

Chapter 1: Introduction to Games and their Representation

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction/Roadmap

- What's game theory.
- Main elements in strategic settings (in games).
- Two main graphical approaches to represent games:
 - Matrices
 - Game trees
- Introducing Imperfect Information in Game Trees
- Identifying Equilibrium Behavior

What is game theory?

Game Theory studies the interaction between a group of rational agents who behave strategically.

Let's analyze each underlined element of this definition.

1. Group of individuals.

- Game theory studies interaction between a group of agents.
- Scenarios with one individual or firms are analyzed with individual decision making, not with games.

What is game theory?

2. *Rationality*. Every agent seeks to maximize her objective function.

- Common knowledge of rationality. In a two-player game:
 - Player 1 seeks to maximize her payoff,
 - Player 2 seeks to maximize her payoff,
 - Player 1 knows that player 2 seeks to maximize her payoff,
 - Player 2 knows that player 1 seeks to maximize her payoff,
 - Player 1 knows that player 2 knows that player 1 seeks to maximize her payoff,...
 - And so on, ad infinitum.
 - Intuitively, everyone can put herself in the shoes of her rival, anticipating her moves.
- Common knowledge of rationality in “The Princess Bride,” movie scene:
www.youtube.com/watch?v=9s0UURBihH8.

What is game theory?

3. *Strategic behavior*. Every agent seeks to maximize a well-defined objective function.
- The objective function can be a utility function of an individual, a profit function of a firm, or a social welfare function of an entire country.
 - This allows for the agent to be selfish (if his objective function only includes his own payoffs) or altruistic (if it contains payoffs from other individuals).
 - Examples of other-regarding preferences:
 - Envy and guilt aversion in Fehr and Schmidt (1999) and Bolton and Ockenfels (2000),
 - meaning that players can be both altruistic and strategic at the same time.

Main elements in a game

1. *Players*. Set of agents, such as individuals, firms, or countries, interacting in a game.
 - For instance, if two firms compete in an industry, we say that the number of players is $N=2$.
 - And if a generic number firms interact, we write that the number of players is $N \geq 2$.
2. *Strategy*. Complete contingent plan, describing which actions a player chooses in each possible situation (each contingency) that she faces along the game.
 - In a two-player game with firms A and B, we denote firm A's strategy as s_A , where $s_A \in S_A$, meaning that firm A selects a specific strategy s_A (such as a price of \$12) from a set of available strategies S_A , known as the "strategy set."

Main elements in a game

2. *Strategy (cont'd).*

- Strategies can be understood as an instruction manual: a player opens the manual, looks for the page describing the actions other players chose, and the manual indicates how to respond, recommending that she chooses a specific action s_A .
- **Discrete strategies.** $S_A = \{6,10\}$ if binary, $S_A = \{1,2, \dots\}$ otherwise.
 - Examples are output levels, or other strategies that require being natural numbers.
- **Continuous strategies.** $S_A = [0,10]$ if bounded meaning that the agent can choose non-integer amounts from 0 to 10 (e.g., $s_A = 5.7$) , or $S_A = [0, +\infty)$ or $S_A = \mathbb{R}_+$ if unbounded.
 - Examples can be prices or bids in an auction, where dollars can be split into cents, or cents of a cent.

Main elements in a game

2. *Strategy (cont'd).*

- **Symmetric strategy sets.** If $S_i = S_j = S$ for every player $j \neq i$.
 - Example: firms competing in the same industry have access to the same technology, so every firm chooses an output level from the same strategy set $S = [0,10]$.
- **Asymmetric strategy sets.** If $S_i \neq S_j$ for at least one player $j \neq i$.
 - Example: firms competing in the same industry have access to different technologies, so $S_i = [0,10]$ but $S_j = [0,75]$.

Main elements in a game

2. *Strategy (cont'd).*

- **Strategy profile.** A list describing the strategies that each player selects, $s = (s_1, s_2, \dots, s_N)$.
 - Examples: $S = (12, 8)$, meaning that firm A chooses $s_A = 12$ and firm B selects $s_B = 8$ units of output.
 - In a setting with N players, a strategy profile is $s \equiv (s_1, s_2, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_N)$.
 - For compactness, we write $s = (s_i, s_{-i})$ where $s_{-i} \equiv (s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_N)$ denotes the strategy profile chosen by player i 's rivals.
- Example:
 - In a setting with four firms, $s = (12, 8, 10, 13)$, so $s = (12, s_{-1})$ where $s_{-1} = (8, 10, 13)$ denotes the output of firm 1's rivals.

Main elements in a game

2. *Strategy (cont'd).*

- Therefore, strategy profile $s = (s_i, s_{-i})$ is an element in the Cartesian product $S_i \times S_{-i}$.
 - In the case of a two-player game where every player has binary strategy sets, this Cartesian product has 2×2 different strategy profiles. Check!
 - Still in a two-player game, where player 1 has k available strategies (i.e., the cardinality of S_1 is k) and player 2 has m available strategies, there are $k \times m$ strategy profiles.
- If players choose from continuous strategy spaces, where every $S_i = \mathbb{R}_+$, then:
 - $S_A \times S_B = \mathbb{R}_+^2$ in the case of two players (A and B); and
 - $S_1 \times S_2 \times \dots \times S_N = \mathbb{R}_+^N$ in the case of N players.

Two graphical approaches

Matrices

		Player 2	
		h	l
Player 1	H	4,4	0,-7
	L	8,1	2,2

Matrix 1.1. Example of a two-player game.

- Player 1 typically chooses rows and is referred to as the “row player” (strategies denoted with uppercase letters, H and L)
- Player 2 chooses columns and is known as the “column player” (strategies denoted with lowercase letters, h and l)

Two graphical approaches

Matrices

		Player 2	
		h	l
Player 1	H	4,4	0,-7
	L	8,1	2,2

Matrix 1.1. Example of a two-player game.

- In Matrix 1.1, if player 1 chooses high prices, H , while player 2 picks low price, l , their payoff corresponds to $(0, -7)$,
 - Indicating that player 1's payoff is 0 and player 2's is -7 .
 - Negative payoffs may represent, for instance, that the player's revenue is lower than her costs, thus losing money.

Two graphical approaches

Matrices

		Player 2	
		h	l
Player 1	H	4,4	0,-7
	L	8,1	2,2

Matrix 1.1. Example of a two-player game.

- This matrix shows that every player has only two available strategies, implying that:

$$S_1 = \{H, L\} \text{ and } S_2 = \{h, l\}$$

so each player faces a binary decision.

Two graphical approaches

Matrices

		Player 2	
		h	l
Player 1	H	4,4	0,-7
	L	8,1	2,2

Matrix 1.1. Example of a two-player game.

- In other games, however:
 - A player could have more strategies than her rival, a rich strategy space with hundreds of possible strategies to choose from, or
 - Continuous strategy spaces (such as when firms compete in prices).

Two graphical approaches

Matrices

		Player 2	
		h	l
Player 1	H	4,4	0,-7
	L	8,1	2,2

Matrix 1.1. Example of a two-player game.

- When players interact choosing their strategies simultaneously but the number of available strategies in S_i or S_j , or both, makes the matrix too large to represent graphically...
 - we will analyze the game keeping in mind that players interact in a context which is, essentially, equivalent to that of a matrix form.

Two graphical approaches

Game tree

- When players act sequentially (one after the other), we frequently use game trees.
 - They can help illustrate who acts when,
 - representing which information a player observes in each part of the game.

Two graphical approaches

Game tree

- Figure 1.1a depicts a game tree, with player 1 acting first at the node, which we can interpret as the “root” of the tree.
- The first mover is often referred to as the leader, and player 2, who observes player 1’s choice and responds with her own action, is known as the follower.




Figure 1.1a. Example of a game tree

Two graphical approaches

Game tree

- In Figure 1.1a, player 1 can choose between H and L .
- If player 1 chooses L , the game is over and payoffs are distributed.
- But if she chooses H , player 2 must respond with either h or l .




Figure 1.1a. Example of a game tree

Two graphical approaches

Game tree

- The nodes at the end of the game tree are denoted as “terminal nodes”.
 - They list the payoffs that players earn at each strategy profile.
- For instance, if the game proceeds through H and then l , we say that the strategy profile is (H, l) .
 - In the context of game trees, it can also be understood as a “path of play.”




Figure 1.1a. Example of a game tree

Two graphical approaches

Game tree, example of strategy profile.





Figure 1.1b. Strategy profile (H, l) .

Two graphical approaches

Game tree, larger tree.



Strategy spaces:

$$S_1 = \{H, M, L\}$$

$$S_2 = \{hh', hl', lh', ll'\}$$

Figure 1.2. Example of a game tree - Larger version.

Two graphical approaches

Game tree, larger tree.

- Player 1, being called on to move at a single node (initial node), has as many strategies as action branches stemming from the initial node,

$$S_1 = \{H, M, L\}.$$

- Player 2 can respond in two different nodes (responding to H or M), entailing four different strategies for this player (two actions in each node times two nodes) which implies

$$S_2 = \{hh', hl', lh', ll'\}$$

- For instance, strategy hl' prescribes that player 2 responds with:
 - h after observing that player 1 chooses H , but
 - with l' when player 1 chooses M .

Two graphical approaches

Strategy spaces:

$$S_1 = \{H, M, L\}$$

$$S_2 = \{hh', hl', lh', ll'\}$$

- The above example illustrates our definition of strategies as complete contingent plans in a sequential context:
 - Each strategy in S_2 serves as an instruction manual for player 2, telling her what to do after player 1 chooses H , after M , and after L .
- But why does player 2's strategy describe how would she respond to M , if for instance player 1 chooses H ?
 - Strategies must be so detailed because player 2's responses at nodes that may not be reached can make player 1 behave in one way or another in the first stage.

Introducing imperfect information in game trees

- **Information set:** An information set connects all nodes that player i cannot distinguish.
 - In addition, these nodes must have the same number of action branches, which must have the same labels.




Figure 1.3. Example of a game tree - Larger version with an information set.

Introducing imperfect information in game trees

Incorrect representation of information sets

- Figure 1.4a:
 - Player 2 chooses between two actions when player 1 selects H (h and l at the top node) but...
 - a different number of actions when player 1 selects M (e.g., h, m , and l at the bottom node).
- Then she would be able to infer the node she is in by looking at the number of available actions on that node.
- In that scenario, player 2 would know what action player 1 chose, so no information set would need to be depicted.




Figure 1.4a. Nodes with different number of actions

Introducing imperfect information in game trees

Incorrect representation of information sets.

- Figure 1.4b: player 2 chooses between h and l after player 1 selects H , but between h and x after player 1 chooses M .
- In this setting, player 2 would be able to infer the node she is at by looking at the specific list of actions she can choose from (either h and l , or h and x), implying that the nodes must not be connected with an information set either.




Figure 1.4b. Nodes with the same number of actions but different action labels

Introducing imperfect information in game trees

- Another interesting point of information sets is that they represent a game which, despite being sequential in nature, has every player choosing her actions “in the dark.”
- In this context, a player cannot observe her rival’s previous choices.
- As a consequence, the sequential-move game in Figure 1.3 becomes strategically equivalent to a simultaneous-game.
- This applies regardless of whatever player 1’s and 2’s choices happen simultaneously (in the same second) or sequentially (days or years apart) as long as player 2 cannot observe whether player 1 chose H or M when called to respond.

Identifying equilibrium behavior

- We now seek to solve games, trying to predict how players will behave.
- We will use different solution concepts:
 - From deleting strictly dominated strategies, to Nash equilibrium, and to Subgame perfect equilibrium in complete-information games; and
 - From Bayesian Nash Equilibrium, to Perfect Bayesian Equilibrium or Sequential Equilibrium in incomplete-information games.

Identifying equilibrium behavior

- We evaluate solution concepts according to different criteria:
 - *Existence:*
 - Does an equilibrium exist for all games when we apply that solution concept?
 - Or are there some games where our equilibrium concept does not yield an equilibrium?
 - *Uniqueness:*
 - Does the solution concept we are using predict a unique equilibrium for *all* games?
 - Or are there games out there where our equilibrium concept predicts two or more equilibria.

Identifying equilibrium behavior

- We evaluate solution concepts according to different criteria:
 - *Robustness*:
 - Is the equilibrium (or equilibria) we found robust to small payoff changes?
 - This is, of course, assuming that the payoff *ranking* is unaltered for each player.
 - If that happens, we seek to test if the originally prescribed equilibrium “suddenly changes” when we change the payoff of one of the players slightly.
 - *Pareto optimality*:
 - Is the equilibrium (or equilibria) we found Pareto optimal, in the sense that:
 - we cannot find another strategy profile that strictly increases the payoff of at least one player without reducing the payoff of all other players.
 - If we can find such a strategy profile, the equilibrium we found is *not* Pareto optimal.
 - If we cannot find such a strategy profile, the equilibrium we found *is* Pareto optimal.

We are ready to play some games!

Chapter 2: Equilibrium Dominance

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- In most solution concepts, every player “searches” for the strategy providing her with the highest payoff, given the time structure and given her available information.
- Strict dominance, instead, will just “rule out” strategies that a rational player would never choose.
- This occurs when a strategy profiles her with a strictly lower payoff than other available strategies *regardless* of her opponents’ behavior.

Introduction

- We apply that “ruling out” strategy iteratively, moving from one player to another, until we cannot rule out any other strategies.
- The result will be the strategy profiles surviving IDSDS (Iterative Deletion of Strictly Dominated Strategies).
- We will apply IDSDS to typical games in economics and social sciences.
- We will finish the chapter talking about two related solution concepts: IDWDS and SDE.

Introduction

Definition: Strictly Dominated Strategy

- Player i finds that a strategy s_i strictly dominates strategy s'_i if and only if

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$$

for every strategy profile $s_{-i} \in S_{-i}$ of player i 's rivals, and where $s_i, s'_i \in S_i$, and $s'_i \neq s_i$.

- Intuition:
 - Strategy s_i that provides a player with a strictly lower payoff than another available strategy, s'_i , regardless of her opponents' strategy.
 - A rational player should, then, never choose s_i , independently on how she believes her rivals behave.
 - We can delete strategy s_i from player i 's strategy set.

Tool 2.1. How to Find Strictly Dominated Strategies

1. Fix your attention on one strategy of the column player, s_2 (i.e., one specific column). Find a strategy s'_1 that yields a strictly lower payoff than some other strategy s_1 for the row player, that is $u_1(s_1, s_2) > u_1(s'_1, s_2)$ for a given strategy of player 2.
2. Repeat step 1, but now fix your attention on a different column. That is, find if the above payoff inequality also holds, $u_1(s_1, s'_2) > u_1(s'_1, s'_2)$, which is now evaluated at a different strategy of the column player, s'_2 .
3. If, after repeating step 1 for all possible strategies of player 2 (all columns), you find that strategy s'_1 yields a strictly lower payoff for player 1 than strategy s_1 , you can claim that strategy s'_1 is strictly dominated by s_1 .

Otherwise, player 1 does not have a strictly dominated strategy.

Analogous argument for player 2, switching subscripts and switching “row” for “column” everywhere.

Example 2.1: Finding Strictly Dominated Strategies

		Firm 2		
		h	m	l
Firm 1	H	2,2	3,1	5,0
	M	1,3	2,2	2.5,2.5
	L	0,5	2.5,2.5	3,3

Matrix 2.1 Strictly Dominated Strategies in the Output game

Example: Finding Strictly Dominated Strategies

		Firm 2		
		<i>h</i>	<i>m</i>	<i>l</i>
Firm 1	H	2,2	3,1	5,0
	M	1,3	2,2	2.5,2.5
	L	0,5	2.5,2.5	3,3

Firm 1 finds row L to be strictly dominated by H since:

- When firm 2 chooses *h*, $2 > 0$.
- When firm 2 chooses *m*, $3 > 2.5$.
- When firm 2 chooses *l*, $5 > 3$.

A similar argument applies to row M, which is strictly dominated by H since:

- When firm 2 chooses *h*, $2 > 0$.
- When firm 2 chooses *m*, $3 > 2$.
- When firm 2 chooses *l*, $5 > 2.5$.

A similar argument applies to firm 2: columns *m* and *l* are both strictly dominated by *h*. Check as a practice.

Tool 2.2. Applying IDSDS

We next seek to delete as many strictly dominated strategies as possible for each player. We use IDSDS, following the next steps:

1. From the definition of rationality, we know that a player would never use strictly dominated strategies,
 - So we can delete them from her original strategy set, S_i ,
 - Obtaining S'_i (which is technically a subset of S_i).
2. We can then proceed by also using common knowledge of rationality.
 - In this context, this entails that every player j can put in her opponent's shoes,
 - Identify all strictly dominated strategies for her opponent, and
 - Delete them from j 's strategy set S_j , obtaining S'_j .

Applying IDSDS

3. Continue again: Player i considers now her rival's reduced strategy set S'_j and finds whether some of his own strategies in $S'_i \times S'_j$ now become strictly dominated.
 - At the end of this step, we obtain a further reduced strategy set S''_i .
 - In some games, this set may coincide with the original set in the third step, $S''_i = S'_i$, if we could not identify more strictly dominated strategies for player i ; while in other games, strategy set S''_i is a strict subset of S'_i .

Applying IDSDS

4. Starting with strategy sets S_i'' and S_j'' , we repeat the above process, seeking to find more strictly dominated strategies for either player.
5. The process continues until no player can identify more strictly dominated strategies to delete.
 - The remaining strategies are referred to as the “strategy profile(s) surviving IDSDS.”

Example 2.2: Applying IDSDS

		Firm 2	
		h	l
Firm 1	H	4,4	0,2
	M	1,4	2,0
	L	0,2	0,0

Matrix 2.2a. When IDSDS yields a unique equilibrium

Starting with firm 1:

Row L is strictly dominated by M . Check!

Example: IDSDS

		Firm 2	
		h	l
Firm 1	H	4,4	0,2
	M	1,4	2,0

Matrix 2.2b. IDSDS yields a unique equilibrium – After one round of IDSDS

Moving to firm 2:

Column l is strictly dominated by h . Check!

Example: IDSDS

		Firm 2	
		h	
Firm 1	H	4,4	
	M	1,4	

Matrix 2.2c. IDSDS yields a unique equilibrium – After two rounds of IDSDS

		Firm 2	
		h	
Firm 1	H	4,4	

Matrix 2.2d. IDSDS yields a unique equilibrium – After three rounds of IDSDS

Then, the unique strategy profile surviving IDSDS is (H, h) , yielding an equilibrium payoff of 4 to each firm.

IDSDS Properties

1. Does the order of deletion matter?

- The order of deletion does not matter when applying IDSDS.
- We would predict different equilibrium results if we started deleting strictly dominated strategies for player 1 or 2.
- See Example 2.3 in the book, where we redo Example 2.2 but starting with firm 2.

2. Deleting more than one strategy at a time

- We can delete all strictly dominated strategies we found for that player, so we do not need to delete on strictly dominated strategy at a time, saving us valuable time.
- We did this in Example 2.3.

3. Multiple equilibrium predictions

- In Examples 2.2 and 2.3 we found unique equilibrium predictions.
- In this case, we say the game is “dominance solvable.”
- That’s not necessarily the case for all games, as we see next.

Example 2.4. Multiple Equilibria

Firm 1 finds H to be strictly dominated by L .

After deleting row H , we obtain the matrix at the bottom of the slide.

		Firm 2		
		h	m	l
Firm 1	H	2,3	1,4	3,2
	M	5,3	2,1	1,2
	L	3,6	4,7	5,4

Matrix 2.3a. When IDSDS yields more than one equilibrium-I

Firm 2 finds l to be strictly dominated by h .

After deleting column l , we obtain the matrix in the next slide

		Firm 2		
		h	m	l
Firm 1	M	5,3	2,1	1,2
	L	3,6	4,7	5,4

Matrix 2.3b. When IDSDS yields more than one equilibrium-II

Example 2.4. Multiple Equilibria

		Firm 2	
		h	m
Firm 1	M	5,3	2,1
	L	3,6	4,7

Matrix 2.3c. When IDSDS yields more than one equilibrium-III

At this point, we find no other strictly dominated strategies for firm 1, or for firm 2. Check!

Then,

$$IDSDS = \{(M, h), (M, m), (L, h), (L, m)\},$$

meaning that IDSDS doesn't provide us a very precise equilibrium prediction.

IDSDS: Prisoner's Dilemma

		Player 2	
		Confess	Not Confess
Player 1	Confess	-4,-4	0,-8
	Not Confess	-8,0	-2,-2

Matrix 2.4a. The Prisoner's Dilemma game

Starting with player 1, he finds that NC is strictly dominated by C since:

- $-8 < -4$ when player 2 confesses, in the left column.
- $-2 < 0$ when player 2 doesn't confess, in the right column.

IDSDS: Prisoner's Dilemma

		Player 2	
		Confess	Not Confess
Player 1	Confess	-4,-4	0,-8
	Not Confess	-8,-4	-5,-5

Matrix 2.4b. The Prisoner's Dilemma game after one round of IDSDS

Player 2 finds NC to be strictly dominated by C since $-8 < -4$.

IDSDS: Prisoner's Dilemma

Deleting the NC column, we are left with just one cell:

		Player 2
		Confess
Player 1	Confess	-4,-4

Matrix 2.4c. The Prisoner's Dilemma game after two rounds of IDSDS

IDSDS then predicts a unique strategy profile, (C, C) , with an equilibrium payoff of -4 to every player. Notice that the outcome is not Pareto optimal, yet no player has *unilateral incentives to deviate*.

Other Prisoner's Dilemma Games

- Similar conflicts between individual and group incentives are common in economics and social sciences:
 - Price wars between firms
 - Tariff wars between countries
 - Use of negative campaigning in politics.
- Some movies with Prisoner's Dilemma game scenes:
 - The dark knight: https://www.youtube.com/watch?v=K4GAQtGtd_0.
 - Murder by numbers, <https://www.youtube.com/watch?v=UAR7WDrL3Ec>.
 - The hunger games: https://www.youtube.com/watch?v=rCWL-pRX_hE.

Prisoner's Dilemma Game, General form

		Player 2	
		Confess	Not Confess
Player 1	Confess	a, a	b, c
	Not Confess	c, b	d, d

Matrix 2.4d. The Prisoner's Dilemma game – General Form

- For players to find Confess to be strictly dominant, we need that:
 - $a > c$, in the left column and
 - $b > d$, in the right column.
- We aren't ranking payoffs a and d , but:
 - If $a < d$ we can say that the strategy profile surviving IDSDS, (C, C) , is not Pareto optimal.
 - Otherwise, the equilibrium of this game according to IDSDS, (C, C) , is also Pareto optimal.

