfare of the three countries combined. Chapter 45 returns to a regulator that sets the tax after firms invest and adds a second investment, cost-reducing R&D. Because R&D raises output and leads the regulator to raise the tax, while abatement leads it to lower the tax, the two investments become strategic complements.

What connects the five chapters is the timing of policy. Whether the government commits to its policy before firms invest (Chapters 43 and 44) or responds to their investment (Chapters 42 and 45) changes how firms invest, and Chapter 41 shows that it also changes the ranking of instruments.

Chapter 41

Taxes, Permits, and the Adoption of Advanced Abatement Technology

Companion note to Requate, T., and W. Unold (2003), "Environmental Policy Incentives to Adopt Advanced Abatement Technology: Will the True Ranking Please Stand Up?," European Economic Review 47(1): 125–146.

41.1 Introduction

New abatement technologies make it cheaper to cut pollution. Consider a coal power plant deciding whether to install a scrubber, which lowers its abatement costs but has an installation cost. Its decision depends on the policy it faces: an emission tax, tradable permits sold at auction, tradable permits handed out for free, or an emission standard. Which policy gives firms the strongest incentive to adopt the new technology?

The key force is how the price of emissions responds to adoption. Under a tax, the price is fixed, so a firm's gain from adopting does not depend on what other firms do. Under permits, every adoption lowers the demand for permits and their price. A firm that does not adopt enjoys this lower price without paying the adoption cost: it free rides on the adoption of others. Earlier studies, which measured the cost savings of the whole industry when all firms adopt, missed this free riding and ranked auctioned permits first.

Requate and Unold (2003) show that, when the regulator keeps the policy it chose before the new technology appeared, taxes give stronger incentives than permits, auctioned and free permits give the same incentives, and standards may give stronger incentives than permits. A stronger incentive is not better policy, however, since firms may adopt too much. When the regulator anticipates the new technology, taxes and permits both induce the socially optimal number of adopters if it sets its policy after firms adopt; if it commits before, only permits are sure to do so.

The toy model has a continuum of identical firms, quadratic abatement costs, and quadratic damage. The first step finds the first-best adoption level, the second step compares the incentives of each

instrument under the old policy, and the third step finds the equilibrium number of adopters. The fourth step studies a regulator who anticipates the new technology, the fifth step considers linear damage, and the sixth step lists what the toy model keeps and drops.

41.2 The Toy Model

Firms and technologies. A continuum of identical firms, with mass one, emits a pollutant. Without regulation, each firm emits one unit. A firm that emits $e \in [0, 1]$ bears the abatement cost

$$C_O(e) = (1-e)^2 \quad \text{with the old technology,} \qquad C_N(e) = \frac{(1-e)^2}{2} \quad \text{with the new technology,} \tag{41.1}$$

so the marginal abatement costs are $-C_O'(e) = 2(1-e)$ and $-C_N'(e) = 1-e$: the new technology halves the marginal abatement cost at every emission level. Adopting the new technology costs $k = \frac{4}{49} \approx 0.082$ per firm, a value at which, as the first step shows, the first best has half of the firms adopting. We denote by $\alpha \in [0, 1]$ the share of firms that adopt. The paper considers n small firms with general convex abatement costs and an adoption cost F. When firms are numerous, it replaces its two equilibrium conditions with a single indifference condition, its equation (6); with a continuum of firms, this condition holds without approximation.

Damage. Aggregate emissions are $E = \alpha e_N + (1-\alpha) e_O$, where e_N and e_O are the emissions of each adopter and each non-adopter. Environmental damage is $D(E) = \frac{E^2}{2}$, with marginal damage $D'(E) = E$, in the first to fourth steps, which satisfies the paper's assumption of increasing and strictly convex damage (its Section 4); the fifth step departs from the paper and considers the linear damage $D(E) = \frac{2}{3}E$. Social cost is the sum of adoption costs, abatement costs, and damage,

$$SC = \alpha k + \alpha\, C_N(e_N) + (1-\alpha)\, C_O(e_O) + D(E). \tag{41.2}$$

Instruments. The regulator uses one of four instruments: an emission tax t per unit; L tradable permits sold at auction, which trade at a competitive price p; L tradable permits handed out for free, L to each firm, which also trade at price p; or an emission standard $\bar{e}$ that every firm must meet. Each firm is small relative to the market and takes the price of emissions as given, as in the competitive setting of the paper.

Timing. In the first stage, each firm decides whether to adopt the new technology. In the second stage, firms choose their emissions and, under permits, the permit market clears. The paper first studies the setting of Milliman and Prince (1989) and Jung et al. (1996), in which the regulator does not anticipate the new technology and keeps its policy after the new technology appears; as a benchmark, its Section 5 considers the policy that is optimal with the old technology, the case the toy model studies. Its Section 6 then studies two scenarios in which the regulator anticipates the new technology: the regulator commits to its policy before firms adopt (Timing A, ex ante optimal policy), or it chooses its instrument, observes how many firms adopt, and then sets the level of its instrument (Timing B, ex post optimal policy).

Emissions and the private gain from adoption. Facing a price t per unit of emissions, a firm with the old technology minimizes $C_O(e) + t\,e$ and emits $e_O = 1 - \frac{t}{2}$, while a firm with the new

technology emits $e_N = 1 - t$. Their costs, including payments for emissions, are $t - \frac{t^2}{4}$ and $t - \frac{t^2}{2}$, so a firm's gain from adopting is

$$g(t) = \frac{t^2}{4}. \tag{41.3}$$

The gain increases in the price of emissions: the more expensive emissions are, the more valuable a lower marginal abatement cost becomes. With a share α of adopters and a common price t, aggregate emissions are

$$E = 1 - \frac{t\,(1+\alpha)}{2}. \tag{41.4}$$

41.3 First Step: The First-Best Adoption Level

Emissions for a given share of adopters. For a given α, the planner allocates abatement so that every firm's marginal abatement cost equals a common price t, and sets this price equal to marginal damage. From (41.4) and $t = E$,

$$E(\alpha) = t(\alpha) = \frac{2}{3+\alpha}. \tag{41.5}$$

Substituting into (41.2), the lowest social cost for a given share of adopters is

$$SC(\alpha) = \alpha k + \frac{1}{3+\alpha}. \tag{41.6}$$

Optimal adoption. The derivative of (41.6) is

$$SC'(\alpha) = k - \frac{1}{(3+\alpha)^2} = k - g\big(t(\alpha)\big), \tag{41.7}$$

and

$$SC''(\alpha) = \frac{2}{(3+\alpha)^3} > 0.$$

The marginal social benefit of one more adopter thus equals the private gain (41.3) of an adopter evaluated at the optimal emission price for that share of adopters.

Result 1 (first-best adoption). The first-best share of adopters is

$$\alpha^* = \begin{cases} 1 & \text{if } k \le \frac{1}{16}, \\ \dfrac{1}{\sqrt{k}} - 3 & \text{if } \frac{1}{16} < k < \frac{1}{9}, \\ 0 & \text{if } k \ge \frac{1}{9}. \end{cases}$$

With $k = \frac{4}{49}$:

- $\alpha^* = \frac{1}{2}$, and the optimal emission price and aggregate emissions are $t^* = E^* = \frac{4}{7}$;
- each adopter emits $\frac{3}{7}$, each non-adopter emits $\frac{5}{7}$, and $SC^* = \frac{16}{49} \approx 0.327$, below $SC(0) = \frac{1}{3}$ and $SC(1) = \frac{65}{196} \approx 0.332$.

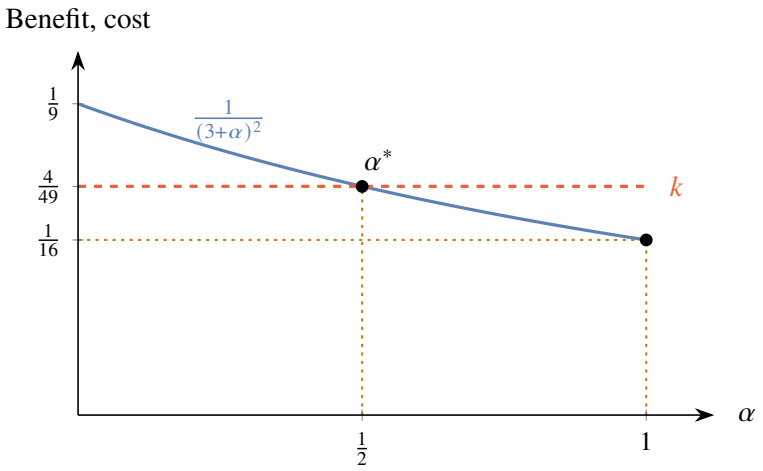


Figure 41.1: Marginal social benefit of adoption and adoption cost $k = \frac{4}{49}$.

Figure 41.1 depicts the marginal social benefit of adoption, $\frac{1}{(3+\alpha)^2}$, and the adoption cost k. The first best is at their intersection.

Intuitively, each adoption lowers the optimal emission price, because the industry can now abate at a lower cost and the planner asks for more abatement. A lower emission price, in turn, reduces the value of a lower marginal abatement cost for the next adopter. The benefit of adoption therefore falls as more firms adopt, and with $k = \frac{4}{49}$ the planner stops when half of the firms have adopted. Identical firms are thus treated differently in the first best: half of them adopt and the other half keep the old technology. Result 1 is the toy-model version of Proposition 2 of the paper: no firm should adopt when the adoption cost is high, all firms should adopt when it is low, and, in between, the optimal number of adopters decreases in the adoption cost.

41.4 Second Step: Adoption Incentives under a Committed Policy

The old policy. Before the new technology appears, $\alpha = 0$ and (41.5) gives the optimal emission price and aggregate emissions $t_0 = E_0 = \frac{2}{3}$, with each firm emitting $\frac{2}{3}$. Four policies implement this outcome: the tax $t_0 = \frac{2}{3}$; $L_0 = \frac{2}{3}$ auctioned permits; $\frac{2}{3}$ free permits to each firm; and the standard $\bar{e} = \frac{2}{3}$. Consider that the regulator keeps its policy after the new technology appears.

Tax. A firm's gain from adopting is $g(t_0) = \frac{1}{9}$, regardless of the share of other firms that adopt.

Auctioned permits. With a share α of adopters, the permit market clears when (41.4) equals $L_0 = \frac{2}{3}$, so the permit price is

$$p(\alpha) = \frac{2}{3(1+\alpha)}, \tag{41.8}$$

which falls from $\frac{2}{3}$ to $\frac{1}{3}$ as the share of adopters rises from 0 to 1. A firm's gain from adopting is $g\big(p(\alpha)\big) = \frac{1}{9(1+\alpha)^2}$.

Free permits. A firm that receives $\frac{2}{3}$ permits for free bears $C(e) + p\left(e - \frac{2}{3}\right)$, which differs from its cost under auctioned permits by the value of its endowment, $\frac{2p}{3}$. The firm receives this value whether or not it adopts, so its gain from adopting is again $\frac{1}{9(1+\alpha)^2}$. Proposition 1 of the paper makes this argument for general abatement costs.

Standard. A firm must emit $\frac{2}{3}$ with either technology, so its gain from adopting is

$$C_O\left(\frac{2}{3}\right) - C_N\left(\frac{2}{3}\right) = \frac{1}{9} - \frac{1}{18} = \frac{1}{18},$$

regardless of the share of other firms that adopt.

Result 2 (individual adoption incentives). When the regulator keeps the policy that is optimal with the old technology, a firm's gain from adopting, given a share α of adopters, is

$$g_T = \frac{1}{9} \text{ (tax)}, \qquad g_P(\alpha) = \frac{1}{9(1+\alpha)^2} \text{ (auctioned or free permits)}, \qquad g_S = \frac{1}{18} \text{ (standard)}. \tag{41.9}$$

- $g_T \geq g_P(\alpha)$, with equality only at $\alpha = 0$, and $g_T > g_S$;
- auctioned and free permits give the same incentive;
- $g_S > g_P(\alpha)$ if and only if $\alpha > \sqrt{2} - 1 \approx 0.414$.

Figure 41.2 depicts the three incentives in (41.9) together with the adoption cost k. The permit incentive starts at the tax incentive when no other firm has adopted and falls to $\frac{1}{36}$ when all firms have adopted, crossing the standard's incentive at $\alpha = \sqrt{2} - 1$.

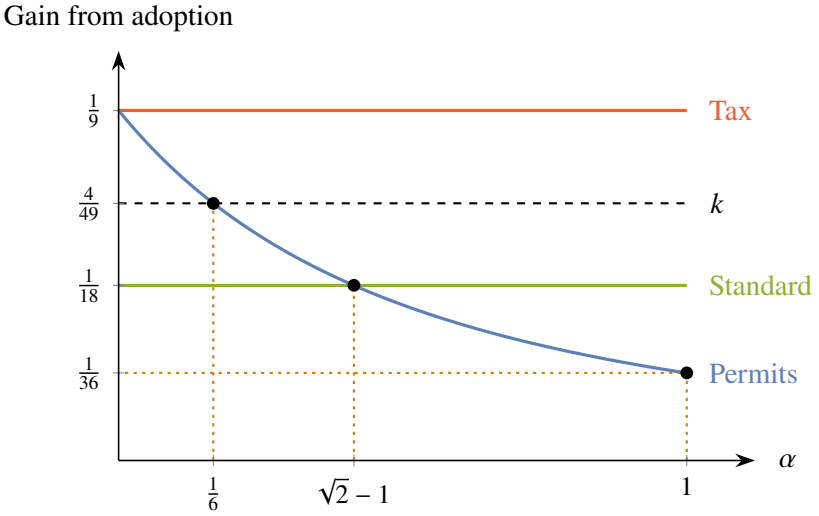


Figure 41.2: Gain from adoption under each instrument when the regulator keeps the policy that is optimal with the old technology, as a function of the share α of other adopters.

Intuitively, under a tax the price of emissions does not depend on what other firms do, so each firm's gain from adopting is the same no matter how many firms adopt. Under permits, every

adoption lowers the demand for permits and their price, and a lower price reduces the value of a lower marginal abatement cost. A firm that does not adopt benefits from the lower price created by other firms' adoption without paying the adoption cost. Whether permits are auctioned or given away does not matter, since the free endowment is a lump-sum transfer that does not depend on the firm's technology. A standard gives a smaller gain than the tax, because the firm cannot use the new technology to cut its emissions and its tax bill; but, unlike permits, the standard's gain does not fall as other firms adopt.

41.5 Third Step: Equilibrium Adoption and the Industry-Wide Ranking

Equilibrium. A share α of adopters is an equilibrium if no firm wants to change its decision: an interior share requires that the gain from adopting equal k; $\alpha = 1$ requires a gain of at least k when all firms adopt; and $\alpha = 0$ requires a gain of at most k when no firm adopts. Under the tax, $g_T = \frac{1}{9} > \frac{4}{49}$, so all firms adopt. Under permits,

$$\frac{1}{9(1+\alpha)^2} = \frac{4}{49}$$

gives $1 + \alpha = \frac{7}{6}$, so one sixth of the firms adopt and the permit price is $p\left(\frac{1}{6}\right) = \frac{4}{7}$. More generally, for $\frac{1}{36} < k < \frac{1}{9}$, the share of adopters under permits is $\frac{1}{3\sqrt{k}} - 1$. As in Lemma 2 of the paper, with partial adoption the permit price equals the price at which firms are indifferent between the two technologies. Under the standard, $g_S = \frac{1}{18} < \frac{4}{49}$, so no firm adopts.

Result 3 (equilibrium adoption under a committed policy). When the regulator keeps the policy that is optimal with the old technology:

- the equilibrium shares of adopters are $\alpha_T = 1$ (tax), $\alpha_P = \frac{1}{6}$ (auctioned or free permits), and $\alpha_S = 0$ (standard), so that
$$\alpha_T > \alpha^* = \frac{1}{2} > \alpha_P > \alpha_S;$$
- social costs are $\frac{317}{882} \approx 0.359$ (tax), $\frac{146}{441} \approx 0.331$ (permits), and $\frac{1}{3} \approx 0.333$ (standard), all above $SC^* = \frac{16}{49} \approx 0.327$.

Figure 41.3 depicts the permit market. The aggregate marginal abatement cost with a share α of adopters is $\frac{2(1-E)}{1+\alpha}$, which pivots down as more firms adopt. With the cap $L_0 = \frac{2}{3}$ fixed, the permit price falls from point A ($\alpha = 0$) to point B ($\alpha = \frac{1}{6}$), where adoption stops, and would fall to point C if all firms adopted. Under the tax, the price stays at $\frac{2}{3}$, all firms adopt, and emissions fall to $\frac{1}{3}$ (point T).

Result 3 is the toy-model version of Propositions 3 to 5 of the paper. With an adoption cost between the two cutoffs of Result 1, $\frac{1}{16} < k < \frac{1}{9}$, the old tax leads to over-investment and the old cap to under-investment. The old standard makes all firms adopt if $k < \frac{1}{18}$ and no firm adopt if $k > \frac{1}{18}$, so with $k = \frac{4}{49}$ it also leads to under-investment.