Coordination Games – Battle of the Sexes

		Wife	
		Football, F	Opera, O
Husband	Football, F	10,8	6,6
	Opera, O	4,4	8,10

Matrix 2.5a. The Battle of the Sexes game

- Husband has no strictly dominated strategy. Check.
- Same argument for wife.
- Then, the game as a whole is our equilibrium prediction after applying IDSDS (i.e., four strategy profiles).

Coordination Games, Other examples

- Interaction of two or more firms investing in different technologies:
 - Firms can share files, or parts, only if they both invest in the same technology.
- Friends buying game consoles that run on different operating systems.
- Friends joining social media platforms.
- Choosing the side of the road to drive (driving conventions).
- More generally, games where players benefit from *positive network externalities* when choosing the same strategy as their rivals.

Coordination Games, General form

		Wife	
		Football, F	Opera, O
Husband	Football, F	a_H, a_W	b_H, c_W
	Opera, O	c_H, b_W	d_H, d_W

Matrix 2.5b. Coordination games – General Form

For this to be a coordination game (no strictly dominated strategy for any player), we need payoffs to satisfy:

- $a_i > c_i$
- $d_i > c_i$

for every player $i = \{H, W\}$.

This general form allows for symmetric payoffs when attending the same event, or asymmetric payoffs.

Pareto Coordination Games – Stag Hunt

Pareto Coordination Game – Stag Hunt

		Player 2	
		Stag	Hare
Player 1	Stag	6,6	1,4
	Hare	4,1	2,2

Matrix 2.6a. The stag hunt game

- Player 1 has no strictly dominated strategies.
- Player 2 doesn't have any strictly dominated strategy either.
- Hence, IDSDS predicts the original four strategy profiles.
 - In these cases, we say that "IDSDS has no bite."

General Form and Pareto Coordination Games

The general form of Pareto coordination games

		Player 2	
		Stag	Hare
Player 1	Stag	a, a	c, b
	Hare	b, c	d, d

Matrix 2.6b. Pareto coordination games – General Form

Payoffs need to satisfy $a > b \geq d > c$.

So, when player 2 selects:

- Stag, player 1 is better off choosing Stag as well since $a > b$
- Hare, player 1 is better off choosing Hare as well since $d > c$.

Anti Coordination Game – The Game of Chicken

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1, -1	-8, 10
	Stay	10, -8	-30, -30

Matrix 2.7a. Anticoordination game

- Player 1 has no strictly dominated strategies.
 - When player 2 chooses to *Swerve*, player 1 is better off choosing *Stay* because $10 > -1$, but if player 2 chooses to *Stay*, player 1 is better off choosing *Swerve* because $-8 > -30$.
- Player 2 doesn't have any strictly dominated strategy either.
- Hence, IDSDS predicts the original four strategy profiles.
 - In these cases, we say that "IDSDS has no bite."

Anti Coordination Game – The Game of Chicken

		Player 2	
		Swerve	Stay
Player 1	Swerve	a, a	c, b
	Stay	b, c	d, d

Matrix 2.7b. Anticoordination game – General form

For players to exhibit incentives to anticoordinate, payoffs need to satisfy $b > a > c > d$.

So, when player 2 selects:

- Swerve, player 1 is better off choosing Stay since $b > a$
- Stay, player 1 is better off choosing Swerve since $c > d$.

Anti Coordination Game – The Game of Chicken

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1, -1	-8, 10
	Stay	10, -8	-30, -30

Matrix 2.7a. Anticoordination game

Movies with Anti-coordination games scenes

- Rebel without a Cause: <https://www.youtube.com/watch?v=BGtEp7zFdrc>
- Footloose: <https://www.youtube.com/watch?v=ZL57muBck0w>.
- Stand by Me: <https://www.youtube.com/watch?v=6L-lbkWHsOE>.
- The hunt for the Red October: <https://www.youtube.com/watch?v=mh599VtMB1c>.
- A beautiful mind: <https://www.youtube.com/watch?v=LJS7lgvk6ZM>

Symmetric Games

- **Definition:** A two-player game is symmetric if both players' strategy sets coincide, $S_A = S_B$, and payoffs are unaffected by the identity of the player choosing each strategy, that is:

$$u_A(s_A, s_B) = u_B(s_A, s_B)$$

for every strategy profile (s_A, s_B) .

- This property is also known as that the payoff function satisfies *anonymity*:
 - Starting from a setting where player A chooses s_A and player B chooses s_B , ...
 - if we were to switch the identities of players A and B , so that player A becomes B and player B becomes A ,...
 - every player's payoff would be unaffected.

Symmetric Games – Visual interpretation

Visually, the matrix of a symmetric game must satisfy the following properties:

1. Same number of rows and columns, so both players have the same number of available strategies.
2. Same action labels in rows and columns, so player A and B face the same strategy sets.
3. Payoffs along the main diagonal must coincide for every player (i.e. both players earn the same payoff when they choose the same strategy).
 1. Formally, this property entails that, if both players choose the same strategy, $s_A = s_B = s$, their payoffs must also coincide, $u_A(s, s) = u_B(s, s)$.
 2. Otherwise, an identity “switch” (formally known as a permutation) would affect players’ payoffs.
4. Cells above the main diagonal must be mirror images of those below the main diagonal.
 1. This implies that when $s_A \neq s_B$, we can satisfy the condition $u_A(s_A, s_B) = u_B(s_A, s_B)$.
 2. Otherwise, an identity “switch” would affect players’ payoffs.

Examples: Prisoner’s Dilemma and Game of Chicken are symmetric. Battle of the Sexes is not.

Symmetric Games - Example

		Player 2	
		Confess	Not Confess
Player 1	Confess	-4,-4	0,-8
	Not Confess	-8,0	-2,-2

Matrix 2.4a. The Prisoner's Dilemma game

- This game satisfies the above 4 properties.
 - Same number of rows and columns.
 - Same action labels.
 - Payoffs along the main diagonal coincide across players.
 - Payoffs away from the main diagonal are mirror images of each other. (next slide explains this point).

Symmetric Games - Example

		Player 2	
		Confess	Not Confess
Player 1	Confess	-4,-4	0,-8
	Not Confess	-8,0	-2,-2

Matrix 2.4a. The Prisoner's Dilemma game

- Consider the strategy profile $(s_1, s_2) = (C, NC)$, where player 1 confesses but player 2 does not.
- In this setting, player 1's payoff is $u_1(C, NC) = 0$, while player 2's is $u_2(C, NC) = -8$.
- If we switch their identities, so that player 1 becomes player 2 (and does not confess), and player 2 becomes player 1 (and confesses), their payoffs are $u_2(NC, C) = 0$ for player 2 (who was originally player 1), and $u_1(NC, C) = -8$ for player 1 (who was originally player 2), thus being unaffected by the identity switch.
- A similar argument applies to other strategy profiles in this game.

Asymmetric Games

- **Definition:** A two-player game is asymmetric if:
 - Players face different strategy sets, $S_A \neq S_B$, or
 - Despite facing the same strategy sets, $S_A = S_B$, their payoffs satisfy

$$u_A(s_A, s_B) \neq u_B(s_A, s_B)$$

for at least one strategy profile (s_A, s_B) .

- In other words, a game is asymmetric if it violates at least one of the 4 properties required for a game to be symmetric.

Asymmetric Games - Example

		Wife	
		Football, F	Opera, O
Husband	Football, F	10,8	6,6
	Opera, O	4,4	8,10

Matrix 2.5a. The Battle of the Sexes game

- Consider the strategy profile $(s_H, s_W) = (F, F)$, where both husband and wife go to the football game.
- In this context, the husband's payoff is $u_H(F, F) = 10$, while his wife's is $u_W(F, F) = 8$.
- If we switch their identities, the wife's payoff becomes $u_w(F, F) = 10$ and $u_H(F, F) = 8$, implying that their payoffs are affected.
- Graphically, this could be seen in Matrix 2.5a:
 - While satisfying Properties #1 and #2 for the game to be symmetric.
 - It violates Property #3 (players' payoffs don't coincide along the main diagonal), and
 - It also violates Property #4 (players' payoffs above and below the main diagonal are not mirror images).

Randomization to bring IDSDS Further

- Sometimes IDSDS has little or no “bite” in the game.
 - Consider Battle of the Sexes game, as an example of Coordination game; or the Game of Chicken, as an examples of anticoordination game).
 - You can also consider the following 3x3 matrix.

		Firm 2		
		<i>h</i>	<i>m</i>	<i>l</i>
Firm 1	H	0,10	4,6	4,6
	M	4,6	0,10	6,4
	L	10,0	6,4	4,6

Matrix 2.8. Applying IDSDS and allowing for randomizations- I

- Firm 1 has no strictly dominated strategies we can delete:
 - While row *H* yields strictly lower payoffs than *L* when firm 2 chooses *h* or *m* ($0 < 10$ and $4 < 6$, respectively), it yields the same payoff when firm 2 chooses *l* (4).
- A similar argument applies to Firm 2 which does not have strictly dominated strategy either.
- What can we do?

Randomization to bring IDSDS Further

Don't despair! When we allow players to randomize between two or more strategies, IDSDS might still have some "bite."

		Firm 2			
		h	m	l	
Firm 1	Prob. $(1 - q) \rightarrow$	H	0,10	4,6	4,6
		M	4,6	0,10	6,4
	Prob. $q \rightarrow$	L	10,0	6,4	4,6

Matrix 2.8. Applying IDSDS and allowing for randomizations- I

Randomization to bring IDSDS Further

		Firm 2			
		h	m	l	
Firm 1	Prob. $(1 - q) \rightarrow$	H	0,10	4,6	4,6
		M	4,6	0,10	6,4
	Prob. $q \rightarrow$	L	10,0	6,4	4,6

Matrix 2.8. Applying IDSDS and allowing for randomizations- I

- However, strategy H for firm 1 seems to produce a lower payoff than a combination of rows M and L .
- Specifically, M and L yield weakly higher payoffs than H does when firm 2 picks columns h or l , but yields extreme payoffs when firm 2 chooses column m (which are clearly above or below that of H).
- This indicates that, if we could create a linear combination between the payoffs in row M and L , player 1 would receive an expected utility that lies strictly above that of strategy H , ultimately helping us to delete this top row from the matrix.

Randomization to bring IDSDS Further

		Firm 2		
		h	m	l
Firm 1	Prob. $(1 - q) \rightarrow$	H	0,10	4,6
		M	4,6	0,10
	Prob. $q \rightarrow$	L	10,0	6,4

Matrix 2.8. Applying IDSDS and allowing for randomizations- I

- Let us first compute the expected utility from assigning:
 - A probability weight $q \in (0,1)$ to strategy L , and
 - The remaining probability weight $1 - q$ to strategy M .
- That is, we create the mixed strategy $\sigma = qL + (1 - q)M$.
- Intuitively, a randomization could be flipping a coin, $q = 1/2$, but that's too specific. At this point, we don't know which probability weight we should assign to L and M to guarantee that their expected payoff exceeds the certain payoff from playing H .

Randomization to bring IDSDS Further

		Firm 2		
		h	m	l
Firm 1	Prob. $(1 - q) \rightarrow$	H	0,10	4,6
		M	4,6	0,10
	Prob. $q \rightarrow$	L	10,0	6,4

Matrix 2.8. Applying IDSDS and allowing for randomizations- I

- Mixed strategy $\sigma = qL + (1 - q)M$ yields the following expected utilities:

$$EU(\sigma|h) = q10 + 4(1 - q) = 4 + 6q$$

$$EU(\sigma|h) = q6 + 0(1 - q) = 6q$$

$$EU(\sigma|h) = q4 + 6(1 - q) = 6 - 2q$$

Randomization

Inserting these expected payoffs into Matrix 2.8 yields

		Firm 2		
		h	m	l
Firm 1	H	0	4	4
	$qL + (1 - q)M$	$4 + 6q$	$6q$	$6 - 2q$

Matrix 2.9. Applying IDSDS and allowing for randomization –II

Randomization

We can then say that, for the randomization $\sigma = qL + (1 - q)M$ to strictly dominate strategy H , we need that Firm 1's expected payoff need to satisfy:

$$4 + 6q > 0, \text{ or } q > -\frac{2}{3} \text{ when firm 2 chooses } h$$

$$6q > 4, \text{ or } q > \frac{2}{3} \text{ when firm 2 chooses } m, \text{ and}$$

$$6 - 2q > 4, \text{ or } q < 1 \text{ when firm 2 chooses } l$$

- The first and last condition on probability q hold because q must satisfy $q \in (0,1)$ by assumption.
- The second condition, $q > \frac{2}{3}$ restricts the range of probability weights we can use in randomization $qL + (1 - q)M$.
- Intuitively, player 1 must assign a sufficiently high weight to strategy L .

Randomization

- Therefore, any randomization with $q > \frac{2}{3}$, helps us show that strategy H is strictly dominated, so we can delete it from Matrix 2.8 to obtain Matrix 2.10.

		Firm 2		
		h	m	l
Firm 1	M	4,6	0,10	6,4
	L	10,0	6,4	4,6

Matrix 2.10. Applying IDSDS and allowing for randomizations –III

Randomization

We can now move to Firm 2 to note that h is strictly dominated by m since h yields a strictly lower payoff than m does:

- When firm 1 chooses M in top row (where $10 > 6$), and
- When firm 1 chooses L in the bottom row (where $4 > 0$).

Randomization

We can now move to Firm 2 to note that h is strictly dominated by m since h yields a strictly lower payoff than m does:

- When firm 1 chooses M in top row (where $10 > 6$), and
- When firm 1 chooses L in the bottom row (where $4 > 0$).

After deleting this strictly dominated strategy for firm 1, we are left with:

		Firm 2	
		m	l
Firm 1	M	0,10	6,4
	L	6,4	4,6

Matrix 2.11. Applying IDSDS and allowing for randomization – IV

Randomization

- We can now move again to firm 1, noticing that it no longer has strictly dominated strategies.
 - In particular, M yields a strictly lower payoff than L when firm 2 chooses m , *but* a strictly higher payoff than L when firm 2 chooses l .
- If we move to firm 2, we cannot find strictly dominated strategies for this player either.
 - When firm 1 chooses row M , firm 2 is better off selecting m ; *but* when firm 1 chooses row L , firm 2 is better off selecting l .

		Firm 2	
		m	l
Firm 1	M	0,10	6,4
	L	6,4	4,6

Randomization

- In summary, the application of IDSDS did not have a bite when we assumed players could not randomize.
 - In that context, our equilibrium prediction would have been the game as a whole, yielding $3 \times 3 = 9$ different outcomes.
- Allowing for randomizations, however, helped us discard one strategy for each player as being strictly dominated, leaving us with a 2×2 matrix (4 different equilibrium outcomes), that is:

$$IDSDS = \{(M, m), (M, l), (L, m), (L, l)\}$$

What if IDSDS has no bite?

In some games,

- Players may have no strictly dominated strategies that we can delete using IDSDS.
- Ok, we can try by allowing for randomizations.
- What if, still, we cannot delete any strictly dominated strategies?
- Examples include the Battle of the Sexes game and the Game of Chicken, among others.

What to do?

- Apply Nash Equilibrium! (Chapters 3-5)

		Criminal	
		<i>Street A</i>	<i>Street B</i>
Police	<i>Street A</i>	10,0	-1,6
	<i>Street B</i>	0,8	7,-1

Evaluating IDSDS as a concept

1. Existence? Yes.
 - When we apply IDSDS to any game, we find that at least one equilibrium exists.
2. Uniqueness? No.
 - IDSDS does not necessarily provide a unique equilibrium outcome for all games since more than one cell may survive IDSDS.
3. Robust to small payoff perturbations? Yes.
 - If we change the payoff of one of the players by a small amount, IDSDS still yields the same equilibrium outcome(s).
4. Socially optimal? No.
 - The application of IDSDS does not necessarily yield socially optimal outcomes.
 - Prisoner's Dilemma game, where IDSDS provides us with a unique equilibrium prediction, (Confess, Confess), which does not coincide with the strategy profile that maximizes the sum of players' payoffs, (Not Confess, Not Confess).

Weakly Dominated Strategies

- **Definition:** Player i finds that the strategy s_i weakly dominates another strategy s'_i if:
 1. $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ for every strategy profile s_{-i} player i 's rivals, and
 2. $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for *at least one* strategy profile s_{-i} .

Therefore, choosing s_i provides player i with:

- the same or higher payoff than s'_i against all her rivals' strategies, but...
- a strictly higher payoff than s'_i for at least one of her rivals' strategies.

Technically, requirement #2 in this definition avoids a “total tie” between player i 's payoffs in s_i and s'_i .

Weakly Dominated Strategies

- **Definition:** Player i finds that the strategy s_i weakly dominates another strategy s'_i if:
 1. $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ for every strategy profile s_{-i} player i 's rivals, and
 2. $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for *at least one* strategy profile s_{-i} .

Graphical interpretation in a two-player game (potentially with more than two rows and columns):

- Player 1 finds that two of her rows yield the same payoff against every column of player 2, except for one column.
- Similar argument applies for player 2.

Weakly Dominated Strategies – Example 2.6

		Firm 2		
		<i>h</i>	<i>m</i>	<i>l</i>
Firm 1	H	2,2	3,2	5,0
	M	2,3	3,2	2.5,2.5
	L	0,5	2.5,2.5	3,3

Matrix 2.12. Finding weakly dominated strategies

- Firm 1 finds that row H weakly dominates M because:
 - When firm 2 chooses *h*, $2=2$.
 - When firm 2 chooses *m*, $3=3$.
 - When firm 2 chooses *l*, $5>2.5$.
- Recall if, firm 1's payoffs coincided across all columns (i.e., $5=5$ in column *l*), then we couldn't claim that H weakly dominates M because Requirement #2 in the definition wouldn't be satisfied.

Does weakly dominated mean strictly dominated?

$$s_i \text{ is strictly dominated} \Rightarrow s_i \text{ is weakly dominated} \\ \nLeftarrow$$

- In Matrix 2.12, strategy L is strictly dominated by H , since L yields an unambiguously lower payoff for firm 1 than H does, regardless of the column that firm 1 chooses.
- This means that the requirement 2 of the weak dominance definition holds for all firm 2's strategies, meaning that L is also weakly dominated by H .
- However, strategy M is weakly dominated by H but M is not strictly dominated by H .

Deletion order matters in IDWDS

IDWDS may yield different equilibrium predictions depending on which player we start with

- Starting with Firm 1

		Firm 2		
		h	m	l
Firm 1	H	2,2	3,2	5,0
	L	0,5	2.5,2.5	3,3

Matrix 2.13a. Applying IDWDS starting with firm 1- First step

Deletion order matters in IDWDS

IDWDS may yield different equilibrium predictions depending on which player we start with

- Starting with Firm 1
- It is straightforward to see that firm 1 finds that H strictly dominates L , which helps us delete row L , and obtain Matrix 2.13b.

		Firm 2		
		h	m	l
Firm 1	H	2,2	3,2	5,0

Matrix 2.13b. Applying IDWDS starting with firm 1- Second step

Deletion order matters in IDWDS

IDWDS may yield different equilibrium predictions depending on which player we start with

- Starting with Firm 1
- Turning now to firm 2, it finds that h strictly dominates l , so we can delete the right-hand column and obtain Matrix 2.13c.

		Firm 2	
		h	m
Firm 1	H	2,2	3,2

Matrix 2.13c. Applying IDWDS starting with firm 1- Third step

- At this point, we cannot delete any further strategies as being weakly or strictly dominates for either firm, yielding an equilibrium outcome:

$$IDWDS = \{(H, h), (H, m)\}$$

Deletion order matters in IDWDS

IDWDS may yield different equilibrium predictions depending on which player we start with

- Starting with Firm 2
- We find that column h weakly dominates m since it yields the same payoff as m does when firm 1 chooses H , but generates a strictly higher payoff when firm 1 selects M or L . After deleting column m , we obtain Matrix 2.14a:

		Firm 2	
		h	l
Firm 1	H	2,2	5,0
	M	2,3	2.5,2.5
	L	0,5	3,3

Matrix 2.14a. Applying IDWDS starting with firm 2- First step

Deletion order matters in IDWDS

IDWDS may yield different equilibrium predictions depending on which player we start with

- Starting with Firm 2
- Still analyzing firm 2, we find that h strictly dominates l , which leaves us with the reduced Matrix 2.14b:

		Firm 2	
		h	
Firm 1	H	2,2	
	M	2,3	
	L	0,5	

Matrix 2.14b. Applying IDWDS starting with firm 2- Second step

Deletion order matters in IDWDS

IDWDS may yield different equilibrium predictions depending on which player we start with

- Starting with Firm 2
- We can now turn to firm 1, which finds that H strictly dominates L , yielding Matrix 2.14c:

		Firm 2	
		h	
Firm 1	H	2,2	
	M	2,3	

Matrix 2.14c. Applying IDWDS starting with firm 2- Third step

- At this point, we cannot delete any other strategies for firm 1, as none are strictly or weakly dominated, which yields the equilibrium prediction:

$$IDWDS = \{(H, h), (M, h)\}$$

IDSDS vs. IDWDS

We say that the set of strategy profiles surviving IDWDS is a subset of those surviving IDS

$$s \text{ survives IDWDS} \Rightarrow s \text{ survives IDS} \\ \nLeftarrow$$

Example

		Firm 2	
		h	m
Firm 1	H	2,2	3,2
	M	2,3	3,2

Matrix 2.15. Applying IDS in Matrix 2.11

IDSDS vs. IDWDS

- Firm 1 cannot delete any strictly dominated strategies from Matrix 2.15, and neither can Firm 2, yielding the four equilibrium outcomes:

$$IDSDS = \{(H, h), (H, m), (M, h), (M, m)\}$$

- Therefore, the application of IDWDS has more bite than IDSDS, producing more precise equilibrium predictions.
- However, IDWDS depends on the order of deletion (e.g. starting with player 1 or 2), thus limiting its applicability.

Strictly Dominant Strategy

Definition: Strictly Dominant Strategy

- Player i finds that strategy s_i is strictly dominant if

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$$

for every strategy $s'_i \neq s_i$, and

for every strategy profile s_{-i} of player i 's rivals.

- Intuitively, all her available strategies yield a strictly lower payoff than s_i regardless of her rivals' behavior.
- This means that we can delete all strategies $s'_i \neq s_i$ as being strictly dominated for player i , leaving her with a single (undominated) strategy.
- In other words, she must be choosing s_i .

Strictly Dominant Strategy

Definition: Strictly Dominant Strategy

- Player i finds that strategy s_i is strictly dominant if

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$$

for every strategy $s'_i \neq s_i$, and

for every strategy profile s_{-i} of player i 's rivals.

- *Examples:*
 - Confess in Prisoner's Dilemma game is strictly dominant for both players.
 - No player has no strictly dominant strategies in BoS or Chicken games.
 - In matrix 2.1:
 - Firm 1 finds H to strictly dominate m and l, so we can claim that H is strictly dominant for firm 2.
 - Same argument with firm 2, which finds h to be strictly dominant.

Strictly Dominant Strategy and Strictly Dominant Equilibrium (SDE)

Definition: Strictly Dominant Equilibrium (SDE)

- A strategy profile $s^{SD} = (s_i^{SD}, s_{-i}^{SD})$ is a strictly dominant equilibrium (SDE) if every player i finds her strategy, s_i^{SD} , to be strictly dominant.
- An SDE is the only strategy surviving IDSDS after just one round of deleting strictly dominated strategies.
 - Example: Matrix 2.1.
 - H (h) is strictly dominant for firm 1 (2, respectively), implying that (H,h) is an SDE.
 - Firm 1 (2) finds that M and L (m and l) are strictly dominated by H (h), implying that firm 1 (2) can delete M and L (m and l) from its rows (columns), leaving us with (H,h) as the unique strategy profile surviving IDSDS.
 - [As a curiosity, note that we only needed to apply Rationality, not Common Knowledge of Rationality]

Example SDEs

- Examples:
 - Prisoner's Dilemma. Confess is strictly dominant strategy for every player, implying that (C,C) is SDE.
 - Battle of the Sexes. No player has strictly dominant strategy, implying that no SDE exists.
 - Game of Chicken. No player has strictly dominant strategy, implying that no SDE exists.

		<i>Firm 2</i>	
		<i>h</i>	<i>l</i>
<i>Firm 1</i>	<i>H</i>	4, 4	0, 2
	<i>M</i>	1, 4	2, 0
	<i>L</i>	0, 2	0, 0

Matrix 2.2a. IDSDS yields a unique equilibrium.

Example SDEs

- More examples:
 - Consider Matrix 2.2 again.
 - Earlier in this chapter, we showed that a unique strategy profile survives IDSDS.
 - What about SDE? There is none, because no player has a strictly dominant strategy.
 - While we found strictly dominated strategies for every player, we cannot find strictly dominant strategies for any player.
 - Same argument applies to Examples 2.3 and 2.4.

		<i>Firm 2</i>	
		<i>h</i>	<i>l</i>
<i>Firm 1</i>	<i>H</i>	4, 4	0, 2
	<i>M</i>	1, 4	2, 0
	<i>L</i>	0, 2	0, 0

Matrix 2.2a. IDSDS yields a unique equilibrium.

Example SDEs

- Therefore, we can conclude that:

$$(s_i, s_{-i}) \text{ is an SDE} \Rightarrow (s_i, s_{-i}) \text{ survives IDSDS} \\ \nLeftarrow$$

- For examples of the first line of implication, consider the Prisoner's Dilemma game.
- For examples of the second line (of not implication), consider (F,F) or (O,O) in the Battle of Sexes game.

Evaluating SDE as a solution concept

1. Existence? No.
 - When we seek to find strictly dominant strategies for each player, in our search of a SDE, we may find games where one or more players do not have a strictly dominant strategy, implying that an SDE does not exist.
2. Uniqueness? Yes
 - While a SDE may not exist in some games, when we find one, it must be unique.
3. Robustness to small payoff perturbation? Yes.
 - If we change the payoff of one of the players by a small amount ($\varepsilon \rightarrow 0$), SDE still yields the same equilibrium prediction.
4. Socially optimal? No.
 - SDE of a game does not need to be socially optimal.
 - Example: Prisoner's Dilemma game, the SDE is (Confess, Confess), which does not coincide with the strategy profile that maximizes the sum of players' payoffs.

Evaluating SDE as a solution concept

	IDSDS	IDWDS	SDE
Existence	Yes	Yes	No
Uniqueness	No	No	Yes
Robustness of payoff changes	Yes	Yes	Yes
Pareto optimal	No	No	No

- In summary:
 - IDSDS and IDWDS exhibit similar properties, although IDWDS has “more bite.”
 - Relative to IDSDS and IDWDS, SDE is so much more demanding that it may not exist. But if it exists, it is unique.

Chapter 3: Nash Equilibrium

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

Nash Equilibrium builds upon the notion that every player finds the strategy that maximizes her payoff against each of his rivals' strategies, which we refer as her “best response”.

- Discrete Setting:

If player i has only two available strategies, “maximizing” his payoff means choosing the strategy that yields the highest payoff, taking her rivals' strategies as given.

- Continuous Setting:

Player i chooses her strategy from a continuum of available strategies, her best response is found by maximizing this player's payoff, literally:

Solving a utility or profit maximization problem, where we will take the strategies of her rivals as given (as if they were a parameter).

Best Response

Definition: Player i regards strategy s_i as a best response to strategy profile s_{-i} , if

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

for every available strategy $s'_i \neq s_i$.

- Intuitively, in a game with two players, i and j , strategy s_i is player i 's best response to player j 's strategy, s_j , if s_i yields a weakly higher payoff than any other strategy s'_i against s_j .
- In other words, when player j chooses s_j , player i maximizes her payoff by responding with s_i than with any other available strategies.

Best Response

Definition: Player i regards strategy s_i as a best response to strategy profile s_{-i} , if

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

for every available strategy $s'_i \neq s_i$.

- If player i finds that the same strategy s_i is a best response against *every strategy profile* of her rivals, all s_{-i} , then s_i must also be strictly dominant.
- Therefore:
 - A strictly dominant strategy is a BR, but...
 - not every BR is a dominant strategy.

Tool 3.1. How to find best responses (BR) in matrix games

1. Focus on the row player by fixing your attention on one strategy of the column player (i.e., one specific column).
 - Cover with your hand (or with a piece of paper) all columns that you are not considering.
 - Find the highest payoff for the row player by comparing the first component of every pair.
 - For future reference, underline this payoff. This is the row player's best response payoff to the column that you considered from the column player.
2. Repeat step 1, but now fix your attention on a different column.
3. For the column player, the method is analogous, but now direct your attention on one strategy of the row player (i.e., one specific row). Cover with your hand all the rows that you are not considering, and compare the payoffs of the column player (i.e., second component of every pair).

Example BR with discrete strategy spaces

		Player 2	
		l	r
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

Example BR with discrete strategy spaces

We can start finding Player 1's best responses.

		Player 2	
		l	r
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

$$BR_1(l) = U ; BR_1(r) = D, \text{ or simply,}$$

$$BR_1(s_2) = \begin{cases} U & \text{if } s_2 = l, \text{ and} \\ D & \text{if } s_2 = r \end{cases}$$

Example BR with discrete strategy spaces

We can start finding Player 1's best responses.

		Player 2	
		l	r
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

		Player 2	
		l	r
Player 1	U	<u>5</u> ,3	2,1
	D	3,6	<u>4</u> ,7

Matrix 3.1b. Finding best responses in a 2x2 matrix – Player 1

Example: BR with discrete strategy spaces

We can start finding Player 2's best responses.

		Player 2	
		l	r
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

$$BR_2(U) = l ; BR_2(D) = r, \text{ or simply,}$$

$$BR_2(s_1) = \begin{cases} l & \text{if } s_1 = U, \text{ and} \\ r & \text{if } s_1 = D \end{cases}$$

Example BR with discrete strategy spaces

We can start finding Player 2's best responses.

		Player 2	
		<i>l</i>	<i>r</i>
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

		Player 2	
		<i>l</i>	<i>r</i>
Player 1	U	5, <u>3</u>	2,1
	D	3,6	4, <u>7</u>

Matrix 3.1c. Finding best responses in a 2x2 matrix – Player 2

Finding BRs with Continuous Strategy Spaces

If player i faces a payoff function $u_i(s_i, s_{-i})$,

- We only need to fix her rival's strategies at a generic s_{-i} , and
- Differentiate with respect to s_i to find the value of the strategy s_i that maximizes player i 's payoff.

Example 3.1: BR under Cournot Quantity Competition

- Consider an industry with two firms, 1 and 2,
- Every firm simultaneously and independently choosing its output level.
- Every firm i faces a symmetric cost function $C(q_i) = cq_i$,
 - where $c > 0$ denotes the marginal cost of additional units of output,
- Firms face inverse demand function $p(Q) = 1 - Q$,
 - where $Q = q_1 + q_2$ represents aggregate output.

Example 3.1: BR under Cournot Quantity Competition

To find BR for firm 1,

$$\max_{q_1 \geq 0} (1 - q_1 - q_2) q_1 - c q_1$$

Differentiating with respect to q_1 , we find:

$$1 - 2q_1 - q_2 = c$$

Solving for q_1 :

$$q_1 = \frac{1 - c}{2} - \frac{1}{2} q_2$$

Firm 1's best response function




Figure 3.1. Firm 1's best response function.

Firm 1's best response function:

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 & \text{if } q_2 < 1-c \\ 0 & \text{otherwise} \end{cases}$$

Because firm 1 responds with a lower output q_1 when its rival increases its output, q_2 , we say that firms' output decisions in this context are "strategic substitutes."

A symmetric best response function applies for firm 2.

Deleting Strategies that are Never a Best Response (NBR)

- Before we use BRs to find the Nash equilibrium of a game...
- We can apply NBR in a similar fashion as in IDSDS.
- Recall that, in IDSDS, we deleted a player's strictly dominated strategies as never being used.
- Now we use a similar approach:
 - If a player never uses a strategy as a BR, we can label it to be "Never a best response" (NBR), and...
 - Delete it from her strategy set.

Deleting Strategies that are Never a Best Response (NBR)

Definition. Never a Best Response (NBR):

Strategy s_i is never a best response if:

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}) \text{ for every } s'_i \neq s_i.$$

does not hold for any strategy profile of her rivals, s_{-i} .

Alternatively, if a strategy s_i is NBR, there are

no beliefs that player i can sustain about how her opponents behave
that would lead her to use strategy s_i as a best response.

Deleting Strategies that are Never a Best Response (NBR)

- If player i finds strategy s_i is strictly dominated by s'_i , then:
 - s_i yields a strictly lower payoff than s'_i
 - *regardless* of the strategy profile her rivals choose.
- As a consequence, s_i cannot be a BR against any strategy profile of i 's rivals.
 - That is, s_i is NBR.
- In summary,

s_i strictly dominated $\Rightarrow s_i$ is NBR.

Rationalizability

- We can go through an iterative process –analogous to IDSDS– but...
 - Identifying strategies that are NBR for either player,
 - rather than strategies that are strictly dominated.
- This iterative process is known as “rationalizability,”
 - Because it finds strategies that player i can rationalize (explain) to be BRs to at least one of her opponents’ strategies.

Tool 3.2. Applying Rationalizability

1. Starting with player i , delete every strategy that is NBR, obtaining the reduced strategy set $S'_i \subset S_i$.
2. Using common knowledge of rationality, we continue the above reasoning, arguing that player $j \neq i$ can anticipate player i 's best responses and, as a consequence, the strategies that are NBR for player i , deleting them from S_i .

Given this reduced strategy space $S'_i \subset S_i$, player j can now examine her own best responses to player i , seeking to identify if one or more are never used, and further restricting her strategy space to $S'_j \subset S_j$.

Applying Rationalizability

3. We obtain the Cartesian product $S'_i \times S'_j$, representing the remaining rationalizable strategies after deleting NBRs for two steps. Player i then finds if some of her strategies in this reduced game are NBR, deleting them from her strategy set S'_i , obtaining $S''_i \subset S'_i$.

4...

k. The process continues until we cannot find more strategies that are NBR for either player.

Applying Rationalizability

- The strategy profile (or set of strategy profiles) that survives this iterative process are referred as “rationalizable” strategy profiles because every player can sustain beliefs about her rival’s behavior.
- In games with two players:
 - Both solution concepts yield identical equilibrium results, that is, IDSDS=Rationalizability.
- In games with three or more players:
 - Equilibrium outcomes do not necessarily coincide.
 - Rationalizability produces more precise equilibrium outcomes than IDSDS.
 - [Technical, see Fudenberg and Tirole, 1995]
- That is, for a given strategy profile s ,

s is rationalizable $\implies s$ survives IDSDS

$\nLeftarrow$

Example 3.2: Rationalizability and IDSDS

		Firm 2	
		h	l
Firm 1	H	4,4	0,2
	M	1,4	2,0
	L	0,2	0,0

Matrix 3.2a. Applying Rationalizability – First Step

Starting from firm 1, row L is NBR:

- When firm 2 chooses h , firm 1's BR is H
- When firm 2 chooses l , firm 1's BR is M

In other words, firm 1 does not have incentives to respond with L regardless of the beliefs it sustains on firm 2's behavior (i.e., the beliefs on firm 2 chooses column h or l).

Example 3.2: Rationalizability and IDSDS

		Firm 2	
		h	l
Firm 1	H	4,4	0,2
	M	1,4	2,0

Matrix 3.2b. Applying Rationalizability – Second Step

Moving to firm 2, we find that:

- When firm 1 chooses H, firm 2's BR is h .
- When firm 1 chooses L, firm 2's BR is h .

Therefore, column l is NBR for firm 2.

Example 3.2: Rationalizability and IDSDS

		Firm 2	
		h	
Firm 1	H	4,4	
	M	1,4	

Matrix 3.2c. Applying Rationalizability –Third Step

Moving back to firm 1, we find that M is its BR to firm 2 choosing l.

Therefore, M is NBR in this step of applying rationalizability.

	h
H	4,4

In conclusion, (H,h) survives rationalizability.

As expected, this coincides with the unique strategy profile surviving IDSDS.

Evaluating Rationalizability as a solution concept

Spoiler: It exhibits the same properties as IDSDS.

1. Existence? Yes
 - Rationalizability satisfies existence, meaning that at least one strategy profile in every game must be rationalizable.
2. Uniqueness? No.
 - One or more strategy profiles may survive rationalizability
 - Does not yield a unique equilibrium prediction in all games with three or more players
3. Robust to small payoff perturbations? Yes.
 - Rationalizability is robust to small perturbations because rationalizability does not change equilibrium outcomes if we alter the payoff of one of the players by a small amount.
4. Socially Optimal? No.
 - Rationalizability does not necessarily yield social optimal outcomes, as illustrated by the Prisoner's Dilemma game, where the only strategy profile surviving rationalizability is (Confess, Confess), which is not socially optimal.

Evaluating Rationalizability as a solution concept

	IDSDS	IDWDS	SDE	Rationalizability
Existence	Yes	Yes	No	Yes
Uniqueness	No	No	Yes	No
Robustness of payoff changes	Yes	Yes	Yes	Yes
Pareto optimal	No	No	No	No

Application of Rationalizability: Finding NBRs in the Beauty Contest

- Consider that your instructor of the Game Theory course shows up in the classroom proposing to play the following “Guess the Average” game (also known as “Beauty Contest”).
- She distributes small pieces of paper and asks each student to write a number (an integer) between 0 and 100.
- There are $N > 2$ students and each of them simultaneously and independently writes her number, taking into account the following rules of the game:
 - The instructor will collect all numbers, write them on the board,
 - Compute the average, and then multiply that average by $\frac{1}{2}$ (half of the average).
 - The instructor, then, declares a winner the student who submitted a number closest to half the average.

Finding NBRs in the Beauty Contest Contd.

Which strategies survive rationalizability?

- In the first step, we can try to eliminate strategies that are NBR for any student.
- Starting with the highest number, $s_i = 100$, we can see that, even if all other students submit $s_j = 100$ for all $j \neq i$, and the number of students is large enough to yield an average of $\bar{s} = 100$, we would obtain that half of the average approaches 50.
- Therefore, student i 's best response would be 50, i.e., $BR_i(s_j) = 50$ where $s_j = 100$ for all $j \neq i$, entailing that i cannot rationalize $s_i = 100$ being a best response, regardless of which numbers she believes her rivals submit.
- In summary, $s_i = 100$ is NBR, helping us restrict her strategy space to $S'_i = \{0, 99\}$, in the first step of rationalizability.

Finding NBRs in the Beauty Contest Contd.

Which strategies survive rationalizability?

- In the second step, even if all other students submit the highest number (in the reduced strategy space S'_i), $s_j = 99$ for all $j \neq i$, the average would be $\bar{s} = 99$, entailing that half of the average is 49.5. As a result, every student i finds that submitting $s_i = 99$ is NBR either, so we can further restrict her strategy space to $S''_i = \{0, 98\}$ in the second step of rationalizability.
- ...
- After 97 more rounds, we are left with a restricted strategy space of $\{0, 1\}$ which cannot be further reduced. Indeed, submitting $s_i = 0$ ($s_i = 1$) is a best response if student i believes that her classmates will all submit $s_j = 0$ ($s_j = 1$), so we cannot rule out strategies as being NBR. Therefore, all symmetric strategy profiles, where every student i submits the same number (all submit $s_i = 0$ or all submit $s_i = 1$) survive rationalizability; and so do all asymmetric strategy profiles, where at least one student submits $s_i = 0$ and all other $j \neq i$ students submit $s_j = 1$.

Experimental Tests

- This game has been tested in experimental labs quite often.
 - See, for instance, Nagel (1995) and Stahl and Wilson (1995), which initiated these studies, and for a literature review, see Camerer (2003) and Crawford et al. (2011).
- If you participate in one of these experiments and believe that all other participants randomly submit a number uniformly distributed between 0 and 100, often referred by the literature as “Level-0” players, then you can participate that the average will be around 50 (assuming a sufficiently large number of participants), entailing that half of the average would be 25.
- Therefore, if you think that your rivals are “Level-0” players, you submit $s_i = 25$.
- However, if you think that your rivals must have gone through that same thought process, and they are submitting $s_j = 25$, so they are “Level-1” players...
 - You could outsmart them anticipating that the average will be 25, and half of it will be 12.5, inducing you to submit $s_i = 12.5$.

Experimental Tests

- If you think that your rivals went through that thought process too, and are submitting $s_j = 12.5$, thus being “Level-2” players.
 - Then half the average will be 6.25, implying that, for you to outsmart them, you must submit $s_i = 3.125$.
- The argument extends to further steps, so that if a player thinks that her rivals are “Level- k ” players, she must submit

$$s_i = \frac{50}{2^{k+1}}.$$

- Interestingly, this argument implies that, as a participant, it is not necessarily optimal for you to submit some of the strategies that survive rationalizability ($s_i = 0$ or $s_i = 1$, as shown above).
 - Instead, you submit a different number depending on the type of Level- k rivals that you deal with, submitting lower numbers as your rivals become more sophisticated (higher k , so they can go through more iterations).

Experimental Tests

- Nagel (1995), for instance, shows that:
 - Many subjects submit $s_i = 25$, suggesting they are Level-1 players; but...
 - Undergraduate students who took Game Theory courses tend to submit $s_i = 12.5$ or $s_i = 6.25$ (thus being Level-2 or -3 players),
 - and so do Caltech students in the Economics major (who generally have a strong Math background)
 - and usual readers of financial newspapers (*Financial Times* in the UK and *Expansion* in Spain).

Finding NBRs in the Cournot Duopoly

From earlier

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 & \text{if } q_2 < 1-c \\ 0 & \text{otherwise} \end{cases}$$

and similarly for firm 2.

- Firm 1's output ranges between $q_1 = \frac{1-c}{2}$ and $q_1 = 0$, implying we can delete all output levels $q_1 > \frac{1-c}{2}$ as NBR for firm 1. So $S'_1 = \left[0, \frac{1-c}{2}\right]$
- Following a similar argument with firm 2's best response function (which is symmetric to firm 1's), we claim that $q_2 > \frac{1-c}{2}$ is NBR for firm 2, yielding a reduced strategy space $S'_2 = \left[0, \frac{1-c}{2}\right]$.
- Therefore, the first step of rationalizability entails that $S'_i = \left[0, \frac{1-c}{2}\right]$ for every firm i .

Finding NBRs in the Cournot Duopoly contd.

- In the second step of rationalizability, we consider firm 1's best response function again, but taking into account that firm 2's rationalizable output levels are those in $S'_2 = \left[0, \frac{1-c}{2}\right]$, that is

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 & \text{if } q_2 < \frac{1-c}{2} \\ 0 & \text{otherwise} \end{cases}$$

- As a consequence, firm 1's output now ranges between:
 - $q_1 = \frac{1-c}{4}$, which occurs when $q_2 = 0$, and
 - $q_1 = \frac{1-c}{2}$, which happens when firm 2 produces $q_2 = \frac{1-c}{2}$.
- This helps us further reduce firm 1's strategy space to $S''_1 = \left[\frac{1-c}{4}, \frac{1-c}{2}\right]$.
 - A similar argument applies to firm 2, yielding $S''_2 = \left[\frac{1-c}{4}, \frac{1-c}{2}\right]$.

Finding NBRs in the Cournot Duopoly contd.

- In the third step of rationalizability, firm 1's best response function becomes:

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 & \text{if } \frac{1-c}{4} < q_2 < \frac{1-c}{2} \\ 0 & \text{otherwise} \end{cases}$$

- Therefore, firm 1's output now ranges between:
 - $q_1 = \frac{1-c}{4}$, which occurs when $q_2 = \frac{1-c}{2}$, and
 - $q_1 = \frac{3(1-c)}{8}$, which happens when firm 2 produces $q_2 = \frac{1-c}{4}$.
- Hence, firm 1's strategy space shrinks to $S_1''' = \left[\frac{1-c}{4}, \frac{3(1-c)}{8} \right]$.
 - Similarly for firm 2's strategy space.

Finding NBRs in the Cournot Duopoly contd.