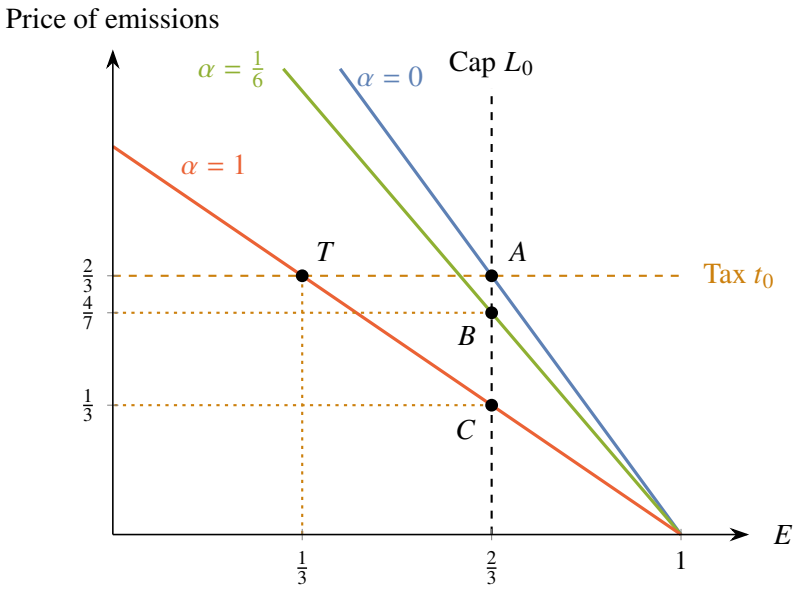


Figure 41.3: Aggregate marginal abatement cost for different shares α of adopters, the permit cap $L_0 = \frac{2}{3}$, and the tax $t_0 = \frac{2}{3}$.

The industry-wide ranking. Earlier studies, such as Milliman and Prince (1989) and Jung et al. (1996), measure an instrument's incentive by the cost savings of the whole industry when all firms adopt, and rank auctioned permits first, taxes second, free permits third, and standards last (Section 1 of the paper). Table 41.1 reports this measure, the saving in abatement costs and payments for emissions when all firms adopt rather than none, next to each firm's own gain from adopting. Under auctioned permits, the industry's cost falls from $\frac{5}{9}$ to $\frac{5}{18}$, a saving of $\frac{5}{18}$: each firm still emits $\frac{2}{3}$, so its abatement cost and its permit bill fall by

$$\frac{1}{9} - \frac{1}{18} = \frac{1}{18} \qquad \text{and} \qquad \left(\frac{2}{3} - \frac{1}{3}\right)\frac{2}{3} = \frac{2}{9},$$

respectively. Under the tax, the industry's cost falls from $\frac{5}{9}$ to $\frac{4}{9}$, a saving of $\frac{1}{9}$. Under free permits and the standard, the saving is $\frac{1}{18}$. The toy model thus reproduces the earlier ranking, except that free permits and the standard tie.

Intuitively, the industry-wide measure ranks auctioned permits first because it counts the fall in the permit bill, $\frac{2}{9}$, as a gain from adoption. This fall is not a saving for society: it is revenue that the regulator no longer collects. Nor can a single firm obtain it by adopting, since each firm is too small to move the permit price; a firm that does not adopt pays the same lower price. Under free permits, firms buy and sell permits only among themselves, so these payments sum to zero for the industry at every price, which is why the industry-wide measure separates auctioned from free permits while firms respond to both in the same way. Once the comparison uses each firm's own gain, the tax gives the strongest incentive. A stronger incentive, however, is not better policy: the tax $\frac{2}{3}$ is above the first-best price $\frac{4}{7}$, so all firms adopt and abate too much, while the cap $\frac{2}{3}$ is above the first-best emissions $\frac{4}{7}$, so the permit price falls too fast and too few firms adopt. In the toy model, permits come closest to the first best; the paper shows that, in general, which instrument comes closest depends on the parameters.

Table 41.1: Industry-wide savings, individual gains from adoption, equilibrium adoption, and social cost when the regulator keeps the policy that is optimal with the old technology.

Instrument	Industry-wide saving	Gain, $\alpha = 0$	Gain, $\alpha = 1$	Equilibrium α	Social cost
Tax	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	1	$\frac{317}{882} \approx 0.359$
Auctioned permits	$\frac{5}{18}$	$\frac{1}{9}$	$\frac{1}{36}$	$\frac{1}{6}$	$\frac{146}{441} \approx 0.331$
Free permits	$\frac{1}{18}$	$\frac{1}{9}$	$\frac{1}{36}$	$\frac{1}{6}$	$\frac{146}{441} \approx 0.331$
Standard	$\frac{1}{18}$	$\frac{1}{18}$	$\frac{1}{18}$	0	$\frac{1}{3} \approx 0.333$
First best	–	–	–	$\frac{1}{2}$	$\frac{16}{49} \approx 0.327$

41.6 Fourth Step: A Regulator Who Anticipates the New Technology

Ex post optimal policy. Consider now that the regulator sets its policy after firms adopt, as in Timing B of the paper. Observing a share α of adopters, it sets the tax $t(\alpha) = \frac{2}{3+\alpha}$ from (41.5), or issues $L(\alpha) = \frac{2}{3+\alpha}$ permits, auctioned or free, which trade at the same price. Firms anticipate this response, so the gain from adopting is $g\big(t(\alpha)\big) = \frac{1}{(3+\alpha)^2}$ under the tax and under both kinds of permits. This gain decreases in α, and it equals $k = \frac{4}{49}$ at $\alpha = \frac{1}{2}$.

Ex ante optimal policy with permits. Consider instead that the regulator knows the new technology and commits to the first-best cap $L^* = \frac{4}{7}$ before firms adopt, as in Timing A of the paper. The permit market clears at

$$p(\alpha) = \frac{6}{7(1+\alpha)}, \tag{41.10}$$

so the gain from adopting is $\frac{9}{49(1+\alpha)^2}$, which falls from $\frac{9}{49}$ at $\alpha = 0$ to $\frac{9}{196}$ at $\alpha = 1$ and equals k only at $\alpha = \frac{1}{2}$.

Ex ante optimal policy with a tax. If the regulator commits to the first-best tax $t^* = \frac{4}{7}$, the gain from adopting is $g\big(\frac{4}{7}\big) = \frac{4}{49} = k$ for every α, so every firm is indifferent. With a share α of adopters, aggregate emissions are $\frac{5-2\alpha}{7}$ and social cost is

$$SC(\alpha; t^*) = \frac{33 - 4\alpha(1-\alpha)}{98}, \tag{41.11}$$

which equals $\frac{16}{49}$ at $\alpha = \frac{1}{2}$ and $\frac{33}{98}$ at $\alpha = 0$ and $\alpha = 1$. A tax above $\frac{4}{7}$ makes all firms adopt, and a tax below $\frac{4}{7}$ makes no firm adopt. With all firms adopting, social cost is

$$k + \frac{t^2}{2} + \frac{(1-t)^2}{2},$$

which increases in t for $t > \frac{1}{2}$; with no firm adopting, it is

$$\frac{t^2}{4} + \frac{1}{2}\Big(1 - \frac{t}{2}\Big)^2,$$

which decreases in t for $t < \frac{2}{3}$. Both expressions equal $\frac{33}{98}$ at $t = \frac{4}{7}$.

Result 4 (ex post optimal policy). If the regulator sets its policy after firms adopt, the tax, auctioned permits, and free permits all induce the unique equilibrium share of adopters $\alpha = \frac{1}{2} = \alpha^*$, with social cost $SC^* = \frac{16}{49}$.

Result 5 (ex ante optimal policy). If the regulator commits to its policy after the new technology becomes known but before firms adopt:

- with the cap $L^* = \frac{4}{7}$, auctioned or free, the unique equilibrium share of adopters is $\alpha = \frac{1}{2} = \alpha^*$;
- with the tax $t^* = \frac{4}{7}$, every $\alpha \in [0, 1]$ is an equilibrium, with social cost (41.11);
- with any tax $t \neq \frac{4}{7}$, the equilibrium is $\alpha = 1$ (if $t > \frac{4}{7}$) or $\alpha = 0$ (if $t < \frac{4}{7}$), and social cost is above $\frac{33}{98} > \frac{16}{49}$.

Figure 41.4 depicts the gain from adopting in the three cases. The curves for ex post optimal policy and for the ex ante cap both cross the adoption cost at $\alpha^* = \frac{1}{2}$, from above, while the gain under the ex ante tax coincides with the adoption cost for every α.

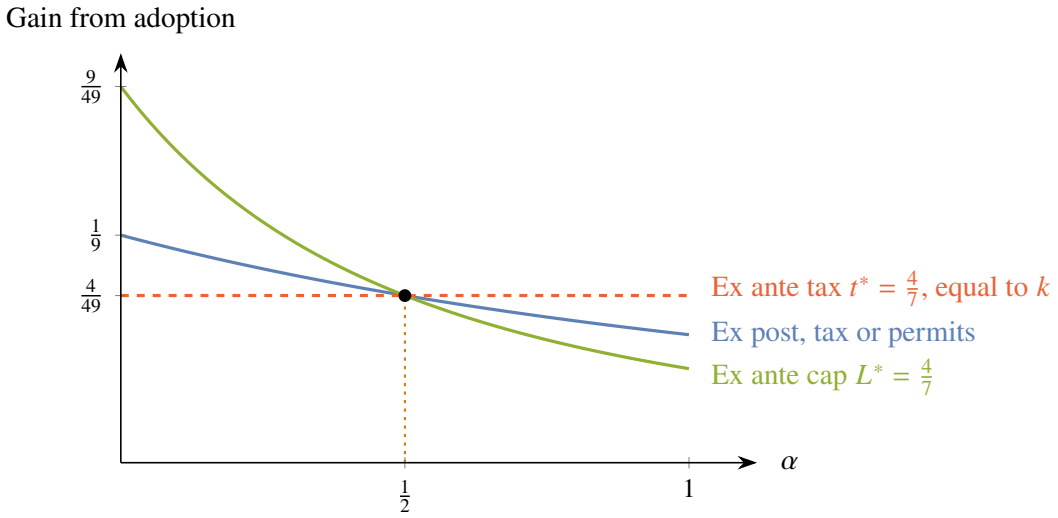


Figure 41.4: Gain from adoption when the regulator anticipates the new technology.

Intuitively, implementing the first best with identical firms requires that the gain from adopting fall as more firms adopt, so that adoption stops at α^*. A fixed cap generates this feedback through the permit price. When the regulator moves after adoption, its response generates the same feedback under the tax and under permits, since it sets a lower price of emissions, through the tax or through the cap, when more firms adopt. A tax set before adoption generates no feedback: either it leaves firms indifferent, and nothing selects the first-best share among all possible shares, or it pushes every firm to the same decision. In the toy model, the cost of an all-or-none outcome is

$$\frac{33}{98} - \frac{16}{49} = \frac{1}{98},$$

about 3% of the first-best social cost. Under the ex ante cap, the share of adopters is unique, although identical firms still need to coordinate on which of them adopt. Results 4 and 5 are the toy-model versions of Propositions 7 and 6 of the paper; Proposition 6 also shows that a tax set before adoption reaches the first best when no firm or all firms should adopt.

41.7 Fifth Step: Linear Damage

Consider now the linear damage $D(E) = \frac{2}{3}E$, so that marginal damage is constant at $\frac{2}{3}$. The paper assumes strictly convex damage, so this step departs from it to show the role of convexity. Before the new technology appears, setting the marginal abatement cost $2(1 - E)$ equal to $\frac{2}{3}$ gives $E_0 = \frac{2}{3}$, so the old policies of the second step are still optimal: the tax $\frac{2}{3}$, the cap $\frac{2}{3}$, and the standard $\frac{2}{3}$. For any share of adopters, the optimal price of emissions is $\frac{2}{3}$, and social cost is

$$SC(\alpha) = \frac{5}{9} + \alpha\left(k - \frac{1}{9}\right), \tag{41.12}$$

which decreases in α because $k < g\left(\frac{2}{3}\right) = \frac{1}{9}$. All firms adopt in the first best.

Result 6 (linear damage). With $D(E) = \frac{2}{3}E$:

- the first best has $\alpha^* = 1$, $E^* = \frac{1}{3}$, and $SC^* = \frac{232}{441} \approx 0.526$;
- if the regulator keeps the policy that is optimal with the old technology, the tax induces $\alpha = 1$ and the first best, permits induce $\alpha = \frac{1}{6}$ with social cost $\frac{244}{441} \approx 0.553$, and the standard induces $\alpha = 0$ with social cost $\frac{5}{9} \approx 0.556$;
- if the regulator anticipates the new technology, the tax $\frac{2}{3}$ and the cap $\frac{1}{3}$ both induce the first best, whether set before or after adoption.

With the ex ante cap $\frac{1}{3}$, the permit price is $\frac{4}{3(1+\alpha)}$, which is at least $\frac{2}{3}$ for every $\alpha \leq 1$, so the gain from adopting is at least $\frac{1}{9} > k$ and all firms adopt. Under ex post optimal policy, the regulator issues the permits that bring the price to $\frac{2}{3}$, and all firms again adopt.

Intuitively, when marginal damage is constant, the optimal price of emissions does not depend on the technology firms use. A tax set for the old technology is therefore still the right tax after the new technology appears, and it gives every firm the right incentive to adopt. A cap, instead, fixes the quantity of emissions, but the optimal quantity falls when the new technology arrives; with the old cap, the permit price falls below marginal damage and adoption stops too early.

41.8 Sixth Step: What the Toy Model Keeps and Drops

When the regulator keeps its old policy, the instrument with the strongest adoption incentive is not the one with the lowest social cost: with convex damage, the tax leads to too much adoption and permits to too little, because the old tax is too strict and the old cap too lax once the new technology

is available. A regulator who anticipates adoption should rely on permits if it must commit before firms adopt, and can use either instrument if it adjusts its policy afterwards. With constant marginal damage, the tax needs no adjustment.

What the toy model keeps.

- The difference between industry-wide savings and each firm's own gain from adoption: the fall in the permit price benefits all firms, adopters or not, so no firm adopts to obtain it.
- The ranking under a committed policy: taxes give stronger incentives than permits, auctioned and free permits give the same incentives, and standards may give stronger incentives than permits once enough firms have adopted.
- The role of anticipation: a regulator who knows the new technology can induce the first-best share of adopters with permits, while a tax set before adoption leaves identical firms indifferent, so inefficient shares of adopters are also equilibria.

What it drops.

- General abatement cost and damage functions. The toy model fixes quadratic costs, a new technology that halves marginal abatement cost, and one value of the adoption cost.
- Asymmetric firms. For this case, Requate and Unold (2001) show that, with an aggregate emission level fixed before the new technology appears, taxes lead to overinvestment and permits to underinvestment, and that a regulator who knows the new technology can induce the optimal adoption whether it moves before or after adoption.
- Subsidies on abatement, which the paper shows give the same incentives as taxes, and old policies other than the one that is optimal with the old technology, for which the paper also compares the instruments (Lemmas 1 to 3).
- Standards when the regulator anticipates the new technology, which cannot reach the first best when partial adoption is optimal (Proposition 7 of the paper), a finite number of firms, and the coordination problem of which identical firms adopt (footnote 3 of the paper).
- Market power in the output or permit market, and the development of the new technology by an innovating sector. With such imperfections, the ranking of instruments can become ambiguous (Requate, 2005).

41.9 Connections to Other Chapters

- Chapter 45 studies firms that invest in abatement and in cost-reducing R&D, and a regulator that sets the emission fee after observing their investments.
- Chapter 44 studies abatement R&D when each government sets its emission tax to favor its own exporting firm.
- Chapter 42 studies abatement R&D by a Cournot duopoly when the regulator cannot commit to its emission tax.

References

Jung, C., Krutilla, K., and Boyd, R. (1996). Incentives for advanced pollution abatement technology at the industry level: An evaluation of policy alternatives. *Journal of Environmental Economics and Management*, 30(1), 95–111.

Milliman, S. R., and Prince, R. (1989). Firm incentives to promote technological change in pollution control. *Journal of Environmental Economics and Management*, 17(3), 247–265.

Requate, T. (2005). Dynamic incentives by environmental policy instruments: A survey. *Ecological Economics*, 54(2–3), 175–195.

Requate, T., and Unold, W. (2001). On the incentives created by policy instruments to adopt advanced abatement technology if firms are asymmetric. *Journal of Institutional and Theoretical Economics*, 157(4), 536–554.

Requate, T., and Unold, W. (2003). Environmental policy incentives to adopt advanced abatement technology: Will the true ranking please stand up? *European Economic Review*, 47(1), 125–146.

Chapter 42

Environmental R&D Cartels and Time-Consistent Emission Taxes

Companion note to Poyago-Theotoky, J. A. (2007), "The Organization of R&D and Environmental Policy," Journal of Economic Behavior & Organization 62(1): 63–75.

42.1 Introduction

Many polluting firms do research on clean technologies together. Consider two chemical producers that run a joint research program to cut their emissions while still competing for customers. Antitrust rules in the United States and the European Union treat such research agreements more leniently than other agreements among competitors. Is cooperation in environmental research good for the environment and for society?

The answer depends on how emission taxes are set. When the regulator cannot commit to a tax before firms invest, it sets the tax after observing their abatement, and more abatement leads to a lower tax. Firms anticipate this response, so research becomes a tool to influence policy. Firms that coordinate their research also care about how the tax affects each other's profit. A lower tax saves the rival taxes on its emissions, but it also leads every firm to produce more, which lowers the price the rival receives.

Poyago-Theotoky (2007) compares independent research with an environmental research cartel, in which firms choose their research to maximize joint profit but still compete in the product market. The cartel abates more and faces a lower tax when environmental damage is small, while independent research abates more when damage is large and research is efficient. Firms always prefer the cartel, and, unless research is very costly, social welfare ranks the two organizations as abatement does. A corrigendum to the paper adds that the tax can be negative, that is, an emission subsidy.