- Repeating the process again, we see that firm 2's output levels that induce firm 1 to remain active (upper part of firm 1's best response function) keeps shrinking, approaching the 45-degree line, until the point where it does not change in further iterations of rationalizability.
- For this to occur, we need that $q_1 = q_2$. Inserting this property in the best response function yields $q_1 = \frac{1-c}{2} - \frac{1}{2}q_1$. Solving for q_1 , we obtain $q_1 = \frac{1-c}{3}$.
- Later, we will confirm that this output level coincides with the Nash equilibrium of the Cournot game, where $q_1 = q_2 = \frac{1-c}{3}$.

Nash Equilibrium

Definition. **Nash Equilibrium (NE).**

A strategy profile $s^* = (s_i^*, s_{-i}^*)$ is a Nash Equilibrium if every player chooses a best response given her rivals' strategies.

- In a two-player game, the above definition says that:
 - Strategy s_i^* is player i 's best response to s_j^* and, similarly,
 - Strategy s_j^* is player j 's best response to s_i^* .
- For compactness, we often write player i 's best response to s_j^* as:
 - $BR_i(s_j^*) = s_i^*$,
 - and player j 's as $BR_j(s_i^*) = s_j^*$.
- Therefore, a strategy profile is a NE if it is a mutual best response.

Tool 3.3 How to find NEs in matrix games

1. Find the best responses to all players, underlining best response payoffs.
2. Identify which cell or cells in the matrix has every payoff underlined, meaning that all players have a best response payoff. These cells are the NEs of the game.

Example 3.3: Finding NEs in a two-player game

		Firm 2	
		h	m
Firm 1	M	<u>5</u> , <u>3</u>	2,1
	L	3,6	<u>4</u> , <u>7</u>

Matrix 3.3. Finding Best Responses and NEs

Also, show that no player has strictly dominated strategies, implying that IDSDS has no bite:

$$NE = \{(M, h), (L, m)\}$$

NE in common games: Prisoner's Dilemma game

		Player 2	
		<i>Confess</i>	<i>Not Confess</i>
Player 1	Confess	-4,-4	0,-8
	Not Confess	-8,0	-2,-2

Matrix 3.4a. The Prisoner's Dilemma Game

NE in common games: Prisoner's Dilemma game

		Player 2	
		<i>Confess</i>	<i>Not Confess</i>
Player 1	Confess	-4,-4	0,-8
	Not Confess	-8,0	-2,-2

Matrix 3.4a. The Prisoner's Dilemma Game

		Player 2	
		<i>Confess</i>	<i>Not Confess</i>
Player 1	Confess	<u>-4</u> , <u>-4</u>	<u>0</u> , -8
	Not Confess	-8, <u>0</u>	-2,-2

Matrix 3.4b. The Prisoner's Dilemma Game – Underlining Best Response payoffs

$$NE = \{(Confess, Confess)\}$$

NE in common games: Battle of the Sexes game

		Wife	
		Football	Opera
Husband	Football	10,8	6,6
	Opera	4,4	8,10

Matrix 3.5a. The Battle of the Sexes Game

NE in common games: Battle of the Sexes game

		Wife	
		Football	Opera
Husband	Football	10,8	6,6
	Opera	4,4	8,10

Matrix 3.5a. The Battle of the Sexes Game

		Wife	
		Football	Opera
Husband	Football	<u>10,8</u>	6,6
	Opera	4,4	<u>8,10</u>

Matrix 3.5b. The Battle of the Sexes Game – Underlining Best Response payoffs

$$NE = \{(F, F), (O, O)\}$$

NE in common games: Stag Hunt game

		Player 2	
		Stag, S	Hare, H
Player 1	Stag, S	6,6	1,4
	Hare, H	4,1	2,2

Matrix 3.6a. The Stag Hunt Game

NE in common games: Stag Hunt game

		Player 2	
		Stag, S	Hare, H
Player 1	Stag, S	6,6	1,4
	Hare, H	4,1	2,2

Matrix 3.6a. The Stag Hunt Game

		Player 2	
		Stag, S	Hare, H
Player 1	Stag, S	<u>6</u> , <u>6</u>	1,4
	Hare, H	4,1	<u>2</u> , <u>2</u>

Matrix 3.6b. The Stag Hunt Game – Underlining Best Response payoffs

$$NE = \{(S, S), (H, H)\}$$

NE in common games: The Game of Chicken

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	-8,10
	Stay	12,-8	-30,-30

Matrix 3.7a. Anticoordination Game

NE in common games: The Game of Chicken

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	-8,10
	Stay	12,-8	-30,-30

Matrix 3.7a. Anticoordination Game

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	<u>-8,10</u>
	Stay	<u>12,-8</u>	-30,-30

Matrix 3.7b. Anticoordination Game – Underlining Best Response payoffs

$$NE = \{(Swerve, Stay), (Stay, Swerve)\}$$

Multiple Nash Equilibria

- In some games, such as the Battle of the Sexes, the Stag Hunt game, and the Game of Chicken, we found more than one NE.
- The literature often relied on some equilibrium refinement tools:
 - Pareto dominance or Risk dominance, which we will cover at the end of this chapter.
 - More technical tools, such as Trembling Hand Perfect Equilibrium or Proper Equilibrium, which we will cover at the end of Chapter 5 (we need to learn mixed strategy NE first).
- An alternative often used are the so-called “focal points” from Schelling (1960), which players tend to choose in the absence of pre-play communication.
- However, it is unclear *how a specific focal point arises* (i.e., what makes a NE more salient than another):
 - it may vary depending on the time or place where players interact,
 - leading many researchers to experimentally test the emergence of focal points in simultaneous-move games with more than one NE.
 - For experimental tests, see Camerer et al. (2004).

Relationship between NE and IDSDS

If a strategy profile $s = (s_i, s_{-i})$ is a NE, then it must survive IDSDS, but the opposite does not necessarily hold:

$$s \text{ is a NE} \Rightarrow s \text{ survives IDSDS} \\ \nLeftarrow$$

In other words, the NEs in a game is a subset of the strategy profiles surviving IDSDS.

- For the first line of implication, think about (C,C) in the PD game (it is a NE and also survives IDSDS).
- For the second line of implication, think about (F,F) in the BoS game (it survives IDSDS, but it's not a NE).

Relationship between NE and IDSDS




Figure 3.2a. Comparing strategy profiles - Two players.

Games with no NE

Do all games have at least one NE in pure strategies?

No, let's see an example.

Games with no NE

		Player 2	
		Heads	Tails
Player 1	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

Matrix 3.8. Matching Pennies Game

Games with no NE

		Player 2	
		Heads	Tails
Player 1	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

Matrix 3.8. Matching Pennies Game

		Player 2	
		Heads	Tails
Player 1	Heads	<u>1</u> ,-1	-1, <u>1</u>
	Tails	-1, <u>1</u>	<u>1</u> ,-1

Matrix 3.9. Matching Pennies Game – Underlining Best Response payoffs

Evaluating NE as a solution concept

1. Existence? Yes.

- A NE may not exist in pure strategies, but then there must be one when we allow for mixed strategies.
- This result is often known as “Nash’s existence theorem.”
- Requires the use of mixed strategies where players randomize their choices (Chapter 5)

2. Uniqueness? No.

- While the above discussion helps us rest assured that a NE will arise in most games, we cannot claim that every game will have a unique NE.
- As an example, recall BoS or Chicken games where more than one NE exists.

3. Robust to small payoff perturbations? Yes.

- If we alter the payoff of one of the players by a small amount, ε , we would find the same NE than before the payoff change.

Evaluating NE as a solution concept

4. Socially optimal? No.

- As discussed, in the Prisoner's Dilemma game, the NE of a game does not need to be socially optimal.
- A similar argument applies to the following "Modified BoS," where (F,F) and (O,O) are both NEs, but (O,O) Pareto dominates (F,F).

		Wife	
		Football	Opera
Husband	Football	8,8	6,6
	Opera	4,4	10,10

Matrix 3.10. Modified Battle of the Sexes Game

Evaluating NE as a solution concept

	IDSDS	IDWDS	SDE	Rationalizability	NE
Existence	Yes	Yes	No	Yes	Yes
Uniqueness	No	No	Yes	No	No
Robustness of payoff changes	Yes	Yes	Yes	Yes	Yes
Pareto optimal	No	No	No	No	No

Then, it's similar to IDSDS or Rationalizability in its properties, but has more “bite” than any of them.

Appendix

Appendix – Equilibrium Selection

- Coordination and anticoordination games give rise to two NEs.
 - In games with two players but N discrete strategies, we can have N strategy profiles sustained as NEs (see examples in the EconS 424 website).
- In games with more than one NE, a typical question is whether we can select one of them as being more “natural” or “reasonable” to occur according to some criteria.
- The literature offers several *equilibrium selection criteria* and we present the most popular below.
- Throughout our presentation, let $s^* = (s_i^*, s_{-i}^*)$ denote a NE strategy profile.

Appendix – Pareto Dominance

Definition. **Pareto Dominated NE.**

A NE strategy profile s^* Pareto dominates another NE strategy profile $s' \neq s^*$ if $u_i(s^*) > u_i(s')$ for every player i .

Example:

		Wife	
		Tech. A	Tech. B
Firm 1	Tech. A	<u>10,10</u>	0,2
	Tech. B	2,0	<u>5,5</u>

Matrix 3.11. Technology Coordination Game

$$NE = \{(A, A), (B, B)\}$$

Here (A, A) Pareto dominates (B, B) , since every firm earns a higher profit at (A, A) than at (B, B) , i.e., $10 > 5$.

Appendix – Pareto Dominance

- Pareto dominance, however, does not help us rank NEs in all games.

Example

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	<u>-8,10</u>
	Stay	<u>12,-8</u>	-30,-30

Matrix 3.7b. Anticoordination Game – Underlining Best Response payoffs

$$NE = \{(Swerve, Stay), (Stay, Swerve)\}$$

In this case, we cannot select $(Stay, Swerve)$ over $(Swerve, Stay)$ or vice versa, since player 1 prefers $(Stay, Swerve)$ to $(Swerve, Stay)$, while player 2 prefers otherwise.

Appendix – Risk Dominance

- Pareto dominance helps us compare NEs where, for instance, both firms simultaneously change their technology decision from B to A , yielding (A, A) .
- But what if each firm unilaterally deviates from NE?
 - If (B, B) is played, every firm's payoff loss when unilaterally switching to A is $5 - 0 = 5$.
 - But if (A, A) is being played, every firm loses $10 - 2 = 8$ when unilaterally deviating to B .
- Informally, at (A, A) , unilateral deviations are “riskier” than at (B, B) .
 - Note that, as opposed to Pareto dominance, we now seek to understand the payoff loss that players can suffer in the path when each of them unilaterally moves from one NE to another.
- How to measure the “riskiness” of deviations in a more formal way?

Appendix – Risk Dominance

- Definition. **Risk Dominated NE.**
- A NE strategy profile $s^* = (s_1^*, s_2^*)$ risk dominates another NE strategy profile $s' = (s_1', s_2')$ if the total payoff loss of moving from s' to s^* ,

$$[u_1(s_1^*, s_2^*) - u_1(s_1', s_2^*)] \times [u_2(s_1^*, s_2^*) - u_2(s_1', s_2^*)]$$

exceeds that of moving from s^* to s' ,

$$[u_1(s_1', s_2') - u_1(s_1^*, s_2')] \times [u_2(s_1', s_2') - u_2(s_1^*, s_2')]$$

Appendix – Risk Dominance

Example:

		Wife	
		Tech. A	Tech. B
Firm 1	Tech. A	<u>10,10</u>	0,2
	Tech. B	2,0	<u>5,5</u>

Matrix 3.11. Technology Coordination Game

We can say that (B, B) risk dominates (A, A) because the total payoff loss of moving from (B, B) to (A, A) ,

$$\begin{aligned}
 & [u_1(A, A) - u_1(B, A)] \times [u_2(A, A) - u_2(A, B)] \\
 &= (10 - 2) \times (10 - 2) = 64
 \end{aligned}$$

exceeds that of moving from (A, A) to (B, B)

$$\begin{aligned}
 & [u_1(B, B) - u_1(A, B)] \times [u_2(B, B) - u_2(B, A)] \\
 &= (5 - 0) \times (5 - 0) = 25.
 \end{aligned}$$

Chapter 4: Nash equilibria in games with continuous action spaces

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- So far, we studied discrete choice games.
 - Players face a discrete set of actions like “left” or “right”, “high” or “low” etc.
- Players often have continuous action spaces.
- Some examples are
 - Firms choosing prices/quantities.
 - Donors choosing how much to contribute.
 - Participants in an auction choosing how much to bid.
- In this chapter, we will look at some such games.

Topics covered

- Cournot model of quantity competition.
- Bertrand model of price competition.
- Competition among firms with differentiated/heterogeneous goods.
 - Examples include clothing brands, drinks brands etc.
 - Here, products are substitutes of one another, but only to a degree, i.e., if a firm sets a slightly higher price than its rival, it does not drive all of its customers away.
- Donations to a public good.
- Electoral competition

Quantity competition

- Two firms simultaneously and independently choose their output levels.
- Commonly known as the Cournot model of quantity competition, after Antoine Augustin Cournot (1838).
- From example 3.1, we know the best response function.

$$q_i(q_j) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_j & \text{if } q_j < 1-c \\ 0 & \text{otherwise.} \end{cases}$$

- Here, firms regard output as “strategic substitutes” as in example 3.1.
- Note that $0 < c < 1$ by assumption throughout this chapter, unless stated otherwise.

Finding Nash equilibrium: First approach

- We have a system of two equations and two unknowns.
- Solving simultaneously will give a solutions (q_i^*, q_j^*) that are mutual best responses.
- **First approach:** We can *insert one firm's BRF into another* as follows

$$q_i = \frac{1-c}{2} - \frac{1}{2} \overbrace{\left(\frac{1-c}{2} - \frac{1}{2} q_i \right)}^{q_j(q_i)}$$

- which is a function of q_i alone, and yields $q_i^* = \frac{1-c}{3}$.

Mutual best response (Figures)





Figure 4.1a. Firm 2's best response function, $q_2(q_1)$.

Figure 4.1b. NE output in the Cournot game.

Finding Nash equilibrium: Second approach

- **Second approach:** An alternative and faster approach is to *recognize the symmetry* in the game (identical demand and cost functions).
- This entails $q_i^* = q_j^* = q^*$.

- Inserting this property into either firm's best response function, yields

$$q^* = \frac{1-c}{2} - \frac{1}{2}q^* \Leftrightarrow q^* = \frac{1-c}{3}$$

- Second approach is useful only if there is symmetry, while the first works regardless of symmetry.
- Second approach:
 - Quick but must be used cautiously (only after being sure the game is symmetric).
 - And don't assume symmetry in the objective function, you can only assume it after FOCs.

Finding payoffs (Equilibrium profits)

$$\pi^* = (1 - q^* - q^*)q^* - cq^*$$

$$= \left(1 - \frac{1-c}{3} - \frac{1-c}{3}\right) \frac{1-c}{3} - c \frac{1-c}{3} = \frac{(1-c)^2}{9}$$

- which coincides with the square of the equilibrium output i.e.,

$$\pi^* = (q^*)^2$$

Extending quantity competition to $N \geq 2$ firms

- The same model can be extended to a setting in which there are $N \geq 2$ firms.
- Objective function:

$$\max_{q_i \geq 0} (1 - q_i - Q_{-i})q_i - cq_i$$

- where $Q_{-i} = \sum_{j \neq i} q_j$ denotes aggregate output by firm i 's rivals, and q_i is firm i 's choice variable.
- Differentiating with respect to q_i yields

$$1 - 2q_i - Q_{-i} = c$$

Solving for q_i , we obtain best response function:

$$q_i(Q_{-i}) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}Q_{-i} & \text{if } Q_{-i} < 1-c \\ 0 & \text{otherwise.} \end{cases}$$

Same vertical intercept and slope as in a duopoly.

Finding Nash equilibrium output

- We can use the symmetric equilibrium method (second approach) here:
- So Q_{-i}^* must be $Q_{-i}^* = (N - 1)q^*$
- Inserting this property into the BRF, yields

$$q^* = \frac{1 - c}{2} - \frac{1}{2}(N - 1)q^* \Leftrightarrow q^*(N) = \frac{1 - c}{N + 1}.$$

which coincides with that in a duopoly since $q^*(2) = \frac{1 - c}{3}$.

Finding Nash equilibrium output

- Therefore, aggregate output in equilibrium is

$$Q^* = N \times q^*(N) = \frac{N(1-c)}{N+1}$$

which is increasing in N since $\frac{\partial Q^*}{\partial N} = \frac{1-c}{(N+1)^2} > 0$ since $c < 1$ by definition.

- Finding price is now easy by using the inverse demand function:

$$p^* = 1 - Q^* = 1 - \frac{N(1-c)}{N+1} = \frac{1+Nc}{N+1}.$$

which is decreasing in N since $\frac{\partial p^*}{\partial N} = -\frac{1-c}{(N+1)^2} < 0$.

Finding Nash equilibrium output

- Inserting these results in the objective function (profits), we obtain that

$$\pi^* = p^* q^* - c q^* = \frac{1+Nc}{N+1} \frac{1-c}{N+1} - c \frac{1-c}{N+1} = \frac{(1-c)^2}{(N+1)^2},$$

or, more compactly, $\pi^* = (q^*)^2$.

Special cases

- When $N = 1$, we get standard results in a monopoly.
 - $q^* = \frac{1-c}{2}$
 - $p^* = \frac{1+c}{2}$
 - $\pi^* = \frac{(1-c)^2}{4}$
- As $N \rightarrow +\infty$, we get the results in perfect competition.
 - q^* approaches 0
 - $p^* = c$ (marginal cost pricing).

Quantity competition with heterogeneous goods

- Consider a duopoly, but a modified inverse demand function.

$$p_i(q_i, q_j) = 1 - q_i - dq_j; 1 \geq d \geq 0$$

- The last property means that own-price effects dominate cross-price effects.
- Objective function:

$$\max_{q_i \geq 0} (1 - q_i - dq_j)q_i - cq_i$$

- We follow the same steps to find the best response functions, and get

$$q_i(q_j) = \begin{cases} \frac{1-c}{2} - \frac{d}{2}q_j & \text{if } q_j < \frac{1-c}{d} \\ 0 & \text{otherwise.} \end{cases}$$

Intuition for d

- As d goes closer to 1, products are more substitutable (closer to homogenous). The opposite happens as d goes closer to 0.





Figure 4.3. Firm i 's best response function.

Finding Nash equilibrium output

- We continue to rely on symmetry, because though the goods are not identical, they face the same demand function and costs, and even the d faced by both is identical.
- Therefore, the BRF becomes $q^* = \frac{1-c}{2} - \frac{d}{2} q^*$. Solving for q^* , yields $q^* = \frac{1-c}{2+d}$.



Finding Nash equilibrium output

- Equilibrium output $q^* = \frac{1-c}{2+d}$ is decreasing in d .
 - d going to 1 yields a standard Cournot output with no differentiation, $q^* = \frac{1-c}{3}$.
 - d going to 0 yields a monopoly output as the two goods then become completely unrelated, and serve different markets, $q^* = \frac{1-c}{2}$.

Finding Nash equilibrium output

- Equilibrium price is

$$p^* = 1 - \frac{1-c}{2+d} - d \frac{1-c}{2+d} = \frac{1+c(1+d)}{2+d},$$

which is decreasing in parameter d since $\frac{\partial p^*}{\partial d} = -\frac{1-c}{(2+d)^2} < 0$ because $0 < c < 1$ by assumption.

- Equilibrium profits are

$$\pi^* = p^* q^* - c q^* = \frac{1 + c(1 + d)}{2 + d} \frac{1 - c}{2 + d} - c \frac{1 - c}{2 + d} = \frac{(1 - c)^2}{(2 + d)^2}$$

which, for compactness, can also be expressed as $\pi^* = (q^*)^2$, thus being decreasing as products become more homogeneous (higher d).

Price competition

- Consider 2 firms competing in prices.
- This form of competition is known as Bertrand competition, after Joseph Louis François Bertrand (1883).
- Demand/sales for each firm:
 - Since products are undifferentiated, the cheaper good attracts all the customers while rivals make 0 sales.
 - If firms set $p_i = p_j$, they evenly share market demand.
- We can summarize this demand function as follows:

$$q_i(p_i, p_j) = \begin{cases} 1 - p_i & \text{if } p_i < p_j \\ \frac{1 - p_i}{2} & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases}$$

Best response function

$$p_i(p_j) = \begin{cases} p^m & \text{if } p_j > p^m \\ p_j - \varepsilon & \text{if } c < p_j \leq p^m \\ c & \text{if } p_j \leq c \end{cases}$$

- Here $p^m = \frac{1+c}{2}$ denotes monopoly price.
- If firm j sets a price above p^m , i can set the price as p^m and get maximum profits.
- If j sets a price between c and p^m , i should undercut j marginally, setting $p_i = p_j - \varepsilon$, where $\varepsilon \rightarrow 0$.
- If j sets a price below c , any price greater than p_j is a best response. Let us say he sets $p_j = c$.

Best response: Graphically

- The graph below represents a case of “strategic complements”, where the best response of a player increases (decreases) as their rival’s choice increases (decreases). This can be seen from the positive slope.




Figure 4.5. Firm i 's best response function when competing in prices.

Best response (continued)

- Firm j 's best response is symmetric to firm i .

$$p_j(p_i) = \begin{cases} p^m & \text{if } p_i > p^m \\ p_i - \varepsilon & \text{if } c < p_i \leq p^m \\ c & \text{if } p_i \leq c \end{cases}$$




Figure 4.6. Firm j 's best response function when competing in prices.

Finding Nash equilibrium prices

- Following a similar method (where the graphs intersect), we find the NE here too.
- $p_i^* = p_j^* = c$, or we can say the NE pair is $(p_i^*, p_j^*) = (c, c)$, and equilibrium profits are 0.
- Super competitive outcome with just two firms.
 - It is actually unaffected by the number of firms (n-firm Bertrand oligopoly), but...
 - It is affected if firms sell differentiated goods, as we examine next.




Figure 4.7. Price equilibrium.

Price competition with heterogeneous goods

- Consider two firms compete in prices but sell a horizontally differentiated product. Demand is given by

$$q_i(p_i, p_j) = 1 - p_i + \gamma p_j; \gamma \in (0, 2)$$

- Intuition for parameter γ :
 - When $\gamma \rightarrow 0$, j 's price does not affect i 's sales.
 - When $\gamma = 1$, a change j 's price affect i 's sales at a similar rate as a change in its own price.
 - When $\gamma > 1$, a change in j 's price affects i 's sales for than a change in i 's own price.

Finding best response functions

- The objective function for every firm i is

$$\max_{p_i \geq 0} \pi_i(p_i) = (p_i - c)(1 - p_j + \gamma p_j)$$

- Differentiating with respect to p_i , yields

$$1 - 2p_i + c + \gamma p_j = 0$$

- and solving for p_i gives us the best response function:

$$p_i(p_j) = \frac{1 + c}{2} + \frac{\gamma}{2} p_j$$

Best response function (graph)

$$p_i(p_j) = \frac{1+c}{2} + \frac{\gamma}{2}p_j$$

- When firm j increases its price by \$1, firm i responds increasing its own by $\frac{\gamma}{2}$, which is less than \$1 since $\gamma < 2$ by assumption.
 - (Less-than-proportional response.)
- When γ decreases, sales are less affected by its rival's price, flattening the BRF.




Figure 4.8. Firm i 's best response function, $p_i(p_j)$, with heterogeneous goods.

Finding Nash equilibrium prices

- Invoking symmetry, $(p^* = p_1 = p_2)$, the BRF becomes

$$p^* = \frac{1+c}{2} + \frac{\gamma}{2}p^*$$

Solving for p^* , yields the equilibrium price

$$p^* = \frac{1+c}{2-\gamma}$$




Figure 4.9. Equilibrium price pair with heterogeneous goods.

Finding Nash equilibrium prices

- $p^* = \frac{1+c}{2-\gamma}$ is always positive since γ is less than 2.
- p^* is increasing in γ .
 - (Take the derivative and check! It is greater than 0).
 - Crossing point in the figure of BRFs still happens along the 45-degree line, but moves northeast as γ increases.
- Intuitively, as firm i 's sales are more positively affected by its rival's price, p_j , firm i can charge a higher price for its product.

Finding Nash equilibrium prices (graph)

- Additionally, equilibrium output can be found by inserting $p^* = \frac{1+c}{2-\gamma}$ into the demand function, as follows

$$q^* = 1 - p^* + \gamma p^* = 1 - \frac{1+c}{2-\gamma} + \gamma \frac{1+c}{2-\gamma} = \frac{1-(1-\gamma)c}{2-\gamma}.$$

- which is increasing in γ (Check).
 - Therefore, both equilibrium price and output are increasing in product differentiation, γ .
- Equilibrium profits are

$$\pi^* = (p^* - c)q^* = \left(\frac{1+c}{2-\gamma} - c \right) \frac{1-(1-\gamma)c}{2-\gamma} = \left(\frac{1-(1-\gamma)c}{2-\gamma} \right)^2,$$

which, for compactness, can be also expressed as $\pi^* = (q^*)^2$.

- This profit is unambiguously positive, and increasing in γ ,
 - exceeding the zero profits under Bertrand competition with homogeneous products.

Public good game, PGG

- Consider a public project (like a highway or a public library),
 - Two individuals simultaneously and independently choose some contribution x_i and x_j ,
 - where $X = x_i + x_j$ denotes aggregate contributions/donations.
- Public goods are “non rivalrous”, meaning one person enjoying a good does not exclude the other from doing so (roads, parks etc.)

$$u_i(x_i, x_j) = (w_i - x_i)\sqrt{mX}$$

- $m \geq 0$ represents returns from aggregate contributions X . Note that the level of wealth for each player shown by w_i is common knowledge.

Finding best response functions

$$\max_{x_i \geq 0} u_i(x_i, x_j) = (w_i - x_i) \sqrt{m(x_i + x_j)}$$

- where $X = x_i + x_j$.
- Differentiating with respect to x_i , we find that

$$-\sqrt{m(x_i + x_j)} + \frac{m(w_i - x_i)}{2\sqrt{m(x_i + x_j)}} = 0$$

- Which, after rearranging, yields

$$\frac{m(w_i - 2x_j - 3x_i)}{2\sqrt{m(x_i + x_j)}} = 0$$

which holds if $3x_i = w_i - 2x_j$.

Finding best response functions

- Thus, the BRF is
- $x_i(x_j) = \begin{cases} \frac{w_i}{3} - \frac{2}{3}x_j & \text{if } x_j < \frac{w_i}{2}, \text{ and} \\ 0 & \text{otherwise} \end{cases}$

which:

- originates at $\frac{w_i}{3}$ and
- decreases in player j 's contributions, at a rate $\frac{2}{3}$.




Figure 4.9. Player i 's best response function, $x_i(x_j)$.

Finding Nash equilibrium (Symmetric wealth)

- In the symmetric case, $w_1 = w_2 = w$. We can invoke $x_i - x_j = x$. Inserting this into the BRF gives us

$$x = \frac{w}{3} - \frac{2}{3}x$$

- Which yields $x^* = \frac{w}{5}$. This means each player donates a fifth of their wealth. The contribution is increasing in w .

Finding Nash equilibrium (Asymmetric wealth)

- We now consider the more general case where $w_i \neq w_j$. We have to solve the BRFs by substitution.

$$x_i = \frac{w_i}{3} - \frac{2}{3} \underbrace{\left(\frac{w_j}{3} - \frac{2}{3} x_j \right)}_{x_j}$$

- Solving yields $x_i^* = \frac{3w_i - 2w_j}{5}$, which is positive if $w_i > \frac{2}{3}w_j$.
- Similarly, $x_j^* = \frac{3w_j - 2w_i}{5}$, which is positive if $w_i < \frac{3}{2}w_j$.
- *Typical trick:*
 - We solve for the same parameter, w_i , in both equations, so we can obtain two different cutoffs, which we can depict in the (w_i, w_j) quadrant (next slide).

Finding Nash equilibrium (Asymmetric wealth)

- When i 's wealth is lower than $\frac{2}{3}w_j$, only j contributes, and vice versa.
- When $\frac{2}{3}w_j \leq w_i < \frac{3}{2}w_j$, both players contribute.
 - A special case of this is symmetric wealth, as seen in the 45° line of the graph.
- When i 's wealth is higher than $\frac{3}{2}w_j$, she is the only one making positive contributions.




Figure 4.10. Contribution profiles in equilibrium.

Finding Nash equilibrium (Asymmetric wealth)

- Aggregate donations,

$$X^* = \frac{3w_i - 2w_j}{5} + \frac{3w_j - 2w_i}{5} = \frac{w_i + w_j}{5} = \frac{W}{5}.$$

- where $W = w_i + w_j$ denotes aggregate wealth.
- Thus, a fifth of total wealth is contributed to the public good.
- *Example:* if $w_i = \$15$ and $w_j = \$12$,

$$x_i^* = \$4.2,$$

$$x_j^* = \$1.2,$$

$$X^* = \$5.4$$

Inefficient equilibrium

- We can show that the equilibrium donations are “Pareto Inefficient”.
- For simplicity, we can show this result in the case of symmetric wealth. To show this, we first need to maximize joint utility, to find which contribution levels are socially optimal:

$$\max_{x_i, x_j \geq 0} u_i(x_i, x_j) = (w - x_i)\sqrt{mX} + (w - x_j)\sqrt{mX}$$

- This expression simplifies to

$$\max_{X \geq 0} (W - X)\sqrt{mX}$$

- where $W = w_i + w_j$. We also changed the choice variable to X .

Inefficient equilibrium

- Differentiating with respect to X yields

$$-\sqrt{mX} + \frac{m(W - X)}{2\sqrt{mX}} = 0$$

which simplifies to $\frac{m(W-3X)}{2\sqrt{mX}} = 0$. Solving for X , we obtain the aggregate socially optimal donation:

$$X^{SO} = \frac{W}{3}$$

- Therefore, every player's socially optimal donation is $x_i^{SO} = \frac{W}{6} = \frac{2w}{6} = \frac{w}{3}$, which exceeds that in equilibrium because

$$x_i^{SO} = \frac{w}{3} > \frac{w}{5} = x^*$$

- Hence, the privately optimal outcome is not Pareto efficient/socially optimal.
 - *Alternative*: equilibrium donations are “socially insufficient.”
- As expected, aggregate donations satisfy $X^{SO} = \frac{W}{3} > \frac{W}{5} = X^*$

Electoral competition

- Consider two politicians running for president who must simultaneously and independently choose their position in a policy $x_i \in [0,1]$.
- Intuitively, x_i is the proportion of the budget allocated to the policy (0 means none and 1 means the entire budget).
- We assume that the candidates
 - Do not per se have any preferences over the policies, and only want to win
 - Will deliver on the promise (i.e., they are not lying).
- Each voter's ideal policy point is located on this $[0,1]$ interval. Additionally, we assume that voters are
 - Uniformly distributed on $[0,1]$.
 - “Non-strategic”, and simply vote for the candidate whose policy is closest to the voters ideal.

Electoral competition

- If $x_i < x_j$
 - All voters to the left of x_i will vote for i .
 - All voters to the right of x_j will vote for j .
 - For those between x_i and x_j , those between x_i and $\frac{x_i+x_j}{2}$ for i while those between $\frac{x_i+x_j}{2}$ and x_j vote for j .
- A candidate wins when
 - If $x_i < x_j$, i wins if $\frac{x_i+x_j}{2} > 1 - \frac{x_i+x_j}{2}$ or $x_i > 1 - x_j$. Otherwise j wins.
 - If $x_i > x_j$, i wins if $x_i < 1 - x_j$. Otherwise j wins.

Finding best response *correspondences*

- For player i
 - If $x_i < x_j$, i sets x_i such that $x_i > 1 - x_j$ and wins as depicted in Region I.
 - If $x_i > x_j$, i sets x_i such that $x_i < 1 - x_j$ and wins as depicted in Region II.
- For player j
 - If $x_j < x_i$, i sets x_j such that $x_j > 1 - x_i$ and wins as depicted in Region III.
 - If $x_j > x_i$, i sets x_j such that $x_j < 1 - x_i$ and wins as depicted in Region IV.




Figure 4.11. Candidate i 's best response correspondence.

Finding the Nash equilibrium of the game

- From the above figure 4.11, it is clear that mutually best responses only occur at $x_i^* = x_j^* = \frac{1}{2}$.
- This implies that both candidates would choose the midpoint, converging on their political announcements.
- Since voters' ideal policies are uniformly distributed, this announcement coincides with the ideal policy of the median voter.
 - This result is known as the “median voter theorem”.
- The theorem was originally described by Hotelling (1929) and formally shown by Downs (1957).
- For a literature review, see Congleton (2003).

Alternative proof to the electoral competition game

- Consider an asymmetric policy announcement where $x_i \neq x_j$.
- Without loss of generality, assume $x_i < x_j$.
- If $x_i > 1 - x_j$, i wins the election.
 j has incentive to deviate (move closer to i).
- Similarly, if $x_i < 1 - x_j$, i loses, and has incentive to deviate (move closer to j).




Figure 4.12. Votes to candidate i and j when $x_i < x_j$ and $x_i > 1 - x_j$.

Alternative proof to the electoral competition game

- Now consider $x_i = x_j = x$.
- If $x < \frac{1}{2}$, one candidate would receive more votes, and hence, the other has incentive to deviate.
- Similar argument holds if $x > \frac{1}{2}$.
- However, if $x = \frac{1}{2}$, both players receive the same number of votes and the winner is randomly selected.
- Thus, this is the Nash equilibrium (no incentive to deviate).

Chapter 5: Mixed Strategy Nash Equilibrium

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- We have considered games that had at least one NE
 - Prisoner's Dilemma, Battle of the Sexes, and Chicken games
 - But do all games have at least one NE?
- If we restrict players to choose a specific strategy with certainty, some games may not have a NE.
 - This occurs when players do not want to be *predictable* –playing in the same way every time they face the game –such as in board games and professional sports.

Introduction

- Example: Penalty kicks in Soccer

		Kicker	
		<i>Aim left</i>	<i>Aim Right</i>
Goalie	Dive Left	0,0	-10,16
	Dive Right	-10,16	0,0

Matrix 5.1a. Anticoordination Game

- $BR_G(L) = L$; $BR_G(R) = R$ and $BR_K(L) = R$; $BR_K(R) = L$
- Intuitively, the goalie tries to move in the same direction as the kicker, so she prevents the latter from scoring. Meanwhile, the kicker seeks to aim to the opposite location of the goalie to score a goal.

Pure Strategy Nash Equilibrium (psNE)

- Definition. Strategy profiles where players use a specific strategy with 100 percent probability are referred to as “**pure-strategy NE**”

		Kicker	
		<i>Aim left</i>	<i>Aim Right</i>
Goalie	Dive Left	<u>0</u> ,0	-10, <u>16</u>
	Dive Right	-10, <u>16</u>	<u>0</u> ,0

Matrix 5.1b. Anticoordination Game underlining best response payoffs

- In the above game, if we restrict players to use a specific strategy, there will be no mutual best response, and this game will have no NE, i.e., $psNE = \{\emptyset\}$.
- However, if we allow players to *randomize*, such as playing left with probability = $1/3$, and right with the remaining probability = $2/3$, we can find the NE of the game. (These strategies are called Mixed Strategy NE (msNE), which we discuss in later slides.)

Another Example

- Consider a game where police department chooses where to locate most of its police patrols and, simultaneously, a criminal organization decides where to run its business
- Matrix 5.2 shows that the police seeks to choose the same action as the criminal, while the latter seeks to miscoordinate by selecting the opposite location as the police patrol to avoid being caught.

		Criminal	
		Street <i>A</i>	Street <i>B</i>
Police	Street <i>A</i>	<u>10</u> ,0	-1, <u>6</u>
	Street <i>B</i>	0, <u>8</u>	<u>7</u> ,-1

Matrix 5.2. Police and Criminal Game

- A similar argument applies to firm monitoring, such as a polluting firm choosing how many emissions to abate and an environmental protection agency deciding the frequency of its inspections.

Mixed Strategy

- Consider an individual i with a binary strategy set $S_i = \{H, L\}$ representing, for instance, a firm choosing between high and low prices.
- Define a player i 's mixed strategy (or randomization) as a probability distribution over her pure strategies (H and L), as follows $\sigma_i = \{\sigma_i(H), \sigma_i(L)\}$, where

$$\sigma_i(H) = p \text{ and } \sigma_i(L) = 1 - p$$

indicating the probability assigned to each strategy S_i . We require that $\sigma_i(H), \sigma_i(L) \geq 0$ and $\sigma_i(H) + \sigma_i(L) = 1$.

Mixed Strategy

- Definition. **Mixed Strategy.** Consider a discrete strategy set $S_i = \{s_1, s_2, \dots, s_m\}$ where $m \geq 2$ denotes number of pure strategies. The mixed strategy

$$\sigma_i = \{\sigma_i(s_1), \sigma_i(s_2), \dots, \sigma_i(s_m)\}$$

is a probability distribution over the pure strategies in S_i , with the property that:

1. $\sigma_i(s_k) \geq 0$ for every pure strategy s_k , and
 2. $\sum_{k=1}^m \sigma_i(s_k) = 1$
- When mixed strategy concentrates all probability weight on a pure strategy $\sigma_i(s_j) = 1$ while $\sigma_i(s_k) = 0$ for all $j \neq k$, it is commonly called a “degenerated mixed strategy” because, graphically, it collapses to a pure strategy.

Mixed Strategy

- However, to avoid unnecessary complications, we only use the term “mixed strategy” to probability distributions over at least two pure strategies.
- As a remark, the above definition can be applied to games where players choose their strategies from a continuous strategy space, e.g., an output level so that $s_i > 0$.
- In this context, player i 's probability distribution over her pure strategies in S_i can be represented with a cumulative distribution function

$$F_i: S_i \rightarrow [0,1]$$

mapping every strategy $s_i \in S_i$ into cumulative probability.

- For instance, if s_i denotes firm i 's output level, the probability that this firm produces an output level equal or lower than $\bar{s}$ is $F_i(\bar{s})$ and, because $F_i(\bar{s})$ is a probability, it must satisfy $F_i(\bar{s}) \in [0,1]$.

Best Response with Mixed Strategies

- Definition. **Best response with Mixed Strategies.** Player i 's mixed strategy σ_i is a best response to her opponents' mixed strategy σ_{-i} if and only if her expected utility from σ_i satisfies

$$EU_i(\sigma_i, \sigma_{-i}) \geq EU_i(\sigma'_i, \sigma_{-i}) \text{ for all } \sigma'_i \neq \sigma_i$$

- Mixed strategy σ_i is player i 's best response to her opponents' mixed strategy σ_{-i} if no other randomization σ'_i (potentially including the use of pure strategies) yields a higher expected utility than σ_i does.
- We use expected utility because player i needs to compute her expected payoff from randomizing over at least two of her pure strategies and, potentially, her rivals also randomize.

Mixed Strategy Nash Equilibrium (msNE)

- **Definition. Mixed Strategy Nash Equilibrium (msNE).** A strategy profile $(\sigma_i^*, \sigma_{-i}^*)$ is a mixed strategy Nash equilibrium if and only if $\sigma_i^* = BR_i(\sigma_{-i}^*)$ for every player i .
 - Therefore, when player i chooses her equilibrium strategy σ_i^* , she is optimally responding to her opponents' strategies, σ_{-i}^* , implying that players are choosing mutual best responses and, thus, have no incentives to unilaterally deviate.
 - This definition is, then, analogous to that of the pure strategy NE in Chapter 3, but using mixed strategy σ_i^* rather than pure strategy s_i^* .

Goalie Example Revisited

			Kicker	
			Prob. q	Prob. $1-q$
			<i>Aim left</i>	<i>Aim Right</i>
Goalie	Prob. p	Dive Left	0,0	-10,16
	Prob. $1-p$	Dive Right	-10,16	0,0

Matrix 5.3. Anticoordination Game –including probabilities

For **mixed strategy**

- The goalie must be indifferent between diving left and right. Her expected utility from diving left must coincide with that of diving right

$$EU_{Goalie}(Left) = EU_{Goalie}(Right)$$

- Let p and $(1 - p)$ denote the probability with which the goalie randomizes, and q and $(1 - q)$ be the probability with which the kicker randomizes
- Intuitively,
 - $p = 1$ means that the goalie dives left with 100 percent probability, whereas $p = 0$ indicates the opposite
 - $0 < p < 1$ means the goalie randomizes her diving decision
 - Similar argument applies to Kicker's probability, q

Tool 5.1. How to Find msNEs in a two-player game

1. Focus on player 1. Find her expected utility of choosing one pure strategy. Repeat for each pure strategy.
2. Set these expected utilities equal to each other.
3. Solve for player 2's equilibrium mixed strategy, σ_2^* .
 - a. If players have two pure strategies, step 2 just entails an equality. Solving for σ_2 in this context yields a probability, σ_2^* which should satisfy $0 < \sigma_2^* < 1$.
 - b. If players have three pure strategies, step 2 entails several equalities, which gives rise to a system of two equations and two unknowns. The solution to this system of equation is, nonetheless, a player 2's equilibrium mixed strategy, σ_2^* .
4. Focus now on player 2. Repeat steps 1-3, to obtain player 1's equilibrium mixed strategy, σ_1^* .

Example 5.1. Finding msNE: Goalie

		Kicker	
		Prob. q <i>Aim left</i>	Prob. $1-q$ <i>Aim Right</i>
Goalie	Prob. p Dive Left	0,0	-10,16
	Prob. $1-p$ Dive Right	-10,16	0,0

Matrix 5.3. Anticoordination Game –including probabilities

- Goalie's expected utility from diving left:

$$EU_{Goalie}(Left) = \underbrace{q \times 0}_{\substack{\text{kicker aims} \\ \text{left}}} + \underbrace{(1 - q) \times (-10)}_{\substack{\text{kicker aims} \\ \text{right}}} = 10q - 10$$

- Goalie's expected utility from diving right:

$$EU_{Goalie}(Right) = \underbrace{q \times (-10)}_{\substack{\text{kicker aims} \\ \text{left}}} + \underbrace{(1 - q) \times 0}_{\substack{\text{kicker aims} \\ \text{right}}} = -10q$$

- $EU_{Goalie}(Left) = EU_{Goalie}(Right)$

$$10q - 10 = -10q \Rightarrow q = 1/2$$

Example 5.1. Finding msNE: Kicker

Following the same steps we have,

- Kicker's expected utility from aiming left:

$$EU_{Kicker}(Left) = \underbrace{p \times 0}_{\substack{\text{goalie dives} \\ \text{left}}} + \underbrace{(1-p) \times 16}_{\substack{\text{goalie dives} \\ \text{right}}} = 16 - 16p$$

- Kicker's expected utility from aiming right:

$$EU_{Kicker}(Right) = \underbrace{p \times 16}_{\substack{\text{goalie dives} \\ \text{left}}} + \underbrace{(1-p) \times 0}_{\substack{\text{goalie dives} \\ \text{right}}} = 16p$$

- $EU_{Kicker}(Left) = EU_{Kicker}(Right)$

$$16 - 16p = 16p \Rightarrow p = 1/2$$

- $msNE = \left\{ \underbrace{\left(\frac{1}{2} \text{ Dive Left}, \frac{1}{2} \text{ Dive Right} \right)}_{\text{Goalie}}; \underbrace{\left(\frac{1}{2} \text{ Aim Left}, \frac{1}{2} \text{ Aim Right} \right)}_{\text{Kicker}} \right\}$

- Remember that players do not need to randomize with the same probability. They only did in this case because payoffs are symmetric in Matrix 5.3.

Graphical Visualization: msNE

- How to graphically represent the best response of each player?

- Let's start with the **Goalie**

- Goalie chooses to dive left if:

$$EU_{Goalie}(Left) > EU_{Goalie}(Right)$$
$$10q - 10 > -10q \Rightarrow q > 1/2$$

- Mathematically, this means that, for all $q > 1/2$, the goalie chooses to dive left (i.e., $p = 1$).
 - In contrast, for all $q < 1/2$, the goalie responds by diving right (i.e., $p = 0$).

Graphical Visualization: msNE

We can summarize the BRF of Goalie as:

$$BR_{Goalie}(q) \begin{cases} \text{Left if } q > \frac{1}{2} \\ \{Left, Right\} \text{ if } q = \frac{1}{2}, \text{ and} \\ \text{Right if } q < \frac{1}{2} \end{cases}$$

Figure 5.1 depicts this best response function




Figure 5.1. The goalie's best responses.