The toy model uses linear demand and quadratic costs. The first step solves the output stage and the regulator's tax, the second step solves for independent research, and the third step solves for the

research cartel and compares abatement and taxes. The fourth step compares profits and welfare, the fifth step adds research spillovers and a research joint venture, and the sixth step lists what the toy model keeps and drops.

42.2 The Toy Model

Firms and demand. Two firms produce a homogeneous good with inverse demand $p = a - Q$, where $Q = q_1 + q_2$, and constant marginal cost $c < a$. We normalize $A \equiv a - c = 1$: outputs, abatement levels and taxes below are multiplied by A in the general case, and profits and welfare by A^2, so no comparison depends on this normalization.

Emissions and abatement. Each unit of output generates one unit of emissions. By investing in environmental R&D, firm i abates $z_i \geq 0$ units, so its emissions are

$$e_i = q_i - z_i, \tag{42.1}$$

and the R&D cost is $\frac{\gamma}{2} z_i^2$, where $\gamma > 0$ measures how costly abatement is; a small γ means efficient R&D. Until the fifth step, a firm's R&D reduces only its own emissions. Total emissions $E = e_1 + e_2$ cause damage $\frac{d}{2} E^2$, where $d > 0$ is the damage parameter.

Policy and welfare. A regulator levies a tax t on each unit of emissions to maximize social welfare, the sum of consumer surplus, profits and tax revenue minus environmental damage. Since tax revenue is a transfer,

$$W = Q - \frac{Q^2}{2} - \frac{\gamma}{2}\left(z_1^2 + z_2^2\right) - \frac{d}{2} E^2. \tag{42.2}$$

Timing. In the first stage, firms choose their abatement z_1 and z_2. In the second stage, the regulator observes z_1 and z_2 and sets the tax t. In the third stage, firms compete in quantities. The regulator cannot commit to a tax before firms invest, so the tax is time consistent: it is optimal given the abatement already in place.

Two organizations of R&D. Under independent R&D (subscript I), each firm chooses z_i to maximize its own profit. In an environmental R&D cartel (subscript C), firms choose z_1 and z_2 to maximize their joint profit, but they still compete in the third stage. We consider damage levels

$$d > \underline{d} \equiv \frac{\sqrt{3} - 1}{2} \approx 0.366,$$

which guarantee positive abatement under both organizations (Results 2 and 3).

42.3 First Step: Output and the Time-Consistent Tax

Third stage. Given the tax t, firm i maximizes

$$(1 - q_i - q_j)\, q_i - t\,(q_i - z_i) - \frac{\gamma}{2} z_i^2.$$

The first-order condition,

$$1 - 2q_i - q_j - t = 0,$$

gives the symmetric output $q = \frac{1-t}{3}$. Since the price-cost margin net of the tax is then $1-2q-t = q$, firm i's profit becomes

$$\pi_i = q^2 + t\, z_i - \frac{\gamma}{2} z_i^2. \tag{42.3}$$

The term q^2 is the margin net of the tax times output, and the term $t\, z_i$ is the tax the firm saves on the units it abates.

Second stage. Let $Z = z_1 + z_2$ denote total abatement. Choosing t is equivalent to choosing total output $Q = \frac{2(1-t)}{3}$, and total emissions are $E = Q - Z$. The first-order condition of (42.2) with respect to Q is

$$1 - Q = d\,(Q - Z), \tag{42.4}$$

so the regulator equates the price-cost margin, $1-Q$, with the marginal damage, dE. Solving (42.4) and using $t = 1 - 3q$,

$$q(Z) = \frac{1 + dZ}{2(1+d)}, \qquad t(Z) = \frac{2d - 1 - 3dZ}{2(1+d)}. \tag{42.5}$$

Result 1 (time-consistent tax). Given total abatement Z, the regulator sets the tax $t(Z)$ in (42.5), and each firm produces $q(Z)$. Then:

- the tax decreases in Z, at rate $\frac{3d}{2(1+d)}$;
- the tax equals marginal damage minus output per firm, $t = dE - q$;
- the tax is negative, an emission subsidy, if and only if $Z > \frac{2d-1}{3d}$, which holds for every $Z \geq 0$ when $d < \frac{1}{2}$.

Intuitively, the regulator has one instrument for two problems. Pollution calls for a tax equal to marginal damage, but Cournot firms already produce too little, and a tax reduces output further. The time-consistent tax balances the two, so it falls short of marginal damage by q, and it becomes a subsidy when damage is small. More abatement lowers emissions and marginal damage, so the regulator responds with a lower tax. Firms anticipate this response when they invest in the first stage.

42.4 Second Step: Independent R&D

Substituting (42.5) into (42.3), firm i's first-stage profit is

$$\pi_i(z_i, z_j) = q(Z)^2 + t(Z)\, z_i - \frac{\gamma}{2} z_i^2.$$

Under independent R&D, firm i chooses z_i taking z_j as given. The first-order condition is

$$\underbrace{\frac{d}{1+d}\, q}_{\text{higher output}} + \underbrace{t}_{\text{tax saved}} - \underbrace{\frac{3d}{2(1+d)}\, z_i}_{\text{lower saving on abated units}} = \gamma z_i. \tag{42.6}$$

A unit of abatement saves the tax t, and through the tax rule it raises output and lowers the tax, which reduces the saving on the units the firm already abates. The second-order condition holds for all parameter values. Imposing symmetry, $z_i = z_j = z_I$, and defining

$$D_I \equiv 2\gamma(1+d)^2 + d(9+7d),$$

we obtain

$$z_I = \frac{2d^2+2d-1}{D_I}, \qquad t_I = \frac{d(2d-3)+2\gamma(2d-1)(1+d)}{2D_I}, \qquad q_I = \frac{2\gamma(1+d)+d(7+4d)}{2D_I}. \tag{42.7}$$

For $\frac{1}{2} < d < \frac{3}{2}$, the tax t_I is zero at the R&D cost

$$\gamma_I(d) \equiv \frac{d(3-2d)}{2(2d-1)(1+d)}.$$

Result 2 (independent R&D). Under independent R&D, abatement, the tax and output are given by (42.7). Moreover:

- abatement is positive if and only if $d > \underline{d} = \frac{\sqrt{3}-1}{2}$;
- the tax is negative if and only if $d(3-2d) > 2\gamma(2d-1)(1+d)$: for every γ when $d \le \frac{1}{2}$, for $\gamma < \gamma_I(d)$ when $\frac{1}{2} < d < \frac{3}{2}$, and never when $d \ge \frac{3}{2}$.[1]

Intuitively, when damage is small, the regulator's main concern is the output restriction of Cournot firms, so it subsidizes emissions, which works as an output subsidy. When R&D is efficient (small γ), firms abate a lot, which also pushes the tax down, by Result 1, so a subsidy arises for a wider range of damage levels. For $d \ge \frac{3}{2}$, damage is large enough for the tax to be positive for every R&D cost.

42.5 Third Step: The Environmental R&D Cartel

In the cartel, firms choose z_1 and z_2 to maximize

$$\pi_1 + \pi_2 = 2q(Z)^2 + t(Z)\,Z - \frac{\gamma}{2}\left(z_1^2 + z_2^2\right).$$

The first-order condition for z_i is the one in (42.6) plus the effect of firm i's abatement on its rival's profit,

$$\frac{d}{1+d}\,q + t - \frac{3d}{2(1+d)}\,z_i \;+\; \underbrace{\frac{d}{1+d}\left(q - \frac{3}{2}z_j\right)}_{\text{effect on the rival}} \;=\; \gamma z_i. \tag{42.8}$$

[1] The paper assumes $d > \frac{1}{2}$ to ensure interior R&D and a positive tax. Its corrigendum (Poyago-Theotoky, 2010) keeps $d > \frac{1}{2}$ for interior R&D only, since for $\frac{1}{2} < d < \frac{3}{2}$ the tax can be negative, an emission subsidy, when R&D is efficient (small γ); Proposition 1(i) of the paper remains valid with the negative tax read as a subsidy. Footnotes 10 and 13 of the paper show that $d > \frac{\sqrt{3}-1}{2}$ is enough for positive abatement under both organizations, and the toy model uses this weaker restriction.

The joint profit is concave in (z_1, z_2). Imposing symmetry, $z_i = z_j = z_C$, and defining

$$D_C \equiv 2\gamma(1+d)^2 + 4d(3+2d),$$

we obtain

$$z_C = \frac{2d^2+3d-1}{D_C}, \qquad t_C = \frac{d(2d-3)+\gamma(2d-1)(1+d)}{D_C}, \qquad q_C = \frac{\gamma(1+d)+d(5+2d)}{D_C}. \tag{42.9}$$

For $\frac{1}{2} < d < \frac{3}{2}$, the tax t_C is zero at the R&D cost

$$\gamma_C(d) \equiv \frac{d(3-2d)}{(2d-1)(1+d)} = 2\gamma_I(d).$$

Result 3 (R&D cartel). In the R&D cartel, abatement, the tax and output are given by (42.9). Moreover:

- abatement is positive for every $d > \underline{d}$, since $2d^2+3d-1 > 0$ for $d > \frac{\sqrt{17}-3}{4} \approx 0.281$;
- the tax is negative if and only if $d(3-2d) > \gamma(2d-1)(1+d)$: for every γ when $d \le \frac{1}{2}$, for $\gamma < \gamma_C(d)$ when $\frac{1}{2} < d < \frac{3}{2}$, and never when $d \ge \frac{3}{2}$.

Figure 42.1 depicts the taxes t_I and t_C as functions of the damage parameter, for $\gamma = \frac{1}{2}$. Under independent R&D the tax is a subsidy for $d < \frac{1+\sqrt{5}}{4} \approx 0.809$, and in the cartel for $d < 1$. The two taxes coincide at $d = 2$, where $t_I = t_C = \frac{1}{10}$; the cartel faces the lower tax for $d < 2$ and the higher tax for $d > 2$.

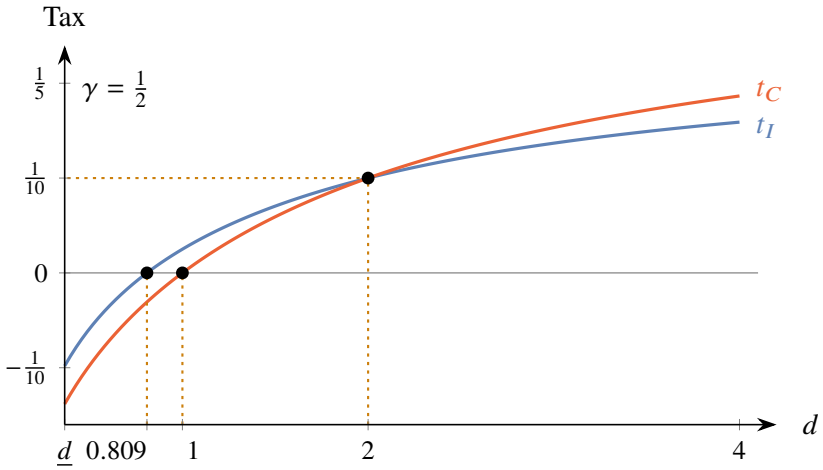


Figure 42.1: Time-consistent tax under independent R&D (t_I) and in the R&D cartel (t_C) for $\gamma = \frac{1}{2}$. Negative values are emission subsidies.

Intuitively, by (42.8), the cartel takes into account that a lower tax changes the rival's profit. By (42.3), the rival's profit changes with the tax at rate

$$\frac{\partial \pi_j}{\partial t} = -\left(\frac{2q}{3} - z_j\right).$$

A lower tax saves the rival the tax on its $q - z_j$ units of emissions, but it also leads firm i to expand its output, which lowers the price and costs the rival $\frac{q}{3}$. The rival gains from a lower tax if and only if its abatement is less than two thirds of its output. When damage is small, abatement is small relative to output, so the cartel invests more than independent firms to lower the tax, and, by Result 1, its tax is lower and becomes a subsidy for a wider range of R&D costs.

Result 4 (abatement and taxes compared). The difference in abatement is

$$z_C - z_I = \frac{(1+d)^2\, d\,(3 + 2\gamma - 2d)}{D_I\, D_C}. \tag{42.10}$$

Therefore:

- $z_C > z_I$ if and only if $d < \gamma + \frac{3}{2}$, and $z_C = z_I$ at $d = \gamma + \frac{3}{2}$;
- the taxes satisfy
$$t_C - t_I = -\frac{3d}{1+d}\,(z_C - z_I),$$
so the cartel faces a lower tax if and only if it abates more;
- total emissions are $E = \frac{1-2z}{1+d}$ when each firm abates z, so the cartel emits less if and only if it abates more.

Figure 42.2 depicts z_I and z_C as functions of d for $\gamma = \frac{1}{2}$, together with the abatement z^* that maximizes welfare given the tax rule, which we derive in the fourth step. The two curves cross at $d = \gamma + \frac{3}{2} = 2$, where $z_I = z_C = \frac{1}{5}$. As d grows, z_I tends to $\frac{1}{4}$ and z_C to $\frac{2}{9}$.

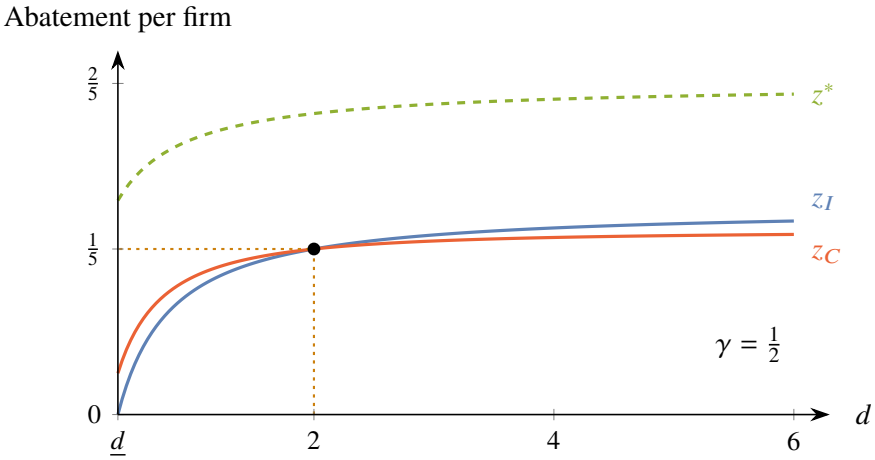


Figure 42.2: Abatement per firm under independent R&D (z_I), in the R&D cartel (z_C), and the welfare-maximizing abatement given the tax rule (z^*), for $\gamma = \frac{1}{2}$.

Figure 42.3 depicts the comparison of Result 4 in the (d, γ) plane. Above the line $\gamma = d - \frac{3}{2}$, the cartel abates more and faces a lower tax; below it, independent R&D abates more. To the left of the curves $\gamma_I(d)$ and $\gamma_C(d)$, the tax under independent R&D and in the cartel, respectively, is a subsidy. The horizontal guide at $\gamma = \frac{1}{2}$ is the slice of Figures 42.1 and 42.2.

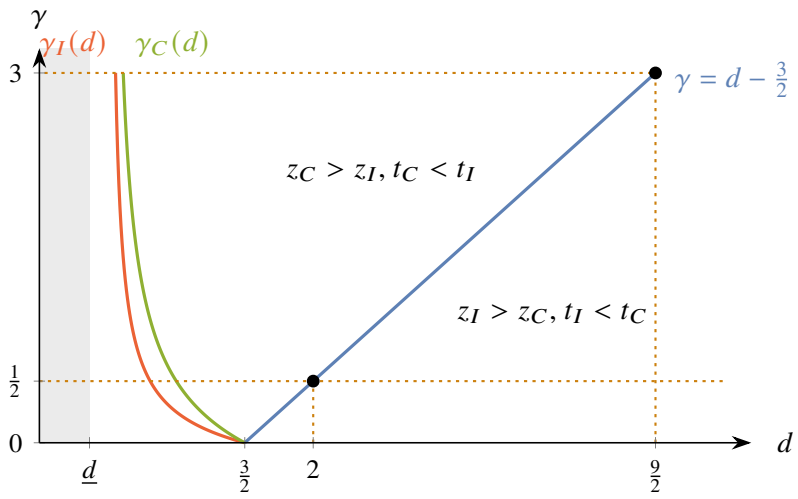


Figure 42.3: Organization of R&D in the (d, γ) plane. The cartel abates more above the line $\gamma = d - \frac{3}{2}$. The tax is a subsidy below $\gamma_I(d)$ under independent R&D and below $\gamma_C(d)$ in the cartel. The shaded area, $d \leq \underline{d}$, is excluded.

Intuitively, the effect on the rival in (42.8) has the sign of $2q - 3z$, and at the independent-R&D equilibrium

$$2q_I - 3z_I = \frac{(1+d)(3+2\gamma-2d)}{D_I}.$$

The cutoff $d = \gamma + \frac{3}{2}$ is therefore the damage level at which each firm abates two thirds of its output. A larger damage leads to more abatement relative to output, and a lower R&D cost does the same. When damage is small, the rival gains from a lower tax, and the cartel internalizes this gain by investing more. When damage is large, each firm emits little, a higher tax mainly restrains the rival's output, and the cartel invests less than independent firms to keep the tax high. Independent firms ignore this effect and invest more. Result 4 is the case without spillovers of Proposition 1 of Poyago-Theotoky (2007), whose critical damage level is then $\gamma + \frac{3}{2}$.

42.6 Fourth Step: Profits and Welfare

Profits. The cartel can always choose the abatement of independent R&D, and the second and third stages do not depend on how R&D was organized, so joint profit cannot be lower in the cartel.

Result 5 (profits). Profit per firm is higher in the cartel,

$$\pi_C - \pi_I = \frac{(1+d)^2\, d^2\, (3+2\gamma-2d)^2}{4\, D_I^2\, D_C} \;\geq\; 0, \tag{42.11}$$

with equality only at $d = \gamma + \frac{3}{2}$, where both organizations choose the same abatement.

Intuitively, firms gain from coordinating R&D in both regions of Figure 42.3: when damage is small, they coordinate to lower the tax, and when damage is large, they coordinate to keep it high.

Welfare. Welfare depends on how R&D is organized only through the abatement level, since the second and third stages are the same. When each firm abates z, output per firm is $q = \frac{1+2dz}{2(1+d)}$, by (42.5), and welfare (42.2) is a concave function of z. Its maximum, the abatement that a planner would choose anticipating the tax rule, is

$$z^* = \frac{d}{2d + \gamma(1+d)}. \tag{42.12}$$

Since welfare is quadratic in z,

$$W_C - W_I = \left(\gamma + \frac{2d}{1+d}\right)(z_C - z_I)(2z^* - z_C - z_I).$$

Define

$$\Omega \equiv 2(2-d)(1+d)^2\gamma^2 + d\gamma\left(14d^3 + 29d^2 + 36d + 29\right) + 2d^2\left(26d^2 + 51d + 21\right). \tag{42.13}$$

Result 6 (welfare). The difference in welfare is

$$W_C - W_I = \frac{d\,(1+d)^2\,(3 + 2\gamma - 2d)\,\Omega}{D_C^2\,D_I^2}. \tag{42.14}$$

Therefore:

- independent R&D abates less than z^* for all parameter values;
- if $d \le 2$, then $\Omega > 0$, and $W_C > W_I$ if and only if $z_C > z_I$;
- if $d > 2$, then $\Omega < 0$ if and only if $\gamma > \gamma_W(d)$, the positive root of $\Omega = 0$; since $\gamma_W(d) > d - \frac{3}{2}$, in this region $z_C > z_I$ but $W_C < W_I$. The cutoff $\gamma_W(d)$ is at least 64.2, its minimum, reached at $d \approx 4.29$.

Table 42.1 reports the equilibrium of both organizations for $\gamma = \frac{1}{2}$ at three damage levels. At $d = \frac{3}{4}$, both taxes are subsidies, and the cartel abates more, emits less and yields higher welfare; at $d = 2$, the two organizations coincide; at $d = 4$, independent R&D abates more, emits less and yields higher welfare, while profits are higher in the cartel at both $d = \frac{3}{4}$ and $d = 4$.

Intuitively, firms abate to save taxes and to influence the tax, not to reduce damage, and independent firms abate less than z^*. When the cartel abates more than independent firms, it moves abatement toward z^* and raises welfare; when it abates less, welfare falls. The exception in the third bullet requires very costly R&D: the cartel then abates more than z^*, which happens if and only if $d > 1$ and $\gamma > \frac{4d^2+2d}{d-1}$, a cutoff of at least $10 + 4\sqrt{6} \approx 19.8$, and it overshoots z^* by more than independent firms fall short of it. For the R&D costs of Figure 42.3, welfare ranks the two organizations as abatement does, as in Proposition 2 of the paper.

42.7 Fifth Step: Spillovers and Research Joint Ventures

Spillovers. Consider now that firm i also benefits from its rival's R&D, so its emissions are $e_i = q_i - z_i - \beta z_j$, with spillover $\beta \in [0, 1]$. With $\beta = 0$ we return to (42.1). The cartel uses

Table 42.1: Equilibrium under independent R&D (I) and in the R&D cartel (C), for $\gamma = \frac{1}{2}$ and $A = 1$.

	$d = \frac{3}{4}$		$d = 2$		$d = 4$	
Variable	I	C	I	C	I	C
Abatement per firm, z	0.1182	0.1434	0.2000	0.2000	0.2254	0.2139
Emission tax, t	−0.0091	−0.0415	0.1000	0.1000	0.1590	0.1866
Output per firm, q	0.3364	0.3472	0.3000	0.3000	0.2803	0.2711
Total emissions, E	0.4364	0.4075	0.2000	0.2000	0.1098	0.1144
Profit per firm, π	0.1086	0.1094	0.1000	0.1000	0.1017	0.1020
Welfare, W	0.3681	0.3807	0.3600	0.3600	0.3540	0.3462

the same spillover as independent firms; a research joint venture in which firms share all their knowledge corresponds to $\beta = 1$, following the terminology of Kamien et al. (1992). Repeating the first step, the tax rule becomes

$$t(Z) = \frac{2d - 1 - 3d(1+\beta)Z}{2(1+d)}, \tag{42.15}$$

and repeating the second and third steps,

$$z_I = \frac{(1+d)(2d-1) + d(1+\beta)}{D_I}, \qquad z_C = \frac{(1+\beta)\left[(1+d)(2d-1) + 2d\right]}{D_C}, \tag{42.16}$$

where

$$\begin{aligned} D_I &\equiv 2\gamma(1+d)^2 + d(1+\beta)\left[9 + 3\beta + d(7+\beta)\right], \\ D_C &\equiv 2\gamma(1+d)^2 + 4d(3+2d)(1+\beta)^2, \end{aligned}$$

which reduce to the expressions of the second and third steps when $\beta = 0$. The difference in abatement is

$$z_C - z_I = \frac{(1+d)^2\varphi}{D_I D_C},$$

where

$$\varphi \equiv d(3-2d)(1+\beta)^2(1-\beta) + 2\gamma\left[d + \beta\left(2d^2 + 2d - 1\right)\right]. \tag{42.17}$$

For $d > \frac{3}{2}$, φ is zero at the R&D cost

$$\hat{\gamma}(d,\beta) \equiv \frac{d(2d-3)(1+\beta)^2(1-\beta)}{2\left[d + \beta\left(2d^2 + 2d - 1\right)\right]}.$$

Result 7 (spillovers). With spillover $\beta \in [0,1]$:

- if $d \leq \frac{3}{2}$, then $z_C > z_I$;
- if $d > \frac{3}{2}$, then $z_C > z_I$ if and only if $\gamma > \hat{\gamma}(d,\beta)$, where $\hat{\gamma}$ increases in d and decreases in β;

- for $\beta > 0$, $\hat{\gamma}(d, \beta)$ is below $\frac{(1+\beta)^2(1-\beta)}{2\beta}$, its limit as d grows;
- with a research joint venture, $\beta = 1$, $\varphi = 2\gamma(2d^2 + 3d - 1) > 0$, so $z_C > z_I$ for every d;
- in all cases,

$$t_C - t_I = -\frac{3d(1+\beta)}{1+d}(z_C - z_I),$$

and profit per firm is higher in the cartel unless $\varphi = 0$.

Figure 42.4 depicts the cutoff $\hat{\gamma}(d, \beta)$ for $\beta = 0$, where it is the line $\gamma = d - \frac{3}{2}$ of Figure 42.3, and for $\beta = \frac{1}{5}$ and $\beta = \frac{1}{2}$, whose limits are $\frac{72}{25}$ and $\frac{9}{8}$. The cartel abates more above each curve. For $\beta = 1$, the curve disappears.

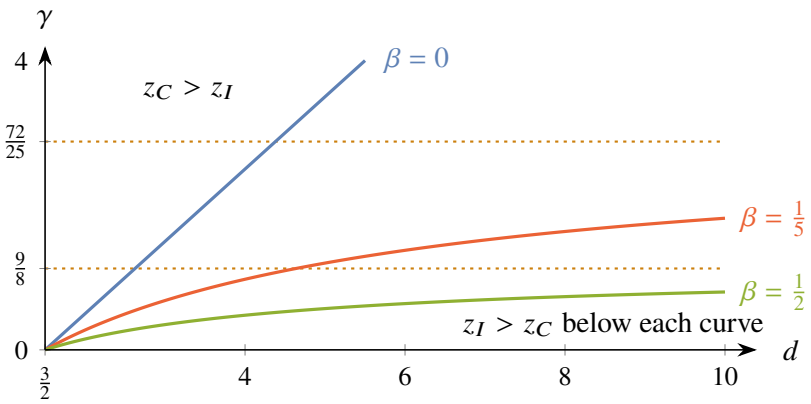


Figure 42.4: Cutoff R&D cost $\hat{\gamma}(d, \beta)$ above which the R&D cartel abates more, for $\beta = 0$, $\beta = \frac{1}{5}$ and $\beta = \frac{1}{2}$. Dotted lines are the limits as d grows.

Intuitively, spillovers add an appropriability problem: part of firm i's R&D lowers its rival's emissions, and an independent firm ignores this benefit, while the cartel takes it into account. A larger β therefore raises abatement in the cartel relative to independent R&D, and the cutoff $\hat{\gamma}$ falls. With full knowledge sharing, each firm's R&D lowers both firms' emissions equally, and the cartel abates more at every damage level. Result 7 reproduces Proposition 1 and Corollary 1 of the paper, whose critical R&D cost is the limit in the third bullet.

42.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- A regulator who cannot commit, so that firms invest in abatement to lower the future tax, and a time-consistent tax that falls short of marginal damage and can become a subsidy.
- The two effects of coordinating R&D: the cartel internalizes the gain of the rival from a lower tax when damage is small, and the gain from a higher tax when damage is large, which delivers the cutoff $d = \gamma + \frac{3}{2}$ without spillovers.

- Profits that always favor the cartel, a welfare ranking that follows the abatement ranking for every $d \leq 2$, and spillovers and research joint ventures that favor the cartel.

What it drops.

- More than two firms, cartels that include only some firms, and the firms' choice of whether to form a research joint venture, which Ouchida and Goto (2016) study in this model.
- A regulator who commits to the tax before firms invest, and other instruments, such as emission standards or tradable permits.
- Competing regulators in different countries, which Iida (2020) studies with international R&D cooperation and strategic emission taxes.
- Asymmetric firms, uncertainty about the outcome of R&D, and abatement technologies that change the production process rather than cleaning emissions at the end of it.

42.9 Connections to Other Chapters

- Chapter 37 studies the four ways of organizing R&D of this chapter without pollution, and shows that an RJV cartel yields the largest cost reduction and the highest total surplus.
- Chapter 43 keeps the emissions, spillovers, and R&D costs of this chapter, lets the number of firms vary, and has the regulator commit to the tax before firms invest; total R&D then has an inverted-U shape in the number of firms when spillovers are present.
- Chapter 41 compares emission taxes, permits, and standards by the incentives they give firms to adopt an advanced abatement technology, including when the regulator adjusts its policy after adoption.
- Chapter 44 extends the comparison between independent R&D and a research joint venture to two exporting countries whose governments set emission taxes strategically.
- Chapter 45 studies firms that invest both in abatement and in cost-reducing R&D under an emission fee that responds to their investments.
- Chapter 15 studies a merger between rival innovators, another setting where coordination changes firms' incentives to invest.

References

Iida, T. (2020). International R&D formations and strategic environmental policy. *Environment and Development Economics*, 25(5), 509–528.

Kamien, M. I., Muller, E., and Zang, I. (1992). Research joint ventures and R&D cartels. *American Economic Review*, 82(5), 1293–1306.

Ouchida, Y., and Goto, D. (2016). Environmental research joint ventures and time-consistent emission tax: Endogenous choice of R&D formation. *Economic Modelling*, 55, 179–188.

Poyago-Theotoky, J. A. (2007). The organization of R&D and environmental policy. *Journal of Economic Behavior & Organization*, 62(1), 63–75.

Poyago-Theotoky, J. A. (2010). Corrigendum to "The organization of R&D and environmental policy" [J. Econ. Behav. Org. 62 (2007) 63–75]. *Journal of Economic Behavior & Organization*, 76(2), 449.

Chapter 43

Competition and Green R&D under a Precommitted Emission Tax

Companion note to Lambertini, L., J. Poyago-Theotoky, and A. Tampieri (2017), "Cournot Competition and 'Green' Innovation: An Inverted-U Relationship," Energy Economics 68: 116–123.

43.1 Introduction

Does competition encourage innovation? Aghion et al. (2005) documented an inverted U: innovation first rises and then falls as competition intensifies. Consider now green research, such as the filters and scrubbers that polluting plants install to clean their emissions. When a regulator taxes emissions, each plant invests in filters to save on the tax. What happens to total green research as more plants enter the market?

Since firms invest only to save on the tax, the regulator's tax is the channel through which competition affects research. Two forces work in opposite directions. More firms produce more output, so there are more emissions to clean, and the regulator keeps the tax high enough for total research to grow. In contrast, when the research of one firm also cleans part of its rivals' emissions, a spillover, each unit of research cleans more emissions as the number of firms grows. The regulator then needs less research and lowers the tax.

Lambertini, Poyago-Theotoky and Tampieri (2017) study a regulator who commits to an emission tax before firms choose their research, after which firms compete in quantities. Without spillovers, total green research always increases in the number of firms. With spillovers, it first increases and then decreases, an inverted U. Social welfare, however, keeps increasing in the number of firms, so the market structure that maximizes green research is not the one that maximizes welfare.

The toy model uses linear demand and quadratic costs. The first step solves the output and research stages for a given tax, and the second step solves for the precommitted tax. The third step shows that total research increases with the number of firms without spillovers, the fourth step adds spillovers

and obtains the inverted U, the fifth step studies welfare, and the sixth step lists what the toy model keeps and drops.

43.2 The Toy Model

Firms and demand. There are $n \geq 1$ identical firms producing a homogeneous good with inverse demand $p = a - Q$, where $Q = q_1 + \cdots + q_n$, and constant marginal cost $c < a$. As in the companion chapter on Poyago-Theotoky (2007), we normalize $A \equiv a - c = 1$: outputs, R&D levels and taxes below are multiplied by A in the general case, and profits and welfare by A^2.

Emissions, R&D and spillovers. Each unit of output generates one unit of emissions. Firm i invests $z_i \geq 0$ in green R&D at cost $\frac{\gamma}{2} z_i^2$, where $\gamma > 0$ measures how costly R&D is. The R&D of each firm also cleans part of its rivals' emissions, so the emissions of firm i are

$$e_i = q_i - z_i - \beta \sum_{j \neq i} z_j, \tag{43.1}$$

where $\beta \in [0, 1]$ is the spillover rate. With $\beta = 0$ and $n = 2$, (43.1) is the emission function of the companion chapter. When every firm invests z, firm i's emissions are $q_i - Bz$, where

$$B \equiv 1 + \beta(n - 1)$$

is the number of units of emissions that one unit of R&D removes from the industry. Total emissions $E = e_1 + \cdots + e_n$ cause damage $\frac{d}{2} E^2$, where $d > 0$. The paper writes the damage as dE^2, so its damage parameter equals $\frac{d}{2}$ here.

Policy and welfare. A regulator levies a tax t on each unit of emissions to maximize social welfare, consumer surplus plus profits plus tax revenue minus damage. Since tax revenue is a transfer,

$$W = Q - \frac{Q^2}{2} - \frac{\gamma}{2} \sum_{i=1}^{n} z_i^2 - \frac{d}{2} E^2. \tag{43.2}$$

Timing. In the first stage, the regulator sets the tax t. In the second stage, firms observe t and choose their R&D. In the third stage, firms compete in quantities. Unlike in the companion chapter, the regulator commits to the tax before firms invest, so firms cannot invest to change the tax. We consider damage levels $d \geq 1$, which guarantee a positive tax for every $n \geq 1$ (Result 2). Figures and Table 43.1 use $\gamma = 4$ and $d = 10$; they treat n as continuous, and dots mark integer values of n.

43.3 First Step: Output and R&D for a Given Tax

Third stage. Given the tax t and the R&D levels, firm i maximizes

$$(1 - q_i - Q_{-i})\, q_i - t\Big(q_i - z_i - \beta \sum_{j \neq i} z_j\Big) - \frac{\gamma}{2} z_i^2,$$

where Q_{-i} is the output of its rivals. The first-order condition,

$$1 - 2q_i - Q_{-i} - t = 0,$$

gives the symmetric output $q = \frac{1-t}{n+1}$. The price-cost margin net of the tax is then $1 - Q - t = q$, so firm i's profit becomes

$$\pi_i = q^2 + t\Big(z_i + \beta \sum_{j \neq i} z_j\Big) - \frac{\gamma}{2} z_i^2. \tag{43.3}$$

Output does not depend on R&D, since R&D cleans emissions at the end of the production process.

Second stage. Firm i chooses z_i to maximize (43.3), taking its rivals' R&D as given. The first-order condition is $t = \gamma z_i$: one more unit of R&D saves the tax on one unit of emissions and costs γz_i.

Result 1 (output and R&D for a given tax). For a given tax t, each firm produces and invests

$$q(t) = \frac{1-t}{n+1}, \qquad z(t) = \frac{t}{\gamma}. \tag{43.4}$$

Then:

- R&D per firm does not depend on n or β, and it is zero if $t = 0$;
- total R&D, $Z(t) = \frac{nt}{\gamma}$, increases in n for every $t > 0$.

Intuitively, a firm invests in green R&D only to save on the tax, and the saving on each unit of abatement is t, regardless of the number of rivals. The R&D of rivals lowers firm i's tax bill in (43.3), but it does not change the return of firm i's own R&D, so spillovers do not affect investment for a given tax. If the tax were fixed, more firms would mean more total R&D, which the paper calls an Arrowian result. Any other relationship between competition and green R&D must therefore come from the response of the regulator's tax to the number of firms, which we study next.