Graphical Visualization: msNE

- For **Kicker**

- Kicker aims left if:

$$EU_{Kicker}(Left) > EU_{Kicker}(Right)$$

$$16 - 16p > 16 \Rightarrow p < 1/2$$

- Mathematically, this means that, for all $p < 1/2$, the kicker chooses to aim left (i.e., $q = 1$).
- In contrast, for all $p > 1/2$, the kicker aims right (i.e., $q = 0$).

Graphical Visualization: msNE

We can summarize the BRF of Kicker as:

$$BR_{Goalie}(q) \left\{ \begin{array}{l} \text{Left if } p < \frac{1}{2} \\ \{Left, Right\} \text{ if } p = \frac{1}{2}, \text{ and} \\ \text{Right if } p > \frac{1}{2} \end{array} \right.$$

Figure 5.2 depicts this best response function




Figure 5.2. Kicker's best responses.

Graphical Visualization: msNE

- Putting together their response
- Figure 5.3 superimposes the goalie's and the kicker's best response functions
- The player's best responses only cross each other at one point, where $p = q = 1/2$




Figure 5.3. Both players' best responses.

Some Lessons

- Indifference.
 - If it is optimal to randomize over a collection of pure strategies, then a player receives the same expected payoff from each of those pure strategies.
- When analyzing player i , ignore her probability to randomize
 - We make player i indifferent between two (or more) of her pure strategies, we write player i 's expected utility from choosing pure strategy, s_i , $EU_i(s_i)$ as a function of her rival's randomization
- Never use strictly dominated strategies
 - If a pure strategy is strictly dominated, then such a pure strategy cannot be part of psNE or a msNE.
- Odd number of equilibria
 - In almost all finite games (games with a finite set of players and available actions), there is a finite and odd number of equilibria

Extensions: Mixed Strategy Equilibria in Games with $k \geq 3$ pure strategies

- Example: Rock-Paper-Scissors game

		Player 2		
		<i>R</i>	<i>P</i>	<i>S</i>
Player 1	<i>R</i>	0,0	-1,1	1,-1
	<i>P</i>	1,-1	0,0	-1,1
	<i>S</i>	-1,1	1,-1	0,0

Matrix 5.4. Rock-Paper-Scissors Game

		Player 2		
		<i>R</i>	<i>P</i>	<i>S</i>
Player 1	<i>R</i>	0,0	-1,1	1,-1
	<i>P</i>	1,-1	0,0	-1,1
	<i>S</i>	-1,1	1,-1	0,0

Matrix 5.4a. Rock-Paper-Scissors Game – Underlining best response payoffs

- No psNE exists in this game
- Player's payoffs are symmetric entailing that they randomize with the same probabilities: r , p , and $1 - r - p$, where r denotes the probability that every player i chooses Rock, p represents the probability she selects Paper, and $1 - r - p$ is the probability she plays Scissors

Example 5.2. Finding msNE when players have three pure strategies

- If player i randomizes, she must be indifferent between all her three pure strategies

$$EU_i(R) = EU_i(P) = EU_i(S)$$

- We separately find player 1's expected utility from each pure strategy
- To compute $EU_1(R)$, focus at the top row of Matrix 5.4a where:
 - If player 2 chooses Rock, which happens with probability r , player 1 earns a payoff of zero,
 - If player 2 chooses Paper, which happens with probability p , player 1 earns a payoff of -1, and
 - If player 2 chooses Scissors, which happens with probability $1 - r - p$, player 1 earns a payoff of 1
 - In summary, player 1's expected utility from Rock is:

$$EU_1(R) = r0 + p(-1) + (1 - r - p)1 = 1 - r - 2p$$

- Working in a similar way:

$$EU_1(P) = r1 + p0 + (1 - r - p)(-1) = 2r + p - 1$$

$$EU_1(S) = r(-1) + p1 + (1 - r - p)0 = p - r$$

Example 5.2. Finding msNE when players have three pure strategies

- We can now set the expected utilities from Rock, Paper, and Scissors equal to each other.
- First, from $EU_1(R) = EU_1(S)$ we find that:

$$1 - r - 2p = p - r \Rightarrow p = \frac{1}{3}$$

- Second, $EU_1(P) = EU_1(S)$ we find that:

$$2r + p - 1 = p - r \Rightarrow r = \frac{1}{3}$$

- Therefore $1 - r - p = \frac{1}{3}$.
- We can then summarize the msNE of this game as that, every player i , randomizes according to the mixed strategy

$$msNE(r^*, p^*, 1 - p^* - r^*) = \left\{ \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3} \right) \right\}$$

which means that every player assigns the same probability weight to each of her pure strategies.

Finding Mixed Strategy Equilibria in games with $N \geq 2$ players

“Snob Effect” game

- Consider that every player i independently and simultaneously chooses between alternatives X and Y , where X can be interpreted as the snob option, while Y is the conforming option.
- If she is the only player choosing X , she earns a payoff of a (she is the “cool girl” in the group), but a payoff of b otherwise (if anyone else chooses X too)
- When she chooses Y , she earns a payoff of c regardless of how many other players choose X .
- Payoffs satisfy $a > c > b$.

Finding Mixed Strategy Equilibria in games with $N \geq 2$ players

“Snob Effect” game

- Before extending the game to N players, let us consider the two-player version.

		Player 2	
		Prob. p	Prob. $1-p$
Player 1	Prob. p	X	b, b
	Prob. $1-p$	Y	<u>c</u> , <u>a</u>
			c, c

- Two pure strategy NEs: (X,Y) and (Y,X) , as in anticoordination games.
- What about msNE?
 - Player 1 is indifferent between X and Y if and only if $pb + (1 - p)a = pc + (1 - p)c$, which yields $p^* = \frac{a-c}{a-b}$. (Same probability for player 2, since payoffs are symmetric.)
 - Probability p^* is positive and smaller than one since $a > c > b$. Check!
 - p^* is increasing in $a - c$, but decreasing in $a - b$. Intuition.

Finding Mixed Strategy Equilibria in games with $N \geq 2$ players

“Snob Effect” game

- We can now extend the game to N players.
- In this setting, we seek to identify a symmetric msNE where every player i chooses X with probability p .
- Therefore, we need to:
 1. Find the expected utility from X and from Y , and
 2. Set these expected utilities equal to each other to obtain the equilibrium probability p^*

Game continued

- **Expected Utility from X .** When player i chooses X , his expected utility is

$$EU_i(X) = \underbrace{(1-p)^{N-1}}_{\substack{\text{No other player} \\ \text{chooses } X}} a + \underbrace{[1 - (1-p)^{N-1}]}_{\substack{\text{At least someone else} \\ \text{chooses } X}} b$$

To understand the probabilities of each event, recall that the probability with which every player chooses X is p , so the probability with which she plays Y is $1 - p$, implying that all other $N - 1$ players (everyone but player i) choose alternative Y with probability

$$\underbrace{(1-p) \times (1-p) \times \cdots \times (1-p)}_{N-1 \text{ times}} = (1-p)^{N-1}$$

$N-1$ times

Game continued

- **Expected Utility from Y .** When player i chooses Y , she earns a payoff of c with certainty, entailing that $EU_i(Y) = c$.
- **Indifference condition.** If player i randomizes between X and Y , she must be indifferent, thus earning the same expected payoff from each pure strategy, $EU_i(X) = EU_i(Y)$, which means:
$$(1 - p)^{N-1}a + [1 - (1 - p)^{N-1}]b = c$$

Rearranging, yields

$$p^* = 1 - \left(\frac{c - b}{a - b} \right)^{\frac{1}{N-1}}$$

Comparative Statics

- For instance, when $a = 2$, $b = 0$, and $c = 1$, the probability simplifies to

$$p^* = 1 - \left(\frac{1}{2}\right)^{\frac{1}{N-1}}$$

which decreases in the number of players, N .

- Intuitively, as the population grows, the probability that someone selects alternative X increases, driving each player to individually decrease their probability of choosing the snob option X , opting instead for conforming option Y , with probability $1 - p^* = \left(\frac{1}{2}\right)^{\frac{1}{N-1}}$, which increase in N




Figure 5.4. Equilibrium probability p^* .

Comparative Statics

$$p^* = 1 - \left(\frac{c - b}{a - b} \right)^{\frac{1}{N-1}}$$

- We can also examine how the above expression p^* changes when:
 - Players are more attracted to the snob option, $a - b$, which captures the payoff gain of snob option relative to the conforming option.
 - An increase in $a - b$ increases the probability of choosing X .
- In contrast, an increase in $c - b$, which measures the payoff loss of not being the only individual choosing X , decreases the probability of choosing X .

Strictly Competitive Games

- Definition. **Strictly Competitive Games.** A two-player game is strictly competitive if, for every two strategy profiles s and s' ,
 - if $u_1(s) > u_1(s')$ then $u_2(s) < u_2(s')$; and
 - if $u_1(s) = u_1(s')$, then $u_2(s) = u_2(s')$.
- Intuitively:
 - If player 1 prefers strategy profile s to s' , then player 2 has the opposite preference order: preferring s' over s ; and
 - if player 1 is indifferent between s and s' , player 2 must also be indifferent between these two strategy profiles.
- *Example:* The penalty kicks game is an example of a strictly competitive game where we can test the above definition (next slide).

Strictly Competitive Games

		Kicker	
		<i>Aim left (l)</i>	<i>Aim Right (r)</i>
Goalie	Dive Left (L)	0,0	-10,16
	Dive Right (R)	-10,16	0,0

- Comparing (L, l) and (L, r) , we see that the goalie prefers the former, since $0 > -10$, while the kicker prefers the latter because $0 < 16$.
- Comparing (L, l) and (R, l) , we find that the goalie prefers the former, since $0 > -10$, while the kicker prefers the latter because $0 < 16$.
- Comparing (L, l) and (R, r) , we see that the goalie is indifferent, and so is the kicker, both players earning a payoff of zero in both strategy profiles.
- Comparing (R, l) and (L, r) , we find that the goalie is indifferent between these two strategy profiles, earning -10 in both of them. A similar argument applies to the kicker, who earns a payoff of 16 in both strategy profiles.
- We can confirm the definition of strictly competitive games (i.e., opposite preferences of players 1 and 2) holds for every two strategy profiles, s and s' .

Games that are not strictly competitive

- A two-player game is not strictly competitive if, for at least two strategy profiles, s and s' , every player i 's utility satisfies $u_i(s) > u_i(s')$.

- Example

		Criminal	
		Street A	Street B
Police	Street A	10,0	-1,6
	Street B	0,8	7,-1

Matrix 5.6. Police and Criminal Game

- Comparing strategy profiles (A, A) and (B, B) , along the main diagonal, we can see that the police prefers (A, A) to (B, B) , since her payoff satisfies $10 > 7$.
 - Similar argument applies for the criminal, as her payoff satisfies $0 > -1$.
- Because we found that players' preferences over strategy profiles are aligned, rather than misaligned, we can already claim that the game is not strictly competitive without having to compare other pairs of strategy profiles.

Constant-sum Games

- Definition. **Constant-sum games.** A two-player game is a constant-sum game if, for every strategy profile s , player's payoffs satisfy

$$u_1(s) + u_2(s) = K, \text{ where } K > 0 \text{ is a constant.}$$

- Then, players' payoffs must add up to the same constant across all cells in the matrix.
- If, instead, players' payoffs add up to a different number in at least one of the cells, then we can claim that the game is *not* constant sum.
 - It can still be strictly competitive, but not constant sum.

Constant-sum Games

Counterexample:

		Player 2	
		l	r
Player 1	U	10,0	9,3
	D	9,3	10,0

Matrix 5.7. A strictly competitive game that is non constant-sum

- The game is strictly competitive (check as practice).
- It is not a constant-sum game since players payoff in strategy profiles like (U, l) and (D, r) add up to 10, while those strategy profiles (U, r) and (D, l) add up to 12.

Constant-sum Games

- Constant-sum games are always strictly competitive:
 - Condition $u_1(s) + u_2(s) = K$ can be rewritten as $u_1(s) = K - u_2(s)$.
 - Then, if player 1's payoff increases when moving from s to s' , then player 2's payoff must decrease.
- We now introduce a special class of constant-sum games, those in which $K=0$, called zero-sum games.

Zero-sum Games

- Definition. **Zero-sum games.** A two-player game is a zero-sum game if, for every strategy profile s , player's payoffs satisfy

$$u_1(s) + u_2(s) = 0.$$

- Alternatively, condition $u_1(s) + u_2(s) = 0$ can be expressed as $u_1(s) = -u_2(s)$.
 - Intuitively, every dollar that player 1 earns comes from the same dollar that player 2 loses and vice versa.

		Player 2	
		<i>Heads</i>	<i>Tails</i>
Player 1	<i>Heads</i>	1,-1	-1,1
	<i>Tails</i>	-1,1	1,-1

Matrix 5.8. Matching Pennies Game

- Matching pennies game is zero-sum game. Rock-paper-scissors is another example.
- Specifically, in Matrix 5.8, we have that either $1 + (-1) = 0$ or $-1 + 1 = 0$.

Security Strategies

- Definition. **Security Strategies.** In a two-player game, player i 's security strategy, i , solves

$$\max_{\sigma_i} \min_{\sigma_j} u_i(\sigma_i, \sigma_j)$$

- Consider the “worst-case scenario” $w_i(\sigma_i) = \min_{\sigma_j} u_i(\sigma_i, \sigma_j)$
- Player i anticipates that player j chooses her strategy σ_j to maximize her own payoff, which entails minimizing i 's payoff, $u_i(\sigma_i, \sigma_j)$.
 - This is because players interact in a strictly competitive game.
- Player i then chooses her strategy σ_i to maximize the payoff across all worst-case scenarios.
- Intuitively, player i seeks to find the strategy σ_i that provides her with the “best of the worst” payoffs, as represented with the max-min problem.
 - This explains why security strategies are sometimes known as max-min strategies.

Tool 5.2. How to find security strategies in a two-player game

1. Find the expected utility of player 1's randomization, fixing player 2's strategy.
2. Repeat step 1 until you considered all strategies of player 2, fixing one at a time.
3. *"Min" part.* Find the lower envelope of player 1's expected utility. That is, for each strategy σ_1 , find the lowest expected utility that player 1 earns.
4. *"Max" part.* Find the highest expected utility of the lower envelope identified in step 3, and the corresponding strategy σ_1 . This is player 1's security strategy, σ_1^{sec} .
5. To find the security strategy for player 2, follow a similar process in steps 1-4 above.

Example 5.3. Finding Security Strategies

Example

		Player 2	
		l	r
Player 1	U	10,0	9,3
	D	9,3	10,0

Matrix 5.9. A Strictly Competitive Game that is non constant-sum

To find the security strategy for player 1, we follow the next steps:

1. We find player 1's expected utility of randomizing between U and D , with associated probabilities p and $1 - p$, respectively. First, we fix player 2's strategy at column l , which yields:

$$EU_1(p|l) = p \times 10 + (1 - p) \times 9 = 9 + p$$

Example 5.3. Finding Security Strategies

2. We now find her expected utility of randomizing, but fixing player 2's strategy at column r , as follows:

$$EU_1(p|r) = p \times 9 + (1 - p) \times 10 = 10 - p$$

3. To find the lower envelope of the previous two expected utilities, we can depict each line as a function of p , as we do in Figure 5.5. The lower envelope is the segment $9 + p$ for all $p \leq \frac{1}{2}$, but segment $10 - p$ otherwise.




Figure 5.5. Lower envelope and security strategies.

Example 5.3. Finding Security Strategies

4. Among all points in the lower envelope, player 1 enjoys the highest utility at $p = \frac{1}{2}$, which yields an expected payoff $EU_1(p|l) = 9 + \frac{1}{2} = 9.5$, as illustrated in Figure 5.5 by the height of the crossing point between $EU_1(p|l)$ and $EU_1(p|r)$. This is player 1's security strategy, $p^{sec} = \frac{1}{2}$.
5. Following the same steps for player 2, we find that, since payoffs are symmetric, her security strategy is $q^{sec} = \frac{1}{2}$.

Security Strategies and NE

- At this point, you may be wondering about the relationship between security strategies and msNE.
- We obtain the same equilibrium result from both solution concepts, but only for two-player strictly competitive games.
- Consider the previous example from Matrix 5.9:

$$(p^{sec}, q^{sec}) = \left(\frac{1}{2}, \frac{1}{2}\right)$$

Let us now confirm that the msNE produces the same result.

Example 5.4. Solved by msNE

- Assuming that player 1 randomizes between U and D with probabilities p and $1 - p$, respectively, and player 2 mixes l and r with probabilities q and $1 - q$, respectively,

- We find that player 1's expected utility from choosing U is:

$$EU_1(U) = q \times 10 + (1 - q) \times 9 = 9 + q$$

- Similarly, player 1's expected utility from choosing D is:

$$EU_1(D) = q \times 9 + (1 - q) \times 10 = 10 - q$$

- Therefore, player 1 randomizes between U and D when she is indifferent between these two pure strategies $EU_1(U) = EU_1(D)$, which entails

$$9 + q = 10 - q \Rightarrow q = \frac{1}{2}$$

Example 5.4. Solved by msNE

- Player 2's expected utilities
 - $EU_2(l) = p \times 0 + (1 - p) \times 3 = 3 - 3p$, and
 - $EU_2(r) = p \times 3 + (1 - p) \times 0 = 3p$
 - $EU_2(l) = EU_2(r) \Rightarrow 3 - 3p = 3p \Rightarrow p = \frac{1}{2}$
- Summarizing, we can claim that the msNE of this game is $(p, q) = \left(\frac{1}{2}, \frac{1}{2}\right)$, which coincides with the security strategies we found in example 5.3.

Example 5.5. Security strategies and msNE yield different equilibrium outcomes

		Player 2	
		l	r
Player 1	U	3,5	-1,1
	D	2,6	1,2

Matrix 5.10. A Game that is not strictly competitive

- The above game is not strictly competitive. We can find strategy profile where players' interests are aligned; both players prefer, for instance, (U, l) to (D, r) .
- Since the game is not strictly competitive, we can expect that security strategies may produce a different equilibrium prediction than msNE.

Example 5.5. Security strategies and msNE yield different equilibrium outcomes

- For player 1:
 - When player 2 chooses l , player 1's expected payoff from randomizing between U and D with probabilities p and $1 - p$ respectively,
$$EU_1(p|l) = p \times 3 + (1 - p) \times 2 = 2 + p$$
 - When player 2 chooses r , player 1's expected utility is
$$EU_1(p|r) = p \times (-1) + (1 - p) \times 1 = 1 - 2p$$

Example 5.5. Security strategies and msNE yield different equilibrium outcomes

- $EU_1(p|l)$ lies above $EU_1(p|r)$ for all $p \in [0,1]$.
- This means that the lower envelope coincides with $EU_1(p|r) = 1 - 2p$ for all values of p .
- The highest point of this lower envelope occurs at $p = 0$, so player 1 assigns no probability weight to U or, alternatively, that she plays D in pure strategies.
- This means that D is player 1's security strategy.




Figure 5.6. Lower envelope and security strategies - Corner solution.

Security strategies and msNE yield different equilibrium outcomes

Similarly, for player 2

- $EU_2(q|U) = q \times 5 + (1 - q) \times 1 = 1 + 4q$, and
- $EU_2(q|D) = q \times 6 + (1 - q) \times 2 = 2 + 4q$.
- Since $EU_2(q|U) < EU_2(q|D)$ for all values of q , we can claim that U is the lower envelope.
- We can, then, notice that the highest point of $1 + 4q$ occurs at $q = 1$, meaning that player 1 puts full probability weight on l , which becomes his security strategy.

In summary, the security strategy profile in this game is (D, l) .

Example contd. & solving by msNE

- For **msNE**

- We can facilitate our analysis by noticing that strategy l strictly dominates r since it yields a strictly higher payoff than r regardless of the row that player 1 chooses ($5 > 1$ and $6 > 2$).
- We know players put no probability weight in strictly dominated strategies, so we can delete column r from the matrix and obtain:

		Player 2	
		l	
Player 1	U	3,5	
	D	2,6	

Matrix 5.10. A Game that is not strictly competitive – After deleting column r

- Turning now to player 1, we do not need to consider his randomization since, at this point, he has a clear best response to l , U . Therefore, the psNE (no msNE) is (U, l) .
- This equilibrium outcome does not coincide with the security strategy profile (D, l) .

Correlated Equilibrium

- Example

		Player 2	
		<i>Swerve</i>	<i>Stay</i>
Player 1	<i>Swerve</i>	6,6	<u>2,7</u>
	<i>Stay</i>	<u>7,2</u>	0,0

Matrix 5.11. Modified Chicken game

- By underlining best response payoffs, we can find two psNEs: (*Swerve*, *Stay*) and (*Stay*, *Swerve*).
- The game also has msNE, where player 1 chooses *Swerve* with probability $p = \frac{2}{3}$, and the same probability applies to player 2, $q = \frac{2}{3}$ (since payoffs are symmetric).

- In this msNE, Player 1's expected utility in equilibrium is

$$EU_1(\sigma^*) = \underbrace{\frac{2}{3} \left(\frac{2}{3} 6 + \frac{1}{3} 2 \right)}_{\text{Player 1 chooses Swerve}} + \underbrace{\frac{1}{3} \left(\frac{2}{3} 7 + \frac{1}{3} 0 \right)}_{\text{Player 1 chooses Stay}} = \frac{28}{9} + \frac{14}{9} = \frac{14}{3} \cong 4.67$$

Player 1 chooses Swerve *Player 1 chooses Stay*

And a similar expected payoff accrues to player 2.

Correlated Equilibrium

- *Natural question:* Can players reach a higher expected payoff if, instead, they rely on a probability distribution, such as a coin toss, that each player privately observes before playing the game, and that informs the player about which action to choose?
- Intuitively, the probability distribution can be interpreted as an external “recommender” who:
 - First, draws one strategy profile, $s = (s_i, s_{-i})$, such as one cell in Matrix 5.11.
 - Second, the recommender makes recommendation s_i to player i , without informing her of the recommendation s_{-i} that her rivals receive.
- **Definition. Correlated Equilibrium.** A probability distribution over strategy profiles is a correlated equilibrium if every player i follows his recommendation, s_i .
 - Intuitively, a probability distribution over strategy profiles is stable in the sense that every player i has no incentives to unilaterally deviate from the recommendation, s_i .
- For simplicity, we first examine public signals, then privately observed signals.