43.4 Second Step: The Precommitted Tax

Anticipating (43.4), total output is $Q = \frac{n(1-t)}{n+1}$, total R&D is $Z = \frac{nt}{\gamma}$, and total emissions are

$$E = Q - BZ.$$

Raising the tax lowers each firm's output by $\frac{1}{n+1}$ and raises each firm's R&D by $\frac{1}{\gamma}$, which removes $\frac{B}{\gamma}$ units of each firm's emissions. The first-order condition of (43.2) with respect to t, divided by n, is

$$\underbrace{dE\Big(\frac{1}{n+1} + \frac{B}{\gamma}\Big)}_{\text{lower damage}} = \underbrace{\frac{1-Q}{n+1}}_{\text{lost margin on output}} + \underbrace{z}_{\text{R\&D cost}}, \tag{43.5}$$

and welfare is concave in t. Define

$$K \equiv \gamma + (n+1)B \qquad \text{and} \qquad \Delta \equiv dnK^2 + \gamma\left[\gamma n + (n+1)^2\right].$$

Result 2 (precommitted tax). The regulator sets the tax

$$t^* = \frac{\gamma\,(dnK - \gamma)}{\Delta}, \tag{43.6}$$

and in equilibrium each firm invests $z^* = \frac{t^*}{\gamma}$, while

$$q^* = \frac{dnBK + \gamma(\gamma + n + 1)}{\Delta}, \qquad E^* = \frac{\gamma n\,(\gamma + n + 1 + B)}{\Delta} > 0. \tag{43.7}$$

Moreover:

- since $dnK - \gamma = \gamma(dn - 1) + dn(n + 1)B$, the tax is positive for every $n \geq 1$ when $d \geq 1$;
- the tax increases in d;
- the tax satisfies

$$t^* = \frac{K}{\gamma + n + 1}\, dE^* - \frac{\gamma}{\gamma + n + 1}\, q^*, \tag{43.8}$$

  where the weight on marginal damage, $\frac{K}{\gamma+n+1} = 1 + \frac{\beta(n-1)(n+1)}{\gamma+n+1}$, equals one if $\beta = 0$ and exceeds one if $\beta > 0$ and $n \geq 2$.

Figure 43.1 depicts the tax t^* against the number of firms for $\beta = 0$ and $\beta = \frac{1}{2}$, with $\gamma = 4$ and $d = 10$. Under monopoly, spillovers play no role and both taxes equal $\frac{4}{7}$. For these parameter values the tax decreases in n, and it decreases faster with spillovers; Proposition 2 of the paper shows that, for other damage levels, the tax can first increase and then decrease in n.

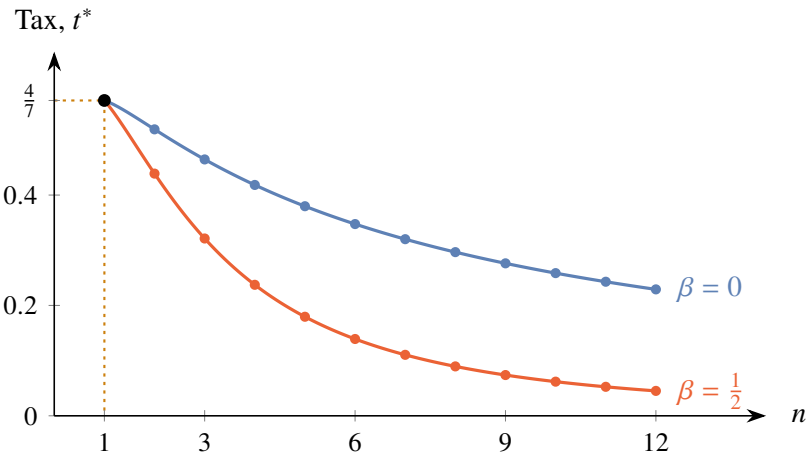


Figure 43.1: Precommitted emission tax against the number of firms, without spillovers ($\beta = 0$) and with spillovers ($\beta = \frac{1}{2}$), for $\gamma = 4$ and $d = 10$.

Intuitively, by (43.5), a higher tax lowers damage through two channels, lower output and more R&D, at the cost of a lower output in a market where Cournot firms already produce too little, and of a higher R&D cost. Since $1 - Q = q + t$ and $z = \frac{t}{\gamma}$, solving (43.5) for t gives (43.8). If R&D were already in place, the tax would equal $dE - q$, as in the time-consistent tax of the companion chapter: marginal damage minus a correction for the output restriction. A precommitted regulator

knows that the tax also induces R&D, so the correction for output in (43.8) is only the fraction $\frac{\gamma}{\gamma+n+1}$ of q. With spillovers, the weight on marginal damage exceeds one: one unit of R&D by firm i removes B units of emissions from the industry, but firm i invests only to save the tax on its own unit, so the regulator gives marginal damage a weight above one, which rewards the units that firm i cleans for its rivals.

43.5 Third Step: Competition and Green R&D without Spillovers

Total R&D is $Z^* = nz^* = \frac{nt^*}{\gamma}$. With $\beta = 0$, we have $B = 1$ and $K = \gamma + n + 1$, so

$$Z^*\big|_{\beta=0} = \frac{n\left[dn(\gamma+n+1) - \gamma\right]}{dn(\gamma+n+1)^2 + \gamma\left[\gamma n + (n+1)^2\right]}. \tag{43.9}$$

Result 3 (no spillovers). Consider that $\beta = 0$. Then:

- total R&D (43.9) increases in n for every $n \geq 1$, $\gamma > 0$ and $d > 0$, since the numerator of its derivative equals $\gamma^2(n^2 - 1)$ plus a polynomial in n, γ and d with positive coefficients;
- total R&D tends to 1 as $n \to \infty$;
- for $\gamma = 4$ and $d = 10$, R&D per firm, z^*, decreases in n for every $n \geq 1$.

Figure 43.2 depicts total R&D, Z^*, and R&D per firm, z^*, against the number of firms for $\beta = 0$, $\gamma = 4$ and $d = 10$. Under monopoly both equal $\frac{1}{7}$. As n grows, each firm invests less, but total R&D keeps increasing toward 1.

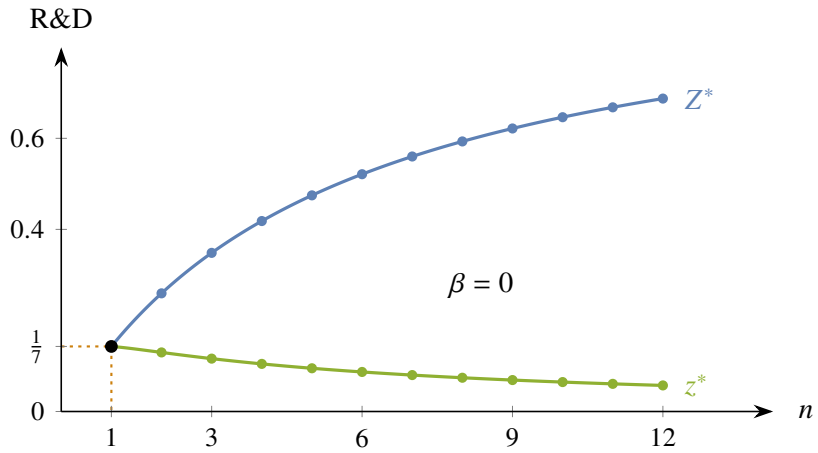


Figure 43.2: Total R&D (Z^*) and R&D per firm (z^*) against the number of firms without spillovers, for $\gamma = 4$ and $d = 10$.

Intuitively, without spillovers each unit of R&D cleans one unit of emissions, and total emissions are $E = Q - Z$. More firms produce more, so there are more emissions to clean, and the regulator

keeps the tax high enough for total R&D to follow total output. The tax falls with n, by Figure 43.1, so each firm invests less, but the tax falls slower than $\frac{1}{n}$, and total R&D increases. This is Lemma 1 of the paper, the case in which competition and green innovation move together.

43.6 Fourth Step: Spillovers and the Inverted U

A zero-emission benchmark. We first consider very large damage, $d \to \infty$. Then the regulator chooses the tax that eliminates emissions: by (43.7), $E^* \to 0$, and each firm's emissions, $q - Bz$, vanish. The limits of (43.6) and (43.7) are

$$t^\infty = \frac{\gamma}{K}, \qquad z^\infty = \frac{1}{K}, \qquad q^\infty = \frac{B}{K}, \qquad Z^\infty = \frac{n}{K} = \frac{n}{\gamma + (n+1)\,[1 + \beta(n-1)]}. \tag{43.10}$$

Result 4 (inverted U in the zero-emission benchmark). Total R&D Z^∞ in (43.10) satisfies:

- if $\beta = 0$, then $Z^\infty = \frac{n}{\gamma+n+1}$, which increases in n;
- if $\beta > 0$, then the derivative of Z^∞ has the sign of $\gamma + 1 - \beta - \beta n^2$, so Z^∞ increases for $n < \hat{n}$ and decreases for $n > \hat{n}$, where
$$\hat{n} = \sqrt{\frac{1 + \gamma - \beta}{\beta}}; \tag{43.11}$$
- $\hat{n} \geq 2$ if and only if $\beta \leq \frac{1+\gamma}{5}$;
- for every n, total R&D Z^* increases in d and tends to Z^∞ as $d \to \infty$.

For $\gamma = 4$, $Z^\infty = \frac{n}{n+5}$ if $\beta = 0$; $Z^\infty = \frac{2n}{n^2+2n+9}$ if $\beta = \frac{1}{2}$, with $\hat{n} = 3$ and $Z^\infty = \frac{1}{4}$ at the peak; and $Z^\infty = \frac{n}{n^2+n+4}$ if $\beta = 1$, with $\hat{n} = 2$ and $Z^\infty = \frac{1}{5}$ at the peak.

Figure 43.3 depicts Z^∞ against the number of firms for $\gamma = 4$ and the three spillover rates. All curves start at $\frac{1}{6}$ under monopoly. Without spillovers, total R&D increases toward 1; with spillovers, it peaks at $n = 3$ when $\beta = \frac{1}{2}$, and at $n = 2$ when $\beta = 1$, and then decreases toward zero.

Intuitively, when emissions are eliminated, each firm invests $z = \frac{q}{B}$, so total R&D is

$$Z^\infty = \frac{Q^\infty}{B}.$$

The numerator is total output, which increases in n because competition expands output; this is the force behind Result 3. The denominator is the number of units of emissions that one unit of R&D removes. Without spillovers it equals one, so total R&D follows total output. With spillovers, $B = 1 + \beta(n-1)$ grows with the number of firms, since each unit of R&D cleans the emissions of more rivals, while total output cannot exceed one. The regulator then needs less R&D to clean the industry and lowers the tax, and for $n > \hat{n}$ the second force dominates. The condition $\beta \leq \frac{1+\gamma}{5}$ is the sufficient condition of Proposition 6 of the paper for the peak to occur at $n \geq 2$.

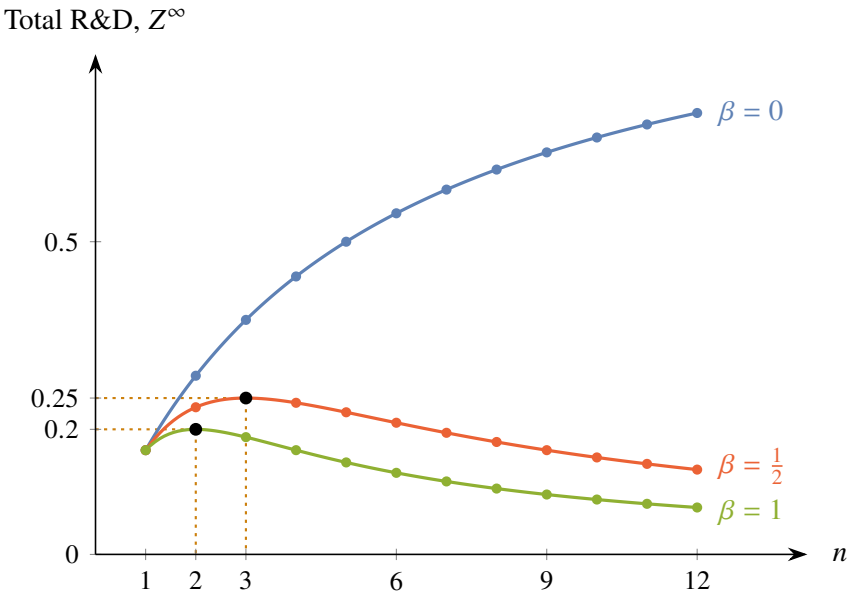


Figure 43.3: Total R&D in the zero-emission benchmark ($d \to \infty$) against the number of firms, for $\gamma = 4$ and $\beta \in \{0, \frac{1}{2}, 1\}$.

Finite damage. For a finite d, total R&D is $Z^* = \frac{nt^*}{\gamma}$, with t^* from (43.6). When $\beta > 0$, its numerator is a polynomial of degree four in n and its denominator one of degree five, so $Z^* \to 0$ as $n \to \infty$.

Result 5 (inverted U with finite damage). Consider that $\gamma = 4$ and $d = 10$. Then:

- if $\beta = 0$, total R&D increases in n;
- if $\beta = \frac{1}{2}$, total R&D increases for $n < \hat{n} \approx 3.25$ and decreases for $n > \hat{n}$; among integers it is highest at $n = 3$, where $Z^* = 0.2410$, against 0.2194 at $n = 2$ and 0.2374 at $n = 4$;
- for every $\beta > 0$, total R&D tends to zero as $n \to \infty$;
- for $\beta = 0$ and $\beta = \frac{1}{2}$, R&D per firm decreases in n.

Figure 43.4 depicts total R&D against the number of firms for $\beta = 0$ and $\beta = \frac{1}{2}$, with $\gamma = 4$ and $d = 10$. The curve without spillovers increases, while the curve with spillovers is an inverted U whose peak, at $\hat{n} \approx 3.25$, lies close to the peak $n = 3$ of the zero-emission benchmark.[1] Table 43.1 reports the tax, total R&D and welfare for $n = 1, \ldots, 8$.

Intuitively, the logic of the benchmark carries over. Since $E = Q - BZ$, total R&D is

$$Z^* = \frac{Q^* - E^*}{B}.$$

With finite damage, some emissions remain, but they are small: at $n = 3$ and $\beta = \frac{1}{2}$, emissions per firm are 0.0090, against 0.0177 without spillovers. Total output increases in n, while B grows with

[1] Proposition 6 of the paper also states that total R&D is concave in n. In the toy model total R&D is concave around its peak but convex for large n, as it falls toward zero: for $\gamma = 4$, $d = 10$ and $\beta = \frac{1}{2}$, it is convex for n above approximately 5.82, which is also the shape of Figure 3 of the paper.

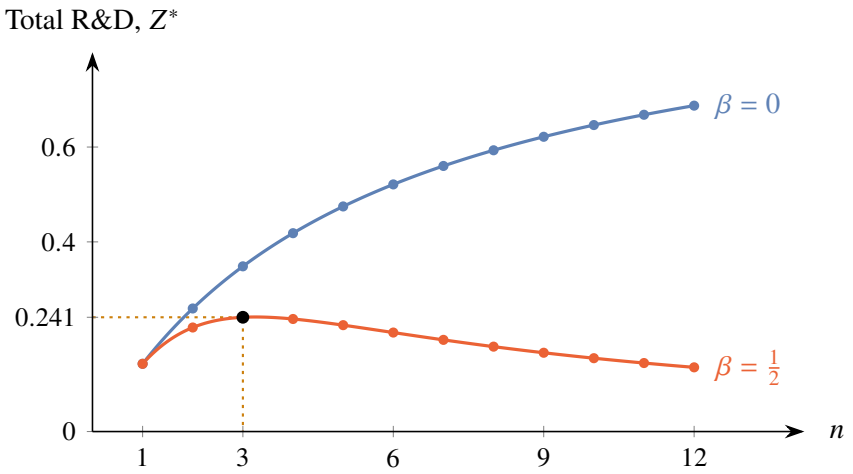


Figure 43.4: Total green R&D against the number of firms without spillovers ($\beta = 0$) and with spillovers ($\beta = \frac{1}{2}$), for $\gamma = 4$ and $d = 10$. With spillovers, total R&D is highest at $n = 3$ among integers.

Table 43.1: Precommitted tax t^*, total R&D Z^* and welfare W^*, for $\gamma = 4$, $d = 10$ and $A = 1$.

	$\beta = 0$			$\beta = \frac{1}{2}$		
Number of firms, n	t^*	Z^*	W^*	t^*	Z^*	W^*
1	0.5714	0.1429	0.1250	0.5714	0.1429	0.1250
2	0.5191	0.2595	0.1832	0.4389	0.2194	0.2459
3	0.4646	0.3484	0.2259	0.3213	0.2410	0.3371
4	0.4183	0.4183	0.2585	0.2374	0.2374	0.3944
5	0.3798	0.4747	0.2843	0.1795	0.2244	0.4293
6	0.3475	0.5212	0.3051	0.1392	0.2088	0.4509
7	0.3201	0.5601	0.3223	0.1105	0.1934	0.4646
8	0.2966	0.5932	0.3367	0.0895	0.1791	0.4737

n when $\beta > 0$, and the second force dominates for $n > \hat{n}$. The precommitted tax is the channel: by Result 1, a fixed tax would give total R&D $\frac{nt}{\gamma}$, which increases in n for every β. The inverted U arises because, with spillovers, the regulator lowers the tax faster than $\frac{1}{n}$ once n exceeds $\hat{n}$, as Figure 43.1 and Table 43.1 show. Result 5 reproduces Proposition 6 and Figure 3 of the paper, and, together with Result 3, its conclusion that spillovers drive the inverted U.

43.7 Fifth Step: Welfare and Market Structure

Substituting (43.4) and (43.6) into (43.2), equilibrium welfare is

$$W^* = \frac{n\left\{dn\left[\gamma(2B-1)+(n+2)B^2\right]+\gamma(\gamma+n+2)\right\}}{2\Delta}, \tag{43.12}$$

which is Eq. (9) of the paper with its damage parameter replaced by $\frac{d}{2}$.

Result 6 (welfare). Consider that $\gamma = 4$ and $d = 10$. For $\beta = 0$ and for $\beta = \frac{1}{2}$:

- welfare (43.12) increases in n for every $n > 0$;
- welfare tends to $\frac{1}{2}$ as $n \to \infty$, the total surplus of a competitive market without pollution.

Table 43.1 reports welfare for $n = 1, \ldots, 8$. With $\beta = \frac{1}{2}$, the market structure that maximizes green R&D, $n = 3$, yields welfare 0.3371, below the limit $\frac{1}{2}$ reached as the number of firms grows. Welfare is also higher with spillovers than without them for every $n \geq 2$ in the table.

Intuitively, more firms expand output toward the competitive level, which raises consumer surplus, and the regulator keeps emissions small through the tax. Green R&D is a cost that society bears only to reduce emissions. With spillovers, a larger industry cleans its emissions with less R&D per unit of output, so the decline in total R&D after $n = 3$ saves R&D costs rather than raising damage. The market structure that maximizes green innovation is therefore not the one that maximizes welfare, as the paper concludes at the end of its Section 3.

43.8 Sixth Step: What the Toy Model Keeps and Drops

What the toy model keeps.

- An n-firm Cournot industry with end-of-pipe green R&D, R&D spillovers, and a regulator who commits to the emission tax before firms invest, so that R&D responds to competition only through the tax.
- The closed forms of the paper for the tax, output, R&D, emissions and welfare, a positive relationship between competition and total R&D without spillovers (Lemma 1), and an inverted U with spillovers (Proposition 6), whose peak occurs at $n \geq 2$ when $\beta \leq \frac{1+\gamma}{5}$ in the zero-emission benchmark.

- Welfare that increases in the number of firms, so that the market structure maximizing green R&D differs from the one maximizing welfare.

What it drops.

- The full comparative statics of the tax with respect to n, γ and β (Propositions 2 to 4 of the paper) and of emissions per firm (Proposition 5), and the simulations of Figure 4 of the paper.
- A regulator who sets the tax after firms invest, as in the companion chapter on Poyago-Theotoky (2007), and R&D cooperation among firms.
- Other measures of competition, such as price competition, product differentiation or concentration indices, which the paper mentions as extensions, and free entry, which would make the number of firms an integer determined in equilibrium.
- Asymmetric firms, uncertainty about the outcome of R&D, and abatement that changes the production process rather than cleaning emissions at the end of it.

43.9 Connections to Other Chapters

- Chapter 42 uses the same emissions, spillovers, and R&D costs, but the regulator sets the tax after firms invest, so firms use R&D to lower the tax. In this chapter the regulator commits to the tax first.
- In Chapter 44, governments also commit to their emission taxes before firms invest, but each government sets its tax to favor its own exporting firm.
- Chapter 45 studies firms that invest both in abatement and in cost-reducing R&D under an emission fee that responds to their investments.
- Chapter 15 studies another link between market structure and innovation: a merger between rival innovators, after which the merging firms innovate less.

References

Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2005). Competition and innovation: An inverted-U relationship. *Quarterly Journal of Economics*, 120(2), 701–728.

Lambertini, L., Poyago-Theotoky, J., and Tampieri, A. (2017). Cournot competition and "green" innovation: An inverted-U relationship. *Energy Economics*, 68, 116–123.

Poyago-Theotoky, J. A. (2007). The organization of R&D and environmental policy. *Journal of Economic Behavior & Organization*, 62(1), 63–75.

Chapter 44

International R&D Cooperation and Strategic Emission Taxes

Companion note to Iida, T. (2020), "International R&D Formations and Strategic Environmental Policy," Environment and Development Economics 25(5): 509–528.

44.1 Introduction

Consider two steel makers in two different countries that sell all their output in a third country. Each pollutes at home, and each can invest in technology that cleans its emissions. Governments often encourage firms to cooperate in this research through a *research joint venture* (RJV), in which the firms share their results. Each government, however, also sets its emission tax to favor its own exporter. Iida (2020) asks how research cooperation and these strategic taxes interact.

When the two firms coordinate their R&D in an RJV cartel, they invest more, because each firm now counts the tax savings that its abatement gives the rival. Governments respond by changing their taxes. When R&D is inexpensive, the cartel invests heavily, and each government cuts its tax to restrain an investment that partly serves the rival. When R&D is costly, each government instead raises its tax to induce more abatement by the rival, whose research results also reduce domestic emissions.

These tax responses can turn cooperation against the countries that encourage it. When R&D is very inexpensive, the low taxes expand output in the third country, which lowers the price and each exporter's revenue, so each exporting country is worse off under the RJV cartel. When R&D is very costly, the cartel can lower the combined welfare of the three countries. Both findings contrast with the standard result that an RJV cartel maximizes welfare (Kamien et al. 1992). Allowing such cooperation is therefore a decision for all governments, including the importing one.

The toy model uses the smallest environmental damage the paper allows, so that every equilibrium object has a closed form. The first to third steps solve the output, R&D, and tax stages and compare R&D across the ways firms can organize research. The fourth step compares national welfare, the

fifth step adds the importing country, the sixth step lets environmental damage be larger, and the seventh step lists what the toy model keeps and drops.

44.2 The Toy Model

Market. Firm 1 is located in country 1 and firm 2 in country 2. Both produce at zero cost and sell all their output in a third country, where inverse demand is $p = 1 - q_1 - q_2$, and compete à la Cournot.

Pollution and abatement. Each unit of output generates one unit of emissions in the producing country. Firm i invests x_i in end-of-pipe abatement at cost $\frac{\eta}{2}x_i^2$, where $\eta > 0$ measures how costly R&D is, and the rival's investment spills over with rate $\beta \in [0, 1]$, so emissions in country i are

$$e_i = q_i - x_i - \beta x_j. \tag{44.1}$$

Environmental damage in country i is $\frac{1}{4}e_i^2$, so marginal damage is $\frac{e_i}{2}$; this is the smallest damage the paper allows.

Policy and welfare. Country i taxes the emissions of its firm at rate t_i. Firm i earns $\pi_i = pq_i - t_i e_i - \frac{\eta}{2}x_i^2$. Since consumers live in the third country, welfare in country i is profit plus tax revenue minus damage,

$$SW_i = pq_i - \frac{\eta}{2}x_i^2 - \frac{1}{4}e_i^2. \tag{44.2}$$

R&D formations. As in Kamien et al. (1992), we compare R&D competition (each firm maximizes its own profit) with an R&D cartel (firms maximize joint profit), for a given β, and the research joint venture (RJV) versions of both, where firms fully share results, $\beta = 1$. We call these four formations N, C, NJ, and CJ.

Timing. (1) Both governments simultaneously set t_1 and t_2; (2) firms choose R&D, competitively or jointly; (3) firms choose outputs. Firms never coordinate outputs, since antitrust law forbids it.

44.3 First Step: Output and R&D

Output. Abatement does not affect the marginal cost of output, which equals the tax. Cournot competition with costs t_1 and t_2 gives

$$q_i = \frac{1 - 2t_i + t_j}{3}. \tag{44.3}$$

R&D. One unit of abatement saves firm i the tax t_i, so under R&D competition firm i chooses x_i so that $\eta x_i = t_i$. In a cartel, firm i's abatement also lowers the rival's tax bill by βt_j, so

$$x_i^N = \frac{t_i}{\eta}, \qquad x_i^C = \frac{t_i + \beta t_j}{\eta}. \tag{44.4}$$

For given taxes, the cartel raises joint profit by

$$\frac{\beta^2\left(t_1^2+t_2^2\right)}{2\eta},$$

so firms always want to form it (Section 6.5 of the paper). Abatement responds only to taxes: without a tax, firms do no end-of-pipe R&D.

44.4 Second Step: Emission Taxes

Consider that $\beta = 1$. Differentiating (44.2) and using (44.3) and (44.4), the first-order condition of country i under RJV competition is

$$\underbrace{\left(t_i-\frac{e_i}{2}\right)\left(-\frac{2}{3}-\frac{1}{\eta}\right)}_{\text{Pigouvian term}} \underbrace{-\frac{q_i}{3}}_{\text{rent shifting}} = 0, \tag{44.5}$$

and under an RJV cartel it gains two terms,

$$\left(t_i-\frac{e_i}{2}\right)\left(-\frac{2}{3}-\frac{1}{\eta}\right)-\frac{q_i}{3} \underbrace{-\frac{t_j}{\eta}}_{\text{over-investment}} \underbrace{+\frac{e_i}{2}\cdot\frac{1}{\eta}}_{\text{environmental spillover}} = 0. \tag{44.6}$$

The rent-shifting term pushes the tax below marginal damage: a lower tax makes firm i more aggressive in the third country and shifts profit from the rival. The over-investment term pushes the tax down further: a higher t_i raises firm i's abatement, which in a cartel already exceeds what country i would choose, because the cartel also counts the savings on the rival's tax bill. The environmental spillover term pushes the tax up, since the rival's extra abatement also cleans country i's air. Solving (44.5) and (44.6) at symmetric taxes gives

$$t^{NJ}=\frac{\eta}{(\eta+2)(4\eta+3)}, \qquad t^{CJ}=\frac{\eta}{(\eta+4)(2\eta+3)}. \tag{44.7}$$

Result 1 (emission taxes, Proposition 1 of the paper). With RJV competition and with an RJV cartel:

- each tax is below marginal damage $\frac{e_i}{2}$;
- the tax is lower under the RJV cartel, $t^{CJ} < t^{NJ}$, if and only if $\eta < \sqrt{3}$, since

$$t^{NJ}-t^{CJ}=\frac{2\eta(3-\eta^2)}{(\eta+2)(4\eta+3)(\eta+4)(2\eta+3)}.$$

Figure 44.1 depicts the two taxes. At $\eta = 1$, $t^{NJ} = \frac{1}{21}$, below marginal damage $\frac{1}{9}$, and $t^{CJ} = \frac{1}{25}$, below marginal damage $\frac{2}{25}$. The curves cross at $\eta = \sqrt{3}$.

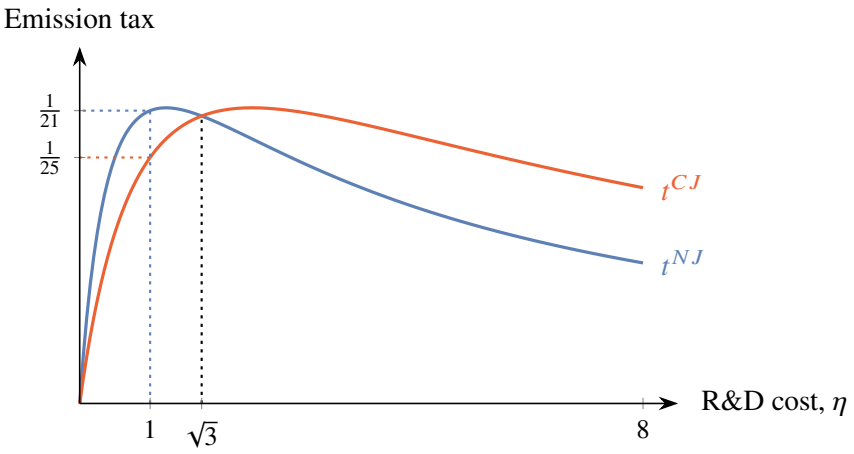


Figure 44.1: Emission taxes under RJV competition (blue) and under an RJV cartel (red).

Intuitively, when R&D is inexpensive, the cartel invests heavily, and each government cuts its tax to restrain an investment that partly serves the rival's tax bill; the over-investment effect dominates. When R&D is costly, investment is small, and each government raises its tax to induce more abatement by the rival, whose results also reduce domestic emissions; the environmental spillover effect dominates.

44.5 Third Step: R&D

Solving for the equilibrium taxes at any spillover β and substituting them into (44.4) gives

$$\begin{aligned} x^N(\beta) &= \frac{1}{2\beta\eta + 3\beta + 4\eta^2 + 9\eta + 3}, \\ x^C(\beta) &= \frac{(1+\beta)(1+\beta^2)}{3\beta^4 + 6\beta^3 + 3\beta^2\eta + 6\beta^2 + 10\beta\eta + 6\beta + 4\eta^2 + 9\eta + 3}, \end{aligned} \tag{44.8}$$

and, with $\beta = 1$,

$$x^{NJ} = \frac{1}{(\eta+2)(4\eta+3)} \qquad \text{and} \qquad x^{CJ} = \frac{2}{(\eta+4)(2\eta+3)}.$$

Result 2 (R&D, Lemma 1 and Proposition 2 of the paper). In equilibrium:

- under R&D competition, abatement falls with the spillover β, so $x^{NJ} < x^N(\beta)$ for every $\beta < 1$;
- the cartel invests at least as much as competing firms, $x^C(\beta) \geq x^N(\beta)$, with equality only at $\beta = 0$;
- the RJV cartel invests more than RJV competition, $x^{CJ} > x^{NJ}$, for every η.

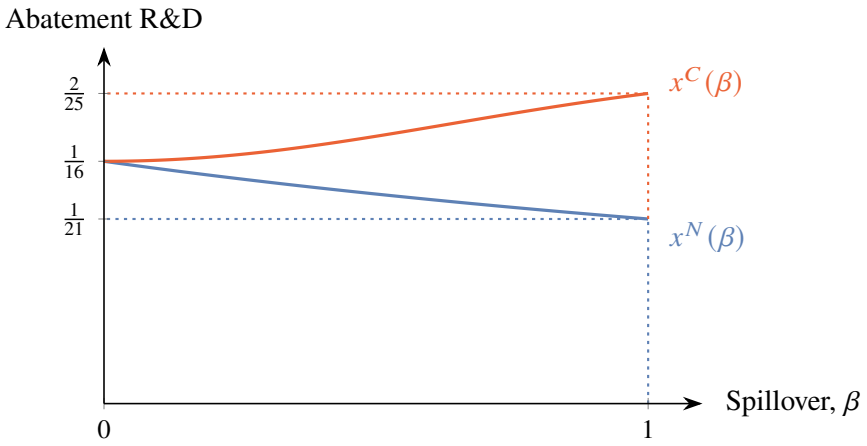


Figure 44.2: Equilibrium abatement R&D under R&D competition (blue) and under an R&D cartel (red), with $\eta = 1$.

Figure 44.2 depicts abatement as a function of the spillover at $\eta = 1$. The two formations coincide at $\beta = 0$, where each firm invests $\frac{1}{16}$. As β rises, competing firms invest less, down to $\frac{1}{21}$, while the cartel invests more, up to $\frac{2}{25}$.

Intuitively, competing firms ignore the tax savings their abatement creates for the rival, and a higher spillover lowers the tax each government sets, so they invest less. The cartel internalizes these savings. In much of the literature without strategic taxes, a cartel invests less than competing firms when spillovers are small (d'Aspremont and Jacquemin, 1988); here, it never does.

44.6 Fourth Step: National Welfare

Substituting the equilibrium into (44.2), the RJV cartel changes welfare in each exporting country by

$$SW^{CJ} - SW^{NJ} = \underbrace{\Delta(pq_i)}_{\text{revenue}} - \underbrace{\Delta\left(\tfrac{\eta}{2}x_i^2\right)}_{\text{R\&D cost, }>0} - \underbrace{\Delta\left(\tfrac{1}{4}e_i^2\right)}_{\text{damage, }<0}, \tag{44.9}$$

so the cartel spends more on R&D and pollutes less, while its effect on revenue depends on taxes.

Result 3 (national welfare, Proposition 3 of the paper). Each exporting country is better off under RJV competition than under an RJV cartel, $SW^{NJ} > SW^{CJ}$, if and only if $\eta < 0.273$, the unique positive root of

$$16\eta^6 + 210\eta^5 + 980\eta^4 + 2076\eta^3 + 1907\eta^2 + 360\eta - 288 = 0. \tag{44.10}$$

Intuitively, when R&D is very inexpensive, governments set a low tax under the cartel to curb investment. The low taxes expand both firms' output in the third country, which lowers the price and each exporter's revenue, and this loss outweighs the lower damage. The RJV cartel then invests more in R&D than RJV competition but yields lower national welfare, in contrast with the conventional result that an RJV cartel maximizes welfare (Kamien et al., 1992).

44.7 Fifth Step: The Importing Country and Global Welfare

Consumers in the third country obtain $CS = \frac{1}{2}(q_1 + q_2)^2$, and total output is $\frac{2-t_1-t_2}{3}$, so they gain from the RJV cartel if and only if it lowers taxes, that is, if and only if $\eta < \sqrt{3}$. Global welfare is $GW = SW_1 + SW_2 + CS$.

Result 4 (global welfare, Propositions 4 and 5 of the paper). Comparing the RJV cartel with RJV competition:

- if $\eta < 0.273$, each exporting country prefers RJV competition, but global welfare is higher with the RJV cartel;
- if $0.273 < \eta < 16.28$, the RJV cartel raises both national and global welfare;
- if $\eta > 16.28$, each exporting country prefers the RJV cartel, but global welfare is higher with RJV competition, where 16.28 is the unique positive root of

$$16\eta^6 - 126\eta^5 - 1772\eta^4 - 6228\eta^3 - 9225\eta^2 - 5616\eta - 864 = 0.$$

Table 44.1 reports one value of η in each interval: at $\eta = \frac{1}{5}$, national welfare is higher under RJV competition and global welfare under the cartel; at $\eta = 1$, both are higher under the cartel; and at $\eta = 20$, national welfare is higher under the cartel and global welfare under RJV competition. Figure 44.3 depicts the percentage gain from the RJV cartel in national welfare (blue) and in global welfare (red); the first is negative only for $\eta < 0.273$, and the second only for $\eta > 16.28$.