Example 5.6. Correlated Equilibrium with Public Signals

- Consider the game in Matrix 5.11 and assume that players observe a public signal that assigns probability:
 - α to one of the psNEs in this game, $(Stay, Swerve)$, and
 - $1 - \alpha$ to the other psNE, $(Swerve, Stay)$.
- A public signal could be:
 - A traffic light,
 - Coin toss,
 - Dice toss,

or any other stochastic mechanism that players agree on before starting the game, that yields this probability distribution (summarized in Matrix 5.12).

		Player 2	
		<i>Swerve</i>	<i>Stay</i>
Player 1	<i>Swerve</i>	6,6	<u>2,7</u>
	<i>Stay</i>	<u>7,2</u>	0,0

Matrix 5.11. Modified Chicken game

		Player 2	
		<i>Swerve</i>	<i>Stay</i>
Player 1	<i>Swerve</i>	0	$1 - \alpha$
	<i>Stay</i>	α	0

Matrix 5.12. Correlated equilibrium with public signals – Probability of each strategy profile

Example 5.6. Correlated Equilibrium with Public Signals

		Player 2	
		<i>Swerve</i>	<i>Stay</i>
Player 1	<i>Swerve</i>	6,6	<u>2</u> , <u>7</u>
	<i>Stay</i>	<u>7</u> , <u>2</u>	0,0

Matrix 5.11. Modified Chicken game

		Player 2	
		<i>Swerve</i>	<i>Stay</i>
Player 1	<i>Swerve</i>	0	$1 - \alpha$
	<i>Stay</i>	α	0

Matrix 5.12. Correlated equilibrium with public signals –
Probability of each strategy profile

- In this context, player 1 does not have incentives to deviate.
- Upon observing $(Stay, Swerve)$, payoff from following the recommendation is 7, and that of unilaterally deviating to *Swerve* decreases to 6.
- Similarly, upon observing $(Swerve, Stay)$, payoff is 2, but decreases to 0 if she unilaterally deviates.
- By symmetry, the same argument applies to player 2.
- As a consequence, a continuum of correlated equilibria can be sustained, where players alternate between the two psNEs of the game with probabilities α and $1 - \alpha$, respectively.

Example 5.7. Correlated Equilibrium with Private Signals

- Consider the game in Matrix 5.11 again, where the recommendations assign the same probability weight to $(Swerve, Swerve)$, $(Swerve, Stay)$, and $(Stay, Swerve)$, as summarized in Matrix 5.13:

		Player 2	
		<i>Swerve</i>	<i>Stay</i>
Player 1	<i>Swerve</i>	$\frac{1}{3}$	$\frac{1}{3}$
	<i>Stay</i>	$\frac{1}{3}$	0

Matrix 5.13. Correlated equilibrium with private signals – Probability of each strategy profile

- Bottom row.** Intuitively, if player 1 receives the recommendation of *Stay*, she knows that the only strategy profile recommended by the public signal is $(Stay, Swerve)$.
- Top row.** However, if she receives the recommendation of *Swerve*, she knows that player 2 may have received:
 - The same recommendation, *Swerve*, with probability $\frac{\frac{1}{3}}{\frac{1}{3} + \frac{1}{3}} = \frac{1}{2}$, or
 - The opposite recommendation, *Stay*, with probability $\frac{\frac{1}{3}}{\frac{1}{3} + \frac{1}{3}} = \frac{1}{2}$.
- A similar interpretation applies to the recommendations that player 2 receives.

Example 5.7. Correlated Equilibrium with Private Signals

We can now show that player 1 does not have incentives to deviate from this recommendation profile:

- If he receives the recommendation of *Stay*, his payoff is 7 at (*Stay*, *Swerve*), which she cannot improve by unilaterally deviating to *Swerve* (earning only 6).
- If instead, she receives the recommendation of *Swerve*, her expected payoff is

$$\frac{1}{2}6 + \frac{1}{2}2 = 4,$$

which he cannot improve by unilaterally deviating to *Stay*, as that only yields

$$\frac{1}{2}7 + \frac{1}{2}0 = 3.5.$$

- Since payoffs are symmetric, a similar argument applies to player 2, making the above recommendation stable.
- We can then say that the recommendation profile can be sustained as a correlated equilibrium, with expected payoff:

$$\frac{1}{3}6 + \frac{1}{3}2 + \frac{1}{3}7 = 5$$

which exceeds that in the msNE of the game, 4.67.

Existence of Correlated equilibrium

- Finally, note that every psNE, $s^* = (s_i^*, s_{-i}^*)$ can be defined as a (trivial) correlated equilibrium:
 - where the probability distribution recommends player i to choose s_i^* with probability 1.
 - That is, every player is recommended to play as she would under the psNE.
- A similar argument applies to msNE $\sigma^* = (\sigma_i^*, \sigma_{-i}^*)$, where the probability distribution recommends:
 - Player i to randomize according to the same mixed strategy that she uses in the msNE of the game.
- In other words, if σ^* is a NE, it must also be a correlated equilibrium, but the converse is not necessarily true.

$$\sigma^* \text{ is a NE} \Rightarrow \sigma^* \text{ is a correlated equilibrium} \\ \nLeftarrow$$

- Examples in the Chicken game:
 - (Stay, Swerve) is a NE, and it can be a correlated equilibrium.
 - Any of the public or private randomizations we saw before are correlated equilibria, but they aren't NE.

Equilibrium Refinements in Strategic-form games (Technical)

- Mixed strategies can help us discard NEs which seem fragile to small strategic mistakes, as if a player's hands could "tremble" when choosing her strategy.

		Player 2	
		l	r
Player 1	U	<u>1</u> , <u>1</u>	<u>0</u> , 0
	D	0, <u>0</u>	<u>0</u> , <u>0</u>

- The above game has two psNE: (U,l) and (D,r).
- The second one seems more fragile to trembles:
 - if player 1 deviates from D to U, even if U only occurs with a small probability, player 2's BR would change from r to l.
- A similar argument applies if player 2 deviates from r to l, by a small prob.
- The issue, of course, is that in (D,r) players use weakly dominated strategies.
- We next seek to rule out psNEs that aren't robust to trembles.

Equilibrium Refinements in Strategic-form games (Technical)

- Definition. **Totally mixed strategy.**
- Player i 's mixed strategy, σ_i , is “totally mixed” if it assigns a strictly positive probability weight on every pure strategy, that is $\sigma_i(s_i) > 0$ for all s_i .
- Therefore, all pure strategies happen, even with small probability.
- This allows for trembles, where D could occur with 0.001 probability or less.

Equilibrium Refinements in Strategic-form games (Technical)

- Definition. **Trembling-Hand Perfect equilibrium.**
- A mixed strategy profile $\sigma_i = (\sigma_i, \sigma_{-i})$ is a Trembling-Hand Perfect Equilibrium (THPE) if:
 1. There exists a sequence of totally mixed strategies for each player i , $\{\sigma_i^k\}_{k=1}^\infty$, that converges to σ_i , and
 2. for which $\sigma_i \in BR_i(\sigma_{-i}^k)$ for every k .
- Informally, these two requirements say that:
 1. Every player i 's totally mixed strategy (which allows for trembles) must converge to σ_i ; and
 2. Strategy σ_i is player i 's BR to her rivals' strategy profile σ_{-i}^k at every point of the sequence (i.e., for all k).
- Second requirement is a bit trickier to show. (Example in a moment.)

Properties of THPE

1. Every THPE must be a NE.
2. Every strategic-form game with finite strategies for each player has a THPE.
3. Every THPE assigns zero probability weight on weakly dominated strategies.

Intuitively, points (1) and (2) show that THPEs are a subset of the set of all NEs in a strategic-form game.

$$\sigma \text{ is a THPE} \Rightarrow \sigma \text{ is a NE}$$
$$\nLeftarrow$$

And point (3) helps us rule out strategies D for player 1 and r for player 2 in the 2x2 game we used as a motivation. Therefore, (D,r) is a NE but cannot be supported as a THPE.

Example 5.9. Trembling-hand Perfect Equilibrium

		Player 2	
		l	r
Player 1	U	<u>1</u> , <u>1</u>	<u>0</u> ,0
	D	0, <u>0</u>	<u>0</u> , <u>0</u>

Matrix 5.14. A Game with two psNEs, but only (U, l) is THPE

- Consider the following sequence of totally mixed strategies

$$\sigma_i^k = \left(1 - \frac{\varepsilon_k}{2}, \frac{\varepsilon_k}{2}\right) \text{ for every player } i, \text{ where } \varepsilon_k = \frac{1}{2^k}.$$

- Example:*

- When $k = 1$, $\varepsilon_1 = \frac{1}{2}$, and σ_i^k becomes $\sigma_i^1 = \left(\frac{3}{4}, \frac{1}{4}\right)$, indicating that every player i makes mistakes with $\frac{1}{4}$ probability.
- When $k = 2$, $\varepsilon_2 = \frac{1}{4}$, and σ_i^k becomes $\sigma_i^2 = \left(\frac{7}{8}, \frac{1}{8}\right)$, representing that mistakes are now less likely.


- In the limit, we find that (see figure in next slide)

$$\lim_{k \rightarrow +\infty} \sigma_i^k = (1, 0) \text{ since } \lim_{k \rightarrow +\infty} \varepsilon_k = \lim_{k \rightarrow +\infty} \frac{1}{2^k} = 0$$

- which implies that player 1 (2) chooses U (l, respectively) in pure strategies, yielding strategy profile (U, l) .

Properties

- Generally, as k increase, mistakes become less likely, and the above totally mixed strategy converges to the psNE (U, l) .
- This leads to the following figure:



Example 5.9. Trembling-hand Perfect Equilibrium

- Therefore, the NE (U, l) can be supported as a THPE because:
 1. The totally mixed strategy $\sigma_1^k(\sigma_2^k)$ converges to $U(l)$; and
 2. $U(l)$ is the best response of player 1 (2) to her rival's totally mixed strategy, $\sigma_2^k(\sigma_1^k, respectively)$ for all k .
- To see point (2), note that:
 - When $k=1$, σ_2^k becomes $\sigma_2^1 = (\frac{3}{4}, \frac{1}{4})$, where U is player 1's best response because $EU_1(U|\sigma_2^1) = \frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 0 = \frac{3}{4}$ and $EU_1(D|\sigma_2^1) = \frac{3}{4} \cdot 0 + \frac{1}{4} \cdot 0 = 0$.
 - When $k=2$, σ_2^k becomes $\sigma_2^2 = (\frac{7}{8}, \frac{1}{8})$, and U is still player 1's best response because $EU_1(U|\sigma_2^2) = \frac{7}{8} \cdot 1 + \frac{1}{8} \cdot 0 = \frac{7}{8}$ and $EU_1(D|\sigma_2^2) = \frac{7}{8} \cdot 0 + \frac{1}{8} \cdot 0 = 0$.
 - Same argument applies to every k since $EU_1(U|\sigma_2^k) = (1 - \frac{\epsilon_k}{2}) \cdot 1 + \frac{\epsilon_k}{2} \cdot 0 = 1 - \frac{\epsilon_k}{2} > 0$ and $EU_1(D|\sigma_2^k) = (1 - \frac{\epsilon_k}{2}) \cdot 0 + \frac{\epsilon_k}{2} \cdot 0 = 0$.
 - Same argument applies to player 2's best response to σ_1^k being l for every k . (Check as a practice.)

Example 5.9. Trembling-hand Perfect Equilibrium

- In contrast, (D,r) cannot be sustained as THPE.
 - While we can find converging sequences of totally mixed strategies (first requirement)...
 - Choosing D (r) is *not* player 1's (2's) best response to her rival's totally mixed strategy for every k (second requirement).

- To see this point, consider this totally mixed strategy:

$$\sigma_i^k = \left(\frac{\varepsilon_k}{2}, 1 - \frac{\varepsilon_k}{2} \right) \text{ for every player } i, \text{ where } \varepsilon_k = \frac{1}{2^k}.$$

- which assigns the opposite probability weights than that converging to (U,l).
- It converges to psNE (D,r). Check!

Example 5.9. Trembling-hand Perfect Equilibrium

- However, U is player 1's BR to σ_2^k for every k.
- To see this point, consider that:
 - When k=1, σ_2^k becomes $\sigma_2^1 = \left(\frac{1}{4}, \frac{3}{4}\right)$, and U is player 1's best response.
 - When k=2, σ_2^k becomes $\sigma_2^2 = \left(\frac{1}{8}, \frac{7}{8}\right)$, and U is still player 1's best response.
 - Same argument applies for every k.
 - Recall that finding that U is player 1's BR, instead of D, for *at least one value of k* and for *at least one player* would have been enough to show that (D,r) cannot be sustained as THPE.

ε – Proper Equilibrium

- THPE helps us rule out NEs that aren't robust to trembles.
- But, which trembles do we allow?
- Myerson (1978) suggested that a rational player, while making mistakes, should put:
 - Higher probability weight on strategies yielding higher payoffs.
 - Lower probability weight on strategies yielding lower payoffs.
- Alternatively, players are less likely to make costly mistakes.

ε – Proper Equilibrium

- Definition. ε – **proper equilibrium**. For any $\varepsilon > 0$, a totally mixed strategy profile $\sigma = (\sigma_i, \sigma_{-i})$ is the ε – proper equilibrium if, for every player i , and for every two pure strategies s_i and $s'_i \neq s_i$ such that

$$u_i(s_i, \sigma_{-i}) > u_i(s'_i, \sigma_{-i}),$$

- we must have that probabilities of playing s_i and s'_i , $\sigma_i(s_i)$ and $\sigma_i(s'_i)$ satisfy

$$\varepsilon \times \sigma_i(s_i) \geq \sigma_i(s'_i)$$

- Intuitively, if player i 's expected payoff from choosing s_i is higher than that from s'_i , then...
 - The probability of playing s_i must be at least “ ε times higher” than the probability of playing s'_i .

Example 5.10. ε – Proper Equilibrium

		Player 2	
		l	r
Player 1	U	$\underline{1}, \underline{1}$	$\underline{0}, 0$
	D	$0, \underline{0}$	$\underline{0}, \underline{0}$

Matrix 5.14. A Game with two psNEs, but only (U, l) is THPE

- Consider $\sigma_i = \left(1 - \frac{\varepsilon}{a}, \frac{\varepsilon}{a}\right)$ for every player i , where $a \geq 2$ and $0 < \varepsilon < 1$.
- This mixed strategy is an ε – proper equilibrium because: (1) it is a totally mixed strategy, assigning a positive probability weight to all players' strategies; and (2) for pure strategies U and D , their expected utilities satisfy

$$u_1(U, \sigma_2) = \underbrace{1 \left(1 - \frac{\varepsilon}{a}\right)}_{\text{Player 2 chooses } l} + \underbrace{0 \left(\frac{\varepsilon}{a}\right)}_{\text{Player 2 chooses } r} = 1 - \frac{\varepsilon}{a} > 0 = u_1(D, \sigma_2)$$

*Player 2
chooses l*

*Player 2
chooses r*

Example 5.10. ε – Proper Equilibrium Example

And the probabilities of applying U and D are

$$\varepsilon \times \sigma_1(U) = \varepsilon \left(1 - \frac{\varepsilon}{a}\right) = \frac{\varepsilon(a-\varepsilon)}{a} \text{ and}$$
$$\sigma_1(D) = \frac{\varepsilon}{a}$$

which satisfy

$$\varepsilon \times \sigma_1(U) = \frac{\varepsilon(a-\varepsilon)}{a} \geq \frac{\varepsilon}{a} = \sigma_1(D)$$

since, after rearranging, this inequality simplifies to $a \geq \varepsilon$, which holds given that $a \geq 2$ and $0 < \varepsilon < 1$ by assumption.

(Since the game is symmetric, a similar argument applies to player 2's utility from choosing l and r , and its associated probabilities.)

Proper Equilibrium

- **Definition. Proper Equilibrium.** A mixed strategy profile $\sigma = (\sigma_i, \sigma_{-i})$ is a proper equilibrium if there exists:
 1. A sequence $\{\sigma_i^k\}_{k=1}^{\infty}$ that converges to σ_i for every player i
 2. A sequence $\{\varepsilon_i^k\}_{k=1}^{\infty}$ where $\varepsilon^k > 0$ for all k , that converges to zero
 3. $\{\sigma_i^k\}_{k=1}^{\infty}$ is an ε_k –proper equilibrium for every k
- Proper equilibrium are also THPE, but the converse is not necessarily true.
- In other words:
 - If σ is a proper equilibrium, it must be robust to a sequence of decreasing trembles where costly mistakes are less likely to occur;
 - while σ being THPE only requires that it is robust to any sequence of decreasing trembles.

Example 5.11. Proper Equilibrium

		Player 2	
		l	r
Player 1	U	<u>1</u> , <u>1</u>	<u>0</u> , 0
	D	0, <u>0</u>	<u>0</u> , <u>0</u>

Matrix 5.14. A Game with two psNEs, but only (U, l) is THPE

- The sequence of totally mixed strategies from example 5.9

$$\sigma_i^k = \left(1 - \frac{\varepsilon_k}{2}, \frac{\varepsilon_k}{2}\right) \text{ for every player } i, \text{ where } \varepsilon_k = \frac{1}{2^k},$$

is a proper equilibrium if it satisfies the three requirements in the above definition:

1. A sequence σ_i^k converges to (U, l)
2. ε^k converges to zero
3. σ_i^k is an ε_k -proper equilibrium for every k (as shown in Example 5.10).

Appendix

Fixed-point theorems, an Introduction

- Consider a function $f: X \rightarrow X$, mapping elements from X into X , where $X \subset \mathbb{R}^N$.
- We then say that a point $x \in X$ is a “fixed point” if $x \in f(x)$. For instance, if $X \subset \mathbb{R}$, we can define the distance function

$$g(x) = f(x) - x,$$

- which graphically measures the distance from $f(x)$ to the 45-degree line, as illustrated in Figure 5.8a and 5.8b.




Figure 5.8a. Function $f(x)$ against the 45-degree line.




Figure 5.8b. Distance function $g(x)$.

Fixed-point theorems, an Introduction

- At points such as x' where $g(x') > 0$, we have that $f(x') > x'$, meaning that $f(x')$ lies above the 45-degree line.
- In contrast, at points $x'' > x'$ where $g(x'') < 0$, we have that $f(x'') < x''$, entailing that $f(x'')$ lies below the 45-degree line.





Figure 5.8a. Function $f(x)$ against the 45-degree line.

Figure 5.8b. Distance function $g(x)$.

Fixed-point theorems, an Introduction

- Since $x'' > x'$, if the distance function is continuous, we can invoke the intermediate value theorem to say that:
 - There must be an intermediate value, $\hat{x}$, between x' and x'' (or more than one) where $g(\hat{x}) = 0$, implying that
 - $g(\hat{x}) = f(\hat{x}) - \hat{x} = 0$, or $f(\hat{x}) = \hat{x}$, as required for a fixed point to exist.
- Note that if $f(x)$ was not continuous, then:
 - $g(x)$ would not be continuous either...
 - Allowing for $g(x') > 0$ and $g(x'') < 0$ to occur,
 - yet we could not guarantee the existence of an intermediate point $\hat{x}$ between x' and x'' for which $g(\hat{x}) = 0$.
- Brouwer's fixed-point theorem formalizes this result (next slide).

Brouwer's fixed-point theorem

Definition. **Brouwer's fixed-point theorem.**

If $f: X \rightarrow X$ is a continuous function, where $X \subset \mathbb{R}^N$, then it has at least one fixed point, that is, a point $x \in X$ where $f(x) = x$.

- While Brouwer's fixed-point theorem is useful when dealing with best response functions, it does not apply to best response correspondences.
 - where player i is, for instance, indifferent between two or more of her pure strategies when her opponent chooses strategy s_j .
- The following theorem generalizes Brouwer's fixed-point theorem to correspondences.
- For a more detailed presentation on fixed-point theorems, see Border (1985).

Kakutani's fixed-point theorem

Definition. **Kakutani's fixed-point theorem.**

A correspondence $F: X \rightarrow X$, where $X \subset \mathbb{R}^N$, has a fixed point, that is, a point $x \in X$ where $F(x) = x$, if these conditions hold:

1. X is a compact, convex, and non-empty set.
2. $F(x)$ is non-empty.
3. $F(x)$ is convex.
4. $F(x)$ has a closed graph.

Nash Existence theorem

- First, define player i 's pure strategy set, S_i , to be finite, i.e., a discrete list of pure strategies, and denote a mixed strategy for this player as σ_i , where $\sigma_i \in \Sigma_i$, meaning that player i chooses her randomization among all possible mixed strategies available to her. Therefore, the Cartesian product

$$\Sigma_i \times \Sigma_{-i} = \Sigma$$

denotes the set of all possible mixed strategy profiles in the game, so that every strategy profile $\sigma = (\sigma_i, \sigma_{-i})$ satisfies $\sigma \in \Sigma$.

- Second, let us define player i 's best response correspondence to her rivals choosing σ_{-i} as $\sigma_i \in BR_i(\sigma_{-i})$.

Nash Existence Theorem

- We now define the joint response correspondence $BR(\sigma)$, as the product of $BR_i(\sigma_{-i})$ and $BR_{-i}(\sigma_i)$, that is,

$$BR(\sigma) \equiv BR_i(\sigma_{-i}) \times BR_{-i}(\sigma_i)$$

- Importantly, if BR has a fixed point, then, a NE exists.
- Therefore, we next check if BR satisfies the four conditions on Kakutani's fixed-point theorem, as that would guarantee the existence of a NE.
- Before doing that, we identify:
 - X in Kakutani's fixed-point theorem with the set of all possible mixed strategy profiles, Σ , and
 - correspondence F with BR .

Nash Existence Theorem

1. Σ is a non-empty, compact, and convex set.
 - a. The set Σ is non-empty as long as players have some strategies, so we can identify pure or mixed strategy profiles.
 - b. Recall that if a set is closed and bounded, it is compact. The set of all possible mixed strategy profiles is closed and bounded, thus satisfying compactness.
 - c. Convexity is satisfied since:
 - For any two strategy profiles, σ and σ' ,
 - their linear combination $\lambda\sigma + (1 - \lambda)\sigma'$ where $\lambda \in [0,1]$,
 - is also a mixed strategy profile, thus being part of Σ .

Nash Existence Theorem

2. $BR(\sigma)$ is nonempty.