Table 44.1: RJV competition (NJ) and RJV cartel (CJ) at three R&D costs.

	$\eta = \frac{1}{5}$		$\eta = 1$		$\eta = 20$	
Outcome	NJ	CJ	NJ	CJ	NJ	CJ
Emission tax	0.0239	0.0140	0.0476	0.0400	0.0110	0.0194
Abatement R&D	0.1196	0.1401	0.0476	0.0800	0.0005	0.0019
Emissions per country	0.0861	0.0486	0.2222	0.1600	0.3286	0.3230
Welfare per exporting country	0.11036	0.11007	0.10242	0.10560	0.08531	0.08706
Consumer surplus, third country	0.21172	0.21604	0.20156	0.20480	0.21738	0.21369
Global welfare	0.43243	0.43619	0.40640	0.41600	0.38799	0.38782

Intuitively, the third country's consumers are left out of each exporting government's calculation. When R&D is inexpensive, the cartel lowers taxes, which raises output and helps these consumers, so global welfare favors the cartel even when national welfare does not. When R&D is very costly, the cartel raises taxes enough that these consumers lose more than the exporting countries gain. Hence the paper's policy conclusion: the decision to allow international RJV cartels should be taken by all governments together, including the importing one.

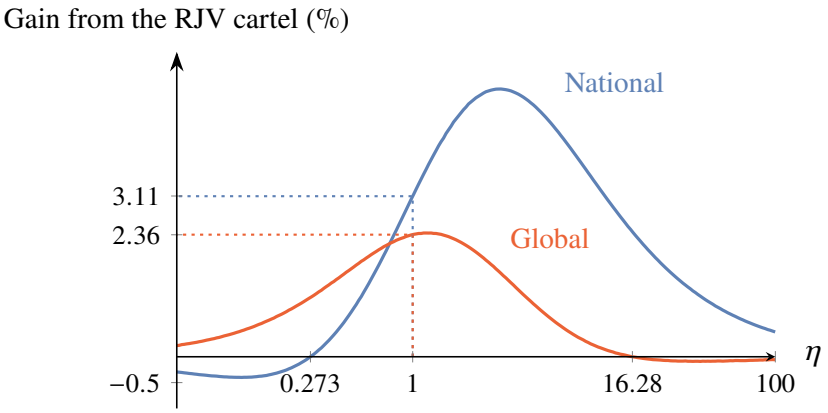


Figure 44.3: Percentage change in welfare per exporting country (blue) and in global welfare (red) when firms form an RJV cartel instead of competing in R&D; logarithmic horizontal axis.

44.8 Sixth Step: Extension to Larger Environmental Damage

Let damage be $\frac{\gamma}{2}e_i^2$ with $\gamma \geq \frac{1}{2}$, as in the paper; the toy model is the case $\gamma = \frac{1}{2}$. The taxes become

$$t^{NJ} = \frac{\eta(2\eta\gamma - \eta + 3\gamma)}{2\eta^2\gamma + 5\eta^2 + 15\eta\gamma + 9\eta + 18\gamma}, \qquad t^{CJ} = \frac{\eta(2\eta\gamma - \eta + 6\gamma)}{2\eta^2\gamma + 5\eta^2 + 30\eta\gamma + 18\eta + 72\gamma}. \tag{44.11}$$

Result 5 (larger damage, Propositions 1, 3, and 4 of the paper). With $\gamma \geq \frac{1}{2}$:

- the RJV cartel lowers the tax if $\eta < \sqrt{3}$, or if γ exceeds a cutoff $\gamma_t(\eta)$ that equals $\frac{1}{2}$ at $\eta = \sqrt{3}$, increases in η, and stays below $\frac{1+\sqrt{7}}{4} \approx 0.911$, so with $\gamma \geq \frac{1+\sqrt{7}}{4}$ the cartel always lowers the tax;
- the range of R&D costs where exporting countries prefer RJV competition widens as γ rises up to about 25.3, from $\eta < 0.273$ at $\gamma = \frac{1}{2}$ to $\eta < 0.686$ at $\gamma = 1$, and never extends beyond $\eta = 9.865$;
- global welfare is higher with RJV competition only if $\gamma < 0.598$ and η is large, at least 16.28.

Figure 44.4 depicts the three cutoffs in the (η, γ) plane for $\gamma \in \left[\frac{1}{2}, 1\right]$. Exporting countries prefer RJV competition to the left of the green curve; the cartel lowers the tax above the blue curve and at every $\eta < \sqrt{3}$; and global welfare is higher with RJV competition below the red curve.

Intuitively, larger damage raises taxes and therefore R&D, which strengthens the over-investment effect, so the cartel lowers the tax unless R&D is costly and damage is small. For much larger γ, outside the range of Figure 44.4, the green curve bends back: the paper shows that for η between $4\sqrt{6} \approx 9.80$ and 9.865, exporting countries prefer RJV competition only for intermediate values of γ.

44.9 Seventh Step: What the Toy Model Keeps and Drops

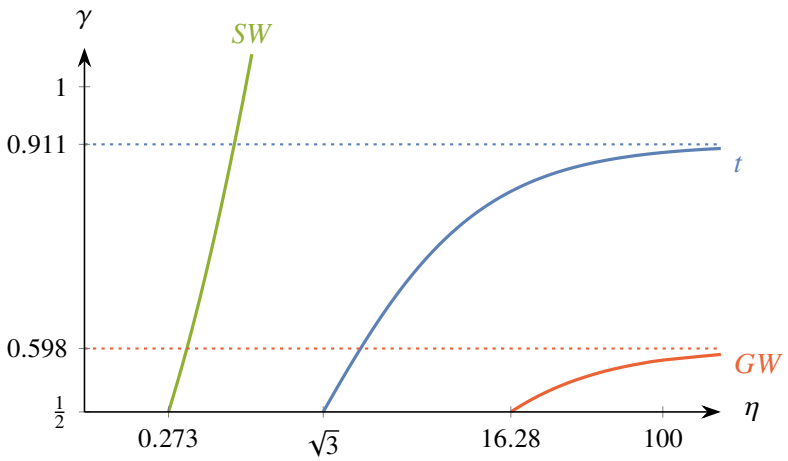


Figure 44.4: Cutoffs for the comparison of taxes (blue), national welfare (green), and global welfare (red) in the (η, γ) plane; logarithmic horizontal axis. Dotted lines mark the limits of the blue and red curves as η grows.

What the toy model keeps.

- The three forces behind each government's tax: rent shifting, over-investment, and environmental spillover.
- The ranking of taxes, with its cutoff $\eta = \sqrt{3}$, and the higher R&D of cartels.
- The welfare results: an RJV cartel does not always maximize national welfare, and national and global welfare can disagree in both directions.

What it drops.

- General damage, except in the sixth step: the toy model fixes $\gamma = \frac{1}{2}$, so several cutoffs of the paper reduce to single numbers.
- Sales in the exporting countries themselves (Section 7 of the paper), where governments also care about their own consumers; the results survive when the third market is larger than the home markets, but not when all three markets are equal.
- The case where governments set taxes after firms choose their R&D formation, in which firms may prefer to compete in R&D to induce higher taxes.
- Transboundary pollution, which the paper reports does not change the results qualitatively.

44.10 Connections to Other Chapters

- Chapter 37 studies the same four arrangements in a closed economy without pollution, where the RJV cartel yields the highest total surplus.
- Chapter 43 studies a committed emission tax in a closed economy with any number of firms, and shows how the number of firms changes total abatement R&D.

- Chapter 45 studies a regulator that responds to firms' investments when it sets the emission fee, which makes cost-reducing and green R&D strategic complements.
- Chapter 15 studies the effect of a merger on innovation, where coordination leads the merging firms to innovate less.
- Chapter 42 studies the closed-economy version: an environmental R&D cartel when the regulator sets the emission tax after firms invest.
- Chapter 41 compares emission taxes, permits, and standards by the incentives they give firms to adopt an advanced abatement technology.

References

d'Aspremont, C., and Jacquemin, A. (1988). Cooperative and noncooperative R&D in duopoly with spillovers. *American Economic Review*, 78(5), 1133–1137.

Iida, T. (2020). International R&D formations and strategic environmental policy. *Environment and Development Economics*, 25(5), 509–528.

Kamien, M. I., Muller, E., and Zang, I. (1992). Research joint ventures and R&D cartels. *American Economic Review*, 82(5), 1293–1306.

Chapter 45

Greener or Cheaper: Economies of Scope in R&D

Companion note to Strandholm, J. C., A. Espinola-Arredondo, and F. Munoz-Garcia (2025), "Greener or Cheaper Goods: Economies of Scope in R&D Investments," Journal of Environmental Economics and Management 130: 103107.

45.1 Introduction

Consider a car maker with one research lab. The lab can make the engine cheaper to build, or it can make the engine cleaner. Firms make both kinds of investment, yet the literature on environmental policy usually studies one at a time: abatement under an emission fee, or cost-reducing R&D without environmental policy. Do the two investments interact, and what does a regulator miss when it models only one? Strandholm, Espinola-Arredondo, and Munoz-Garcia (2025) study these questions.

The two investments interact through two channels. The first is the emission fee. A firm that invests in R&D produces more and pollutes more, so the regulator raises the fee; a firm that invests in abatement pollutes less, so the regulator lowers it. Under regulation, each investment therefore raises the firm's return to the other: the two are *strategic complements*, even when their costs are independent. The second channel is *economies of scope*: progress in one investment lowers the cost of the other. Economies of scope give firms a reason to abate even without regulation, since abating makes their R&D cheaper.

These links matter for policy. A regulator who sets the fee from a model with abatement only ignores the R&D that expands output, so it underestimates the pollution of the industry. When economies of scope are weak, its fee is too lax (undertaxation); when they are strong, the error reverses and its fee is too stringent.

The toy model has two Cournot firms, linear demand, and quadratic investment costs, with one parameter that measures economies of scope. The first step solves the last two stages, output and

the emission fee. The second and third steps solve investments without and with regulation, and the fourth step compares the fee with the one from a model with abatement only. The fifth step lets firms choose abatement before R&D, and the sixth step lists what the toy model keeps and drops.

45.2 The Toy Model

Market. Two firms, i and j, sell a homogeneous good with inverse demand $p(Q) = 1 - Q$, where $Q = q_i + q_j$. Firm i's marginal cost is $\frac{1}{2} - k_i$, where k_i denotes its investment in cost-reducing R&D. Each unit of output generates one unit of emissions, and firm i's investment in abatement, z_i, lowers its emissions to $e_i = q_i - z_i$. Environmental damage is

$$ED = \frac{3}{2}(Q - Z)^2, \tag{45.1}$$

where $K = k_i + k_j$ and $Z = z_i + z_j$ denote aggregate R&D and aggregate abatement.

Investment costs. Firm i's cost of investing (k_i, z_i) is

$$C(k_i, z_i) = k_i^2 + \frac{3}{4}z_i^2 - \lambda k_i z_i, \qquad \lambda \in [-1, 1]. \tag{45.2}$$

When $\lambda = 0$, the two investments have independent costs. When $\lambda > 0$, they exhibit economies of scope: the marginal cost of R&D, $2k_i - \lambda z_i$, decreases in abatement, and the marginal cost of abatement, $\frac{3}{2}z_i - \lambda k_i$, decreases in R&D. When $\lambda < 0$, they exhibit diseconomies of scope. In the paper's notation, the toy model sets $n = 2$ firms, a marginal cost $c = \frac{1}{2}$, an R&D cost parameter $\gamma = 2$, an abatement cost parameter $\alpha = \frac{3}{2}$, and pollution severity $d = \frac{3}{2}$. The interval $\lambda \in [-1, 1]$ guarantees that every firm's investment problem is strictly concave and that all equilibrium investments and fees under regulation are positive.

Timing. In the first stage, both firms simultaneously choose (k_i, z_i). In the second stage, the regulator observes all investments and sets an emission fee t per unit of emissions. In the third stage, firms compete à la Cournot. The regulator maximizes social welfare,

$$SW = CS + PS + T - ED, \tag{45.3}$$

where $CS = \frac{1}{2}Q^2$ is consumer surplus, PS denotes profits net of fees and investment costs, and $T = t(Q - Z)$ is fee revenue. As a benchmark, the second step considers an unregulated industry, $t = 0$.

45.3 First Step: Output and the Emission Fee

Third stage. Given investments and fee t, firm i chooses q_i to maximize

$$(1 - Q)q_i - \Big(\frac{1}{2} - k_i\Big)q_i - t(q_i - z_i).$$

Solving the two first-order conditions yields

$$q_i = \frac{\frac{1}{2} - t + 2k_i - k_j}{3}, \qquad \text{and} \qquad Q = \frac{1 - 2t + K}{3}. \tag{45.4}$$

Firm i's output increases in its own R&D, decreases in its rival's R&D, and decreases in the fee. Its profit, including investment costs, is

$$\pi_i = q_i^2 + tz_i - C(k_i, z_i), \tag{45.5}$$

where tz_i captures the fee that firm i avoids by abating.

Second stage. An increase in t lowers every firm's output by $\frac{1}{3}$, so the regulator's first-order condition equates price to the average marginal cost plus the marginal damage of emissions,

$$1 - Q = \left(\frac{1}{2} - \frac{K}{2}\right) + 3(Q - Z), \qquad \text{which yields} \qquad Q = \frac{1 + K + 6Z}{8}. \tag{45.6}$$

Inserting Q from (45.4) and solving for t gives the emission fee.

Result 1 (emission fee). The welfare-maximizing fee is

$$t(K, Z) = \frac{5}{16} + \frac{5}{16}K - \frac{9}{8}Z, \tag{45.7}$$

which increases in aggregate R&D and decreases in aggregate abatement. The fee is separable in K and Z: the effect of one more unit of abatement on the fee does not depend on R&D.

Intuitively, R&D lowers marginal costs and expands output, so the industry pollutes more and the regulator responds with a more stringent fee. Abatement, instead, lowers emissions at every output level, so the regulator can afford a less stringent fee. This gives rise to two externalities across firms that a model with a single investment cannot capture: firm i's R&D raises the fee that its rival pays, and firm i's abatement lowers it. For the general model, the fee is

$$t(K, Z) = \frac{(2dn - 1)[n(1 - c) + K] - 2dn(n + 1)Z}{(2d + 1)n^2},$$

which reduces to (45.7) at the toy-model parameters.

45.4 Second Step: Investments without Regulation

Without regulation, $t = 0$, so firm i chooses (k_i, z_i) to maximize

$$\left(\frac{\frac{1}{2} + 2k_i - k_j}{3}\right)^2 - \left(k_i^2 + \frac{3}{4}z_i^2 - \lambda k_i z_i\right).$$

The first-order conditions are

$$\frac{4}{9}\left(\frac{1}{2} + 2k_i - k_j\right) - 2k_i + \lambda z_i = 0 \qquad \text{and} \qquad \lambda k_i - \frac{3}{2}z_i \le 0. \tag{45.8}$$

Abatement only lowers the cost of R&D, so $z_i = \frac{2\lambda}{3} k_i$ when $\lambda > 0$, and $z_i = 0$ otherwise. Imposing symmetry yields the next result.

Result 2 (no regulation). Equilibrium investments without regulation are:

- If $\lambda \leq 0$, then $k^{NR} = \frac{1}{7}$ and $z^{NR} = 0$.
- If $\lambda > 0$, then
$$k^{NR} = \frac{1}{7 - 3\lambda^2} \qquad \text{and} \qquad z^{NR} = \frac{2\lambda}{3(7 - 3\lambda^2)},$$
both increasing in λ.

Intuitively, an unregulated firm does not pay for its emissions, so abatement has no value of its own. With economies of scope, however, abatement lowers the marginal cost of R&D, and the firm abates to lower the cost of its R&D. This motive for green investment does not rely on emission fees, green consumers, or corporate social responsibility.

45.5 Third Step: Investments under Regulation

In the first stage, firm i anticipates output (45.4) and fee (45.7), and chooses (k_i, z_i) to maximize profit (45.5). Solving the first-order conditions for k_i and z_i, respectively, yields the best responses

$$k_i = \frac{9}{175} + \frac{94 + 128\lambda}{175} z_i - \frac{9}{25} k_j + \frac{54}{175} z_j, \tag{45.9}$$

$$z_i = \frac{23}{222} + \frac{47 + 64\lambda}{222} k_i - \frac{1}{222} k_j - \frac{9}{37} z_j. \tag{45.10}$$

Firm i's two investments are strategic complements: R&D increases in own abatement, and abatement increases in own R&D, whenever $\lambda > -\frac{47}{64}$, which includes the case without economies of scope, $\lambda = 0$. Intuitively, abatement lowers the fee, which raises output and, thus, the return of every unit of cost reduction; and R&D raises the fee, which raises the fee that every unit of abatement avoids. Economies of scope strengthen both effects. In contrast, each investment is a strategic substitute of the rival's investment of the same type.

Imposing symmetry, $k_i = k_j = k$ and $z_i = z_j = z$, in (45.9) and (45.10) gives the two loci

$$k = \frac{9 + (148 + 128\lambda)z}{238} \qquad \text{and} \qquad z = \frac{23 + (46 + 64\lambda)k}{276}, \tag{45.11}$$

whose intersection is the equilibrium. Figure 45.1 depicts both loci at $\lambda = 0$ and $\lambda = \frac{1}{2}$. Their intercepts, $k = \frac{9}{238}$ when $z = 0$ and $z = \frac{1}{12}$ when $k = 0$, are the investments of models with a single type of R&D, which the fourth step uses as benchmarks. When $\lambda > -\frac{23}{32}$, as in the figure, both loci are positively sloped, so the equilibrium lies above and to the right of both intercepts. An increase in λ rotates both loci, moving the equilibrium to the northeast.

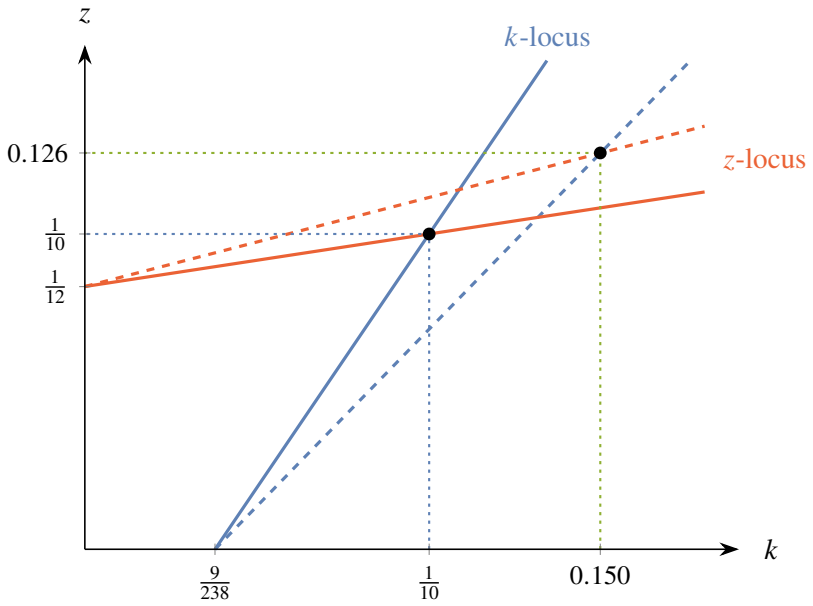


Figure 45.1: Symmetric best-response loci under regulation, at $\lambda = 0$ (solid) and $\lambda = \frac{1}{2}$ (dashed). Black dots mark the equilibria.

Result 3 (investments under regulation). Each firm's equilibrium investments and the resulting fee are

$$k^R = \frac{23(2+\lambda)}{4(115-30\lambda-16\lambda^2)}, \qquad z^R = \frac{92+9\lambda}{8(115-30\lambda-16\lambda^2)}, \qquad t^R = \frac{276-133\lambda-80\lambda^2}{16(115-30\lambda-16\lambda^2)}. \tag{45.12}$$

Without economies of scope, $\lambda = 0$, they become $k^R = z^R = \frac{1}{10}$ and $t^R = \frac{3}{20}$. Both investments increase in λ. Relative to no regulation:

- Abatement is higher under regulation, $z^R > z^{NR}$, for all $\lambda \in [-1, 1]$.
- R&D is lower under regulation, $k^R < k^{NR}$, if and only if

$$\lambda < \lambda_k \equiv \frac{\sqrt{58537} - 143}{138} \approx 0.717.$$

Figure 45.2 depicts equilibrium investments with and without regulation as a function of λ. Without regulation, abatement is zero until economies of scope appear, and R&D is flat at $\frac{1}{7}$ for $\lambda \leq 0$. Under regulation, both investments are positive at every λ, and R&D is lower than without regulation unless economies of scope are strong.

Intuitively, the fee makes output more expensive, which lowers the return of cost reduction; but it also makes abatement more valuable, and, when λ is high, more abatement lowers the cost of R&D enough to offset the first effect.

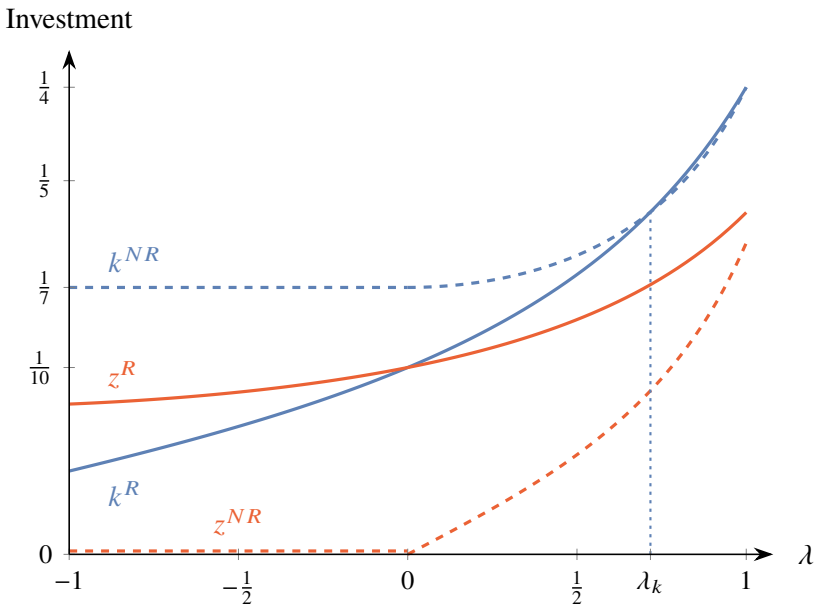


Figure 45.2: Equilibrium R&D and abatement under regulation (solid) and without regulation (dashed).

45.6 Fourth Step: Undertaxation

Consider now a regulator who designs the fee using a standard model where firms only invest in abatement, as in Poyago-Theotoky (2007). Setting $k = 0$ in (45.11) gives abatement $\bar{z} = \frac{1}{12}$ per firm and, from (45.7), the fee

$$\bar{t} = \frac{5}{16} - \frac{9}{8} \cdot \frac{1}{6} = \frac{1}{8}, \tag{45.13}$$

which does not depend on λ, since the cross term $\lambda k_i z_i$ vanishes when $k_i = 0$. Similarly, a model where firms only invest in R&D predicts $\bar{k} = \frac{9}{238} \approx 0.038$. At $\lambda = 0$, both one-investment models underestimate investments, $\bar{z} < z^R$ and $\bar{k} < k^R = \frac{1}{10}$, because they omit the complementarity between them.

Result 4 (undertaxation). The fee gap is

$$t^R - \bar{t} = \frac{(2+\lambda)(23-48\lambda)}{16(115-30\lambda-16\lambda^2)}, \tag{45.14}$$

so that:

- If $\lambda < \frac{23}{48}$, then $t^R > \bar{t}$: the abatement-only model prescribes a fee that is too lax (undertaxation).
- If $\lambda > \frac{23}{48}$, then $t^R < \bar{t}$: the abatement-only model prescribes a fee that is too stringent (overtaxation).

In addition, t^R is decreasing in λ for all $\lambda > -0.718$.

Figure 45.3 depicts both fees. At $\lambda = 0$, the gap is $\frac{1}{40}$, and it can be decomposed using (45.7),

which, per firm, reads

$$t = \frac{5}{16} + \frac{5}{8}k - \frac{9}{4}z.$$

R&D raises the fee by $\frac{5}{8} \cdot \frac{1}{10}$, while the extra abatement, $\frac{1}{10} - \frac{1}{12} = \frac{1}{60}$ per firm, lowers it by $\frac{9}{4} \cdot \frac{1}{60}$, so that

$$\frac{5}{8} \cdot \frac{1}{10} - \frac{9}{4} \cdot \frac{1}{60} = \frac{1}{16} - \frac{3}{80} = \frac{1}{40}.$$

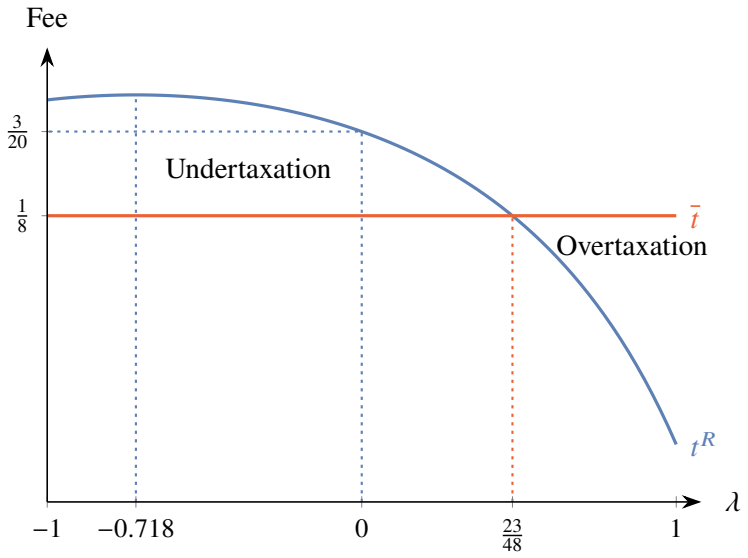


Figure 45.3: Equilibrium fee t^R and the fee from an abatement-only model, $\bar{t}$.

Intuitively, the abatement-only model ignores the R&D that lowers production costs and raises emissions, so it expects less pollution than the industry generates, and sets a lax fee. As λ increases, both investments increase, but abatement lowers the fee more than R&D raises it: from $\lambda = 0$ to $\lambda = \frac{1}{2}$, k^R grows by 0.050 and z^R by 0.026, which changes the fee by

$$\frac{5}{8} \cdot 0.050 - \frac{9}{4} \cdot 0.026 \approx -0.027.$$

Economies of scope, therefore, ameliorate undertaxation, and reverse it when they are strong enough.

45.7 Fifth Step: Sequential Investments

Consider now that firms choose abatement first, and R&D second, before the regulator sets the fee and firms compete. Abatement then works as a commitment: when firm i chooses z_i, it anticipates how both firms' R&D responds.

Second stage: R&D given abatement. Solving both firms' first-order conditions for R&D given (z_i, z_j) yields

$$k_i^{NR}(z_i, z_j) = \frac{2 + 15\lambda z_i - 6\lambda z_j}{14} \qquad \text{and} \qquad k_i^{R}(z_i, z_j) = \frac{18 + (233 + 400\lambda)z_i + (63 - 144\lambda)z_j}{476}, \tag{45.15}$$

without and with regulation, respectively. Without regulation, firm i's abatement lowers its rival's R&D when $\lambda > 0$,

$$\frac{\partial k_j^{NR}}{\partial z_i} = -\frac{3\lambda}{7} < 0.$$

Its abatement lowers the cost of firm i's R&D, firm i invests more in R&D, and its rival responds with less. With regulation,

$$\frac{\partial k_j^{R}}{\partial z_i} = \frac{9(7 - 16\lambda)}{476}, \tag{45.16}$$

which is positive if and only if $\lambda < \frac{7}{16}$. Firm i's abatement now has a second effect: it lowers the fee that both firms pay, which raises the rival's output and its return to R&D. This effect dominates when economies of scope are weak.

First stage: abatement. By the envelope theorem, firm i's first-order condition for abatement is

$$\underbrace{\frac{\partial \pi_i}{\partial z_i}}_{\text{Direct effect}} + \underbrace{\frac{\partial \pi_i}{\partial k_j}\frac{\partial k_j}{\partial z_i}}_{\text{Strategic effect}} = 0, \tag{45.17}$$

where the direct effect alone gives the simultaneous-investment condition. Since the rival's R&D hurts firm i, $\frac{\partial \pi_i}{\partial k_j} < 0$, the strategic effect has the opposite sign of $\frac{\partial k_j}{\partial z_i}$, as in Brander and Spencer (1983).

Result 5 (sequential investments). When firms choose abatement before R&D:

- Without regulation and with $\lambda > 0$, abatement is

$$z^{NR,Seq} = \frac{20\lambda}{3(49 - 23\lambda^2)},$$

  which exceeds simultaneous abatement z^{NR}; both are zero when $\lambda \leq 0$.

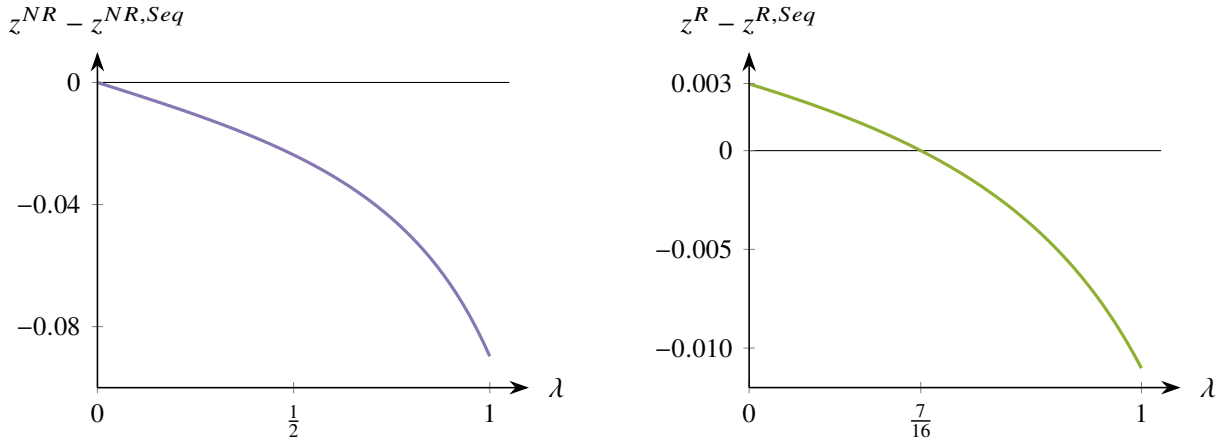
- With regulation, abatement is

$$z^{R,Seq} = \frac{5(613 + 90\lambda)}{31721 - 9105\lambda - 4496\lambda^2}.$$

  If and only if $\lambda < \frac{7}{16}$, it is below simultaneous abatement z^R and yields a more stringent fee, $t^{R,Seq} > t^R$.

Figure 45.4 depicts the abatement differential between simultaneous and sequential investments without regulation (panel a) and with regulation (panel b). Without regulation, the differential is

negative: firms over-invest in abatement to commit to more R&D, which discourages their rival's R&D. With regulation and weak economies of scope, the differential is positive: an extra unit of abatement lowers the fee and encourages the rival's R&D, so firms abate less when they anticipate this response. At $\lambda = 0$, sequential abatement is $z^{R,Seq} = \frac{3065}{31721} \approx 0.097$, against $z^{R} = \frac{1}{10}$, and the fee rises from 0.150 to 0.156. When $\lambda > \frac{7}{16}$, the commitment effect of the unregulated industry prevails again.



(a) No regulation.

(b) Regulation.

Figure 45.4: Abatement under simultaneous minus sequential investments.

45.8 Sixth Step: What the Toy Model Keeps and Drops

Table 45.1 reports equilibrium outcomes at $\lambda = 0$ and $\lambda = \frac{1}{2}$.

Table 45.1: Equilibrium outcomes ($n = 2$, $c = \frac{1}{2}$, $\gamma = 2$, $\alpha = d = \frac{3}{2}$); k and z are per firm.

Setting	λ	k	z	t	$Q - Z$	SW
No regulation	0	0.1429	0	0	0.4286	−0.1327
Regulation, simultaneous	0	0.1000	0.1000	0.1500	0.1000	0.0850
Regulation, sequential	0	0.0979	0.0966	0.1563	0.1012	0.0842
No regulation	$\frac{1}{2}$	0.1600	0.0533	0	0.3333	−0.0200
Regulation, simultaneous	$\frac{1}{2}$	0.1497	0.1257	0.1234	0.0996	0.1018
Regulation, sequential	$\frac{1}{2}$	0.1503	0.1263	0.1222	0.0994	0.1020
Abatement only	any	0	0.0833	0.1250	0.0833	0.0729

What the toy model keeps.

- The emission fee increases in R&D and decreases in abatement, creating two externalities across firms.
- Without regulation, firms abate only when the investments exhibit economies of scope.
- Under regulation, R&D and abatement are strategic complements, even without economies of scope.
- A regulator who uses a model with abatement only sets a lax fee (undertaxation), unless economies of scope are strong.
- Regulation raises abatement and, unless economies of scope are strong, lowers R&D.

What the toy model drops.

- The general number of firms n and the general parameters c, γ, α, and d, so it cannot show how undertaxation varies with market concentration or pollution severity.
- Asymmetric economies of scope, and spillovers or cooperation in R&D, which the paper lists as further research.

45.9 Connections to Other Chapters

- Chapter 44 studies abatement R&D in an international setting, where each government sets its emission tax to favor its own exporting firm.
- Chapter 38 shows that a cost reduction can lower welfare when the firm that lowers its cost is inefficient enough.
- Chapter 39 studies how patent length and forward protection affect the timing of R&D.
- Chapter 42 studies firms that coordinate their abatement R&D when the regulator sets the emission tax after observing their investment.
- Chapter 41 compares emission taxes, permits, and standards by the incentives they give firms to adopt an advanced abatement technology.

References

Brander, J. A., and Spencer, B. J. (1983). Strategic commitment with R&D: The symmetric case. *Bell Journal of Economics*, 14(1), 225–235.

Poyago-Theotoky, J. A. (2007). The organization of R&D and environmental policy. *Journal of Economic Behavior & Organization*, 62(1), 63–75.

Strandholm, J. C., Espinola-Arredondo, A., and Munoz-Garcia, F. (2025). Greener or cheaper goods: Economies of scope in R&D investments. *Journal of Environmental Economics and Management*, 130, 103107.