Since every player i 's payoff, $u_i(\sigma_i, \sigma_{-i})$, is linear in both σ_i and σ_{-i} (expected utility is linear in the probabilities)...

She must find a maximum (a best response to her rivals choosing σ_{-i}) among her available strategies, Σ_i , which we know it is a compact set from point 1b.

Because $\sigma_i \in BR_i(\sigma_{-i})$ and $\sigma_{-i} \in BR_i(\sigma_i)$ are non-empty (best response exists), then their product, $BR(\sigma)$, must also be non-empty.

3. $BR(\sigma)$ is convex.

To prove this point, consider two strategies for player i , σ_i and σ'_i , that are best responses to her rivals choosing σ_{-i} , that is, $\sigma_i, \sigma'_i \in BR_i(\sigma_{-i})$.

Because both σ_i and σ'_i are best responses, they must both yield the same expected payoff; otherwise, one of them cannot be a best response.

Therefore, a linear combination of σ_i and σ'_i , $\lambda\sigma_i + (1 - \lambda)\sigma'_i$ where $\lambda \in [0,1]$, must yield the same expected payoff as σ_i and σ'_i , thus being a best response as well, that is, $\lambda\sigma_i + (1 - \lambda)\sigma'_i \in BR_i(\sigma_{-i})$.

Nash Existence Theorem

4. $BR(\sigma)$ has a closed graph.

This property means that the set $\{(\sigma_i, \sigma_{-i}) \mid \sigma_i \in BR_i(\sigma_{-i})\}$ is “closed,” meaning that every player i 's best response correspondence has no discontinuities.

The best responses depicted in this chapter, for instance, showed no discontinuities. Because every player i 's payoff, $u_i(\sigma_i, \sigma_{-i})$, is continuous and compact, the set $\{(\sigma_i, \sigma_{-i}) \mid \sigma_i \in BR_i(\sigma_{-i})\}$ is closed.

Nash Existence Theorem

- The previous 4 properties guarantee that a NE exists when players face finite strategy spaces (i.e., a list of pure strategies).
- What if they choose their strategies from a continuous strategy space, as when firm set their prices or output levels?
 - Glicksberg (1952) extended the above result to setting with continuous strategy spaces, where $S_i \subset \mathbb{R}^N$, showing that, if:
 1. Every player i 's strategy space, S_i , is compact, and
 2. Her utility function, $u_i(\cdot)$, is continuous,
 - Then, a NE exists, in pure or mixed strategies.
 - (For generalization of this result to non-continuous utility functions, see Dasgupta and Maskin (1986).)

Chapter 6: Subgame Perfect Equilibrium

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Sequential Move Game

- Why not just solve them using NE?
 - We would find too many NEs
 - Some NEs would be “sequentially irrational,” as we define in this chapter.
- We then need players to behave rationally:
 - Every player maximizes her payoff when called to move (on a node or on an information set)...
 - given her position on the game tree, and
 - given her available information.
- Strategy profiles where every player is sequentially rational will be called Subgame Perfect Equilibria (SPEs).
- We will consider games with:
 - Discrete and continuous strategies.
 - When players observe their rivals’ previous actions, and when they do not.
- Before we start, we need to specify some “game tree rules.”

Tree Rules

1. Every node is the successor of the initial node.

- Figure 1 illustrates this property.
- Game tree in the left satisfies it, since there is only one initial node.
- The game tree in the right panel violates it.
- If a modeler wanted to indicate that two different players act simultaneously...
 - Then the figure 6.1a at the bottom is the correct way to represent this.
 - Player 1 selects between B or C , then player 2 chooses c or d without observing player 1's choice (as they are simultaneous).
- If, instead, the modeler considers the same player in two initial nodes (right panel in Figure 6.1), there must be a mistake because Player 1 cannot choose simultaneously between A and B (on the top of the figure), and between C and D (on the bottom), as if he had multiple personalities!
 - In this case, player 1's decision should happen at the subsequent stage.




Figure 6.1. Games with one or two initial nodes.




Figure 6.1a: two different players acting simultaneously

Tree Rules

2. Every node has exactly one immediate predecessor; except the initial node that has no predecessor.

- The left panel in Figure 6.1 satisfies this property.
- But Figure 6.2 depicts a game tree that violates it.
- If a tree has nodes with more than one predecessor, we could run into misunderstandings:
 - In Figure 6.2, we do not know if player 4 is called to move after player 2 chose b or because player 3 chose c




Figure 6.1. Games with one or two initial nodes.





Figure 6.2: Game with two predecessors


Tree Rules

3. Multiple branches extending from the same node have different action labels.

- Left panel in Figure 6.3 satisfies this property
- The right panel violates it because label *A* is on two branches for player 1.
- If the modeler seeks to represent that:
 - Player 1 chooses *A* in both the top and middle branches that stem from the initial node, then these two branches should be collapsed into a single branch.
 - If, instead, the modeler tries to represent that player 1 has two different actions in each of these branches, then they should be labelled differently to avoid misunderstandings.



Correct



Incorrect

No, you must be referring to a different action. Otherwise collapse everything under the same name.

Figure 6.3. Games with same/different action labels.

Tree Rules

4. Each information set contains nodes for only one player.

- The left panel in Figure 6.4 satisfies this property.
- The right panel violates it.
- Intuitively, if player 2 is called to move at the top node of the game tree in the right panel, she can infer that player 1 selected A , leaving her no uncertainty about player 1's choices in the previous stage.
 - A similar argument applies to player 3: if she is called to move at the bottom node, she knows that player 1 selected B .
- Therefore, the information set connecting the top and bottom nodes is incorrect (unnecessary).




Figure 6.4. Information sets with one or more players.

*Incorrect, otherwise P3 knows he is called on to move after player 1 chose B.
(no uncertainty)*

Tree Rules

5. Same number of branches and labels.

- *1st requirement.* All nodes in a given information set have the same number of branches stemming from them (i.e., the same number of immediate successors)
 - Figure 6.5a depicts a game tree that violates this:
 - Player 2 would be able to infer whether player 1 chose *Invest* or *Not Invest* by just observing the number of strategies.
 - Drawing the information set is unnecessary in this setting!
- *2nd requirement.* The labels in the branches must coincide across all nodes connected by the same information set.
 - Figure 6.5b depicts a game tree that violates this:
 - Player 2 would be able to infer that:
 - Player 1 must have chosen *Invest* if she must select between *A* and *B*, but...
 - Player 1 must have chosen *Not Invest* if she must select between *A* and *C*.




Figure: 6.5a. Different number of branches




Figure 6.5b. Different immediate successors.

Actions vs. Strategies

- Graphically, an **action** denoted as a_i or $s_i(h_i)$, is just a specific branch that player i chooses when it is her turn to move.
 - Note that the notation is to emphasize one action among those available in information set $h_i, A_i(h_i)$.
- A **strategy** is the list of all the branches that she would choose along the game tree, both in nodes that are reached in equilibrium and those that are not reached.
 - Complete contingent plan.
- **Definition. Pure Strategy.**
 - In a sequential-move game, a pure strategy for player i is a mapping
$$s_i: H_i \rightarrow A_i$$
that assigns an action $s_i(h_i) \in A_i(h_i)$ at information set $h_i \in H_i$, where $A_i(h_i)$ denotes the set of available actions at information set h_i .

Actions and Strategies

- Importantly, the above definition of pure strategy applies to *every* information set, $h_i \in H_i$, describing how player i behaves once she reaches each information set h_i .
- In contrast, an action only describes how player i behaves when reaching a specific information set h_i .
- Example (Figure 6.3):
 - In denotes player 2's action at the top node of the game tree.
 - (In, Out', In'') denotes her strategy, indicating that she responds with:
 - In after player 1 chooses A ,
 - Out' after player 1 selects B , and
 - In'' after player 1 chooses C .




Figure 6.3. Games with same/different action labels.

Why don't we just find the Nash Equilibria of the Game Tree?

- Example. **Entry Game.**

- Consider the Entry game in Figure 6.6.
- A firm (potential entrant) chooses whether to enter into an industry, where an incumbent firm operates as a monopolist, or stay out of the market.
- If the potential entrant stays out (at the bottom of the figure):
 - The incumbent remains a monopolist, earning \$10 while the potential entrant earns zero.
- However, if the entrant joins the industry (top of the figure), the incumbent observes this decision and responds either:
 - Accommodating the entry (e.g., setting moderate prices) which yields a payoff of \$4 for each firm, or
 - Fighting it (e.g., starting a price war against the entrant) leading to a payoff of -\$2 for each firm.




Figure 6.6. The entry game.

Entry Game (contd.)

- To find the NEs of this game tree, we first need to represent the game in matrix form.

		Potential Entrant	
		In	Out
Incumbent	Accommodate	<u>4</u> , <u>4</u>	<u>10</u> ,0
	Fight Entry	-2,-2	<u>10</u> , <u>0</u>

Matrix 6.1a. Finding NEs in the Entry Game

- There are two psNE for this game:
 $NE = \{(Accommodate, In) \text{ and } (Fight, Out)\}$

Example 6.1. Entry Game (contd.)

- Specifically, for the incumbent, we find that:
 - $BR_{inc}(In) = Acc$ because $4 > -2$
 - $BR_{inc}(Out) = \{Acc, Fight\}$ because both yield a profit of 10
- For the potential entrant,
 - $BR_{ent}(Acc) = In$ because $4 > 0$
 - $BR_{ent}(Fight) = Out$ because $0 > -2$
- Therefore:
 - In the first NE, entry occurs and the incumbent responds accommodating.
 - In the second NE, the entrant does not enter because it believes that the incumbent will start a price war.

Example 6.1. Entry Game (contd.)

- But is this belief credible?
 - NO! The entrant beliefs about the incumbent's decision to *Fight* after he enters are not rational (in a sequential way):
 - Once the entrant is in, the best thing that the incumbent can do it to accommodate ($4 > -2$).
 - Hence, the incumbent would never have incentives to start a price war after entry has already occurred.
 - Therefore, among the two NEs:
 - Only (*Accommodate*, *In*) is sequentially rational,
 - (*Fight*, *Out*) is not sequentially rational.

Introducing a new solution concept

- In the Entry Game, we found NEs that were sequentially irrational.
- Can we find only those equilibria that are sequentially rational? Yes!
- We can guarantee sequential rationality by using a new solution concept, **Subgame Perfect Equilibrium (SPE)**.
- For this solution concept, we first need to define what we mean by a “subgame.”
- Definition. **Subgame:**
- A subgame is a tree structure defined by a node and all its successors.
 - This definition means that:
 - if nodes a and b are connected with an information set (so player i does not know whether he is at node a or b), then both nodes must be part of the *same* subgame.
 - Examples in next slide.

Subgames - Example

- Graphically, a subgame can be identified by drawing a rectangle or a circle around a section of the game tree without “breaking” any information set.
- Figure 6.7a depicts the Entry Game, identifying only two subgames:
 - that initiated when the incumbent is called on to move; and
 - the game as a whole.




Figure: 6.7a. Subgames in the entry game

Subgames – Example

- There are 4 subgames in Figure 6.7b:
 1. After player 2 is called to move after player 1 chooses *Up*
 2. After player 2 is called to move after player 1 chooses *Down*
 3. After player 3 is called to move, which happens when player 2 responds to *Down* with *C*
 4. The game as a whole




Figure 6.7b. Four subgames in a game tree.

What if the Game Tree has Information Sets?

- Figure 6.8 depicts a game tree with an information set.
 - The smallest subgame must include player 2's information set (otherwise we would be breaking it!)
- The presence of information sets reduces the number of subgames.
- In Figure 6.8, player 2 does not observe whether player 1 selected a or b .
 - This entails that player 2, when choosing whether to respond with c or d , operates "as if" player 1 was selecting a or b at the same time.
 - This is equivalent to a setting where players interact in a *simultaneous-move game*.




Figure: 6.8. Finding subgames in a game with an information set

Subgame Perfect Equilibrium

- Definition. **Subgame Perfect Equilibrium (SPE).**
- A strategy profile (s_i^*, s_{-i}^*) is a Subgame Perfect Equilibrium if it specifies a NE in each subgame.
- To find SPEs in a sequential-move game, we just need to apply the notion of *backward induction*.

Tool 6.1. Applying backward induction

1. Go to the farthest right side of the game tree (where the game ends), and focus on the last mover.
2. Find the strategy that yields the highest payoff for the last mover.
3. Shade the branch that you found to yield the highest payoff for the last mover.
4. Go to the next-to-last mover and, following the response of the last mover that you found in step 3, find the strategy that maximizes her payoff.
5. Shade the branch that you found to yield the highest payoff for the next-to-last mover.
6. Repeat steps 4-5 for the player acting before the previous-to-the-last mover, and then for each player acting before her, until you reach the first mover at the root of the game.

Example 6.2. Applying Backward Induction – Entry Game

Before we start, recall this game has two subgames.

1st step: Applying backward induction, we first focus on the last subgame (closest to the terminal nodes).

- This corresponds to the last mover, the incumbent.

2nd step: Comparing its payoff from accommodating entry, 4, and starting a price war, -2 , we find that its best response to entry is to accommodate.

We shade the corresponding branch in Figure 6.9 to keep in mind the optimal response of the incumbent in this subgame.




Figure: 6.9. Applying Backward Induction in the entry game –Last Move

Example 6.2. Applying Backward Induction – Entry Game

3rd step: We move to the first mover, the entrant, who anticipates that, if it enters, the incumbent will *accommodate*.

- This means that the entrant expects that, upon entry, the game will proceed through the shaded branch in Figure 6.10 yielding a payoff of 4 from entering.
- If, instead, the entrant stays out, its payoff is only 0. As a consequence, the optimal strategy for the entrant is to enter.




Figure: 6.10. Applying Backward Induction in the entry game –First Mover

Example 6.2. Applying Backward Induction – Entry Game

- Hence:

$$SPE = \{(In, Accommodate)\}$$

- Among the *two psNE* we found, i.e., $(In, Accommodate)$ and $(Out, Fight)$,
 - only the former is sequentially rational.




Figure: 6.10. Applying Backward Induction in the entry game –First Mover

Example 6.2. Applying Backward Induction – Entry Game

- Figure 6.10 shades the branches that players choose in equilibrium, known as the “*equilibrium path*”.
- The equilibrium path of play is a visual tool to understand how players behave in equilibrium:
 - From the initial node to one of the terminal nodes in the tree,
 - but does not coincide with the SPE in more involved games.




Figure: 6.10. Applying Backward Induction in the entry game –First Mover

Example 6.2. Applying backward induction – Entry Game

- Recall that $(In, Accommodate)$ was also one of the NEs in the Entry game, as shown in Example 6.1.
- Then Examples 6.1 and 6.2 illustrate that:
 - Every SPE must be a NE, since $(In, Accommodate)$ is both a SPE and one of the NEs in the Entry game, but...
 - The converse is not necessarily true.
 - More formally, for a strategy profile s^* ,

$$s^* \text{ is a SPE} \Rightarrow s^* \text{ is a NE}$$

$\nLeftarrow$

- Alternatively, the set of strategy profiles that can be supported as SPEs of a game is a subset of those strategy profiles that can be sustained as NE, as depicted in Figure 6.11.




Figure 6.11. SPEs as a subset of NEs.

Example 6.3. Applying Backward Induction in the Modified Entry Game

- Consider the modified version of the Entry Game depicted in Figure 6.12.
- The top part of the game tree coincides with that in the original Entry game.
- However, if the entrant chooses to stay out of the industry, the incumbent can now respond investing in a new technology or not, with associated payoffs $(0,12)$ and $(0,10)$, respectively.




Figure 6.12. Modified entry game.

Example 6.3. Applying Backward Induction in the Modified Entry Game

Three subgames:

1. one initiated after the entrant joins the industry
2. another initiated after the entrant remains out and
3. the game as a whole




Figure 6.13. Modified entry game - Subgames.

Example 6.3. Applying Backward Induction in the Modified Entry Game




Figure 6.13. Modified entry game - Subgames.

- Anticipating *Acc* upon entry and *Invest* upon no entry, the entrant can, in the initial node, compare its payoff from entry, 4, and from no entry 0, thus choosing to enter.

- Therefore:

$$SPE = \{(In, Accommodate/Invest)\}$$

which denotes that the entrant chooses *In*, and the incumbent responds with *Acc* after *In*, but with *Invest* after *Out*.

Equilibrium Path vs. SPE




Figure: 6.14 Equilibrium Path vs. SPE

- Figure 6.14 shades the branches that players choose in equilibrium.
- The equilibrium path, coincides with that in the original Entry game.
- The SPE, however, is more intricate because it specifies the incumbent's equilibrium behavior both in the node that it reaches in equilibrium (when the entrant joins the market) and in the node that the incumbent does not even reach in equilibrium (when the entrant stays out).
- Thus, when describing the SPE of a sequential-move game, it must specify the equilibrium behavior for every player *at every node* where she is called to move, even in nodes that may not be reached in equilibrium.

Finding SPEs in Game Trees with Information Sets

- Consider Figure 6.15.
- If firm 1 chooses *Up*, this firm gets to play again, choosing between *A* and *B*.
- Firm 2 is then asked to respond, but without seeing whether firm 1 chose *A* or *B*.
- Firm 2's uncertainty is graphically represented by the dotted line connecting the end of the branches that it doesn't distinguish, *A* and *B*.
- This dotted line is formally known as an "information set" for Firm 2, because this firm doesn't know which of these two actions was chosen by Firm 1.




Figure 6.15. A more involved game tree.

Finding SPEs in Game Trees with Information Sets

- Before applying backward induction to this game, a usual trick is to find all the subgames (i.e., circling the portions of the tree that do not break any information set)




Figure: 6.15(a) Finding subgames




Figure: 6.15(b) Not subgames

Finding SPEs in Game Trees with Information Sets

- **Subgame 1.**

- Firm 2 does not observe which action firm 1 chose (either A or B). Therefore, subgame 1 can be represented using Matrix 6.2.
- Matrix 6.3 underlines the best response payoffs.
- $NE of \text{subgame 1} = \{(A, X)\}$

		Firm 2	
		X	Y
Firm 1	A	3,4	1,4
	B	2,1	2,0

Matrix 6.2. Representing Subgame 1 in matrix form

		Firm 2	
		X	Y
Firm 1	A	<u>3</u> , <u>4</u>	1, <u>4</u>
	B	2, <u>1</u>	<u>2</u> ,0

Matrix 6.3. Finding the NE of Subgame 1

Finding SPEs in Game Trees with Information Sets

- **The Game as a whole.**

- Firm 1 must choose between *Up* and *Down*, anticipating that if it chooses *Up*, subgame 1 will start.
- Firm 1 can anticipate equilibrium behavior in subsequent stages of the game; that is, the NE of subgame 1 is (A, X) , with payoff $(3, 4)$.
- Firm 1 can then simplify its decision problem to the tree depicted in Figure 6.16. Therefore, firm 1 only needs to compare the following: if it chooses *Down*, the game is over and its payoff is 2, whereas if it chooses *Up*, subgame 1 is initiated, obtaining a payoff of 3.
- Because $3 > 2$, firm 1 prefers to choose *Up* rather than *Down*, illustrated by the thick arrows on the branch corresponding to *Up*.
- The SPE of this game is $(Up, (A, X))$.




Figure 6.16. The reduced game tree from figure 6.15a.

Evaluating SPE as a Solution Concept

1. Existence? Yes.

- When we apply backward induction in any game tree, we find that at least one equilibrium exist.
 - This is often known as “Zermelo’s theorem”, after Zermelo (1913) article on chess, and later on extended by Kuhn (1953).
- Intuitively, this means that every player selects on strategy at every node where she is called to move.
- If a player is indifferent between two available actions a and b , for instance, then her best response is to choose either.
 - But every player selects on strategy at every node, yielding at least one SPE strategy profile.

Evaluating SPE as a Solution Concept

2. Uniqueness? No.

- While the examples in this chapter display game trees that produce a unique SPE, we cannot guarantee that applying backward induction will always induce a unique SPE, as required by this criterion.
- For instance, if in the Entry game of Example 6.2 the incumbent's payoff from both accommodating and fighting entry was 4, then this firm would be indifferent between *Acc* and *Fight*, leading to two SPEs:

$$SPE = \{(In, Acc), (In, Fight)\}$$

- In this case, the SPEs exactly coincide with those that are NEs.
- Interesting point:
 - If players are not indifferent about any of their actions at any of the nodes (or information sets) where they are called to move...
 - equilibrium behavior in each subgame must specify a *unique strategy for every player in that subgame*,
 - so the SPE must be unique.

Evaluating SPE as a Solution Concept

3. Robust to small payoff perturbations? Yes.

- If we change the payoff of one of the players by a small amount (e.g. 0.001, but generally, for any ε that approaches zero), backward induction still provides us with the same equilibrium outcome, implying that SPE is robust to small payoff perturbations.
- This is due the fact that:
 - If a strategy s_i yields a strictly higher payoff than another strategy s'_i ,
 - it must still yield a strictly higher payoff than s'_i after we apply a small payoff perturbation (remember that ε can be infinitely small),
 - meaning that play i still chooses s_i in the subgame where this action was optimal.

Evaluating SPE as a Solution Concept

4. Socially optimal? No.

- The application of backward induction does not necessarily produce a socially optimal outcome, which goes in line with our evaluation of IDSDS and NE, which did not yield socially optimal outcomes either.

Application: Stackelberg Game of Sequential Quantity Competition

- Consider two firms produce a homogeneous good facing a linear demand function $p(Q) = 1 - Q$, where $Q = q_1 + q_2$ denotes aggregate output.
- Assume that all firms have a constant marginal cost of production c , where $1 > c \geq 0$.
- Firms interact in a sequential-move game:
 1. In the first stage, firm 1 (the industry leader) chooses its output q_1 .
 2. In the second stage, firm 2 (the industry follower) observes that the leader's output q_1 and responds with its own output level q_2 .
- Solving the game by backward induction, we start by solving the second stage.

Application: Stackelberg Game of Sequential Quantity Competition

- **Second stage, follower – firm 2.**

Firm 2 maximizes its profit as follows

$$\max_{q_2 \geq 0} \pi_2 = (1 - q_1 - q_2)q_2 - cq_2$$

Differentiating with respect to q_2 :

$$1 - 2q_2 - q_1 - c = 0$$

Solving for q_2 yields:

$$q_2 = \frac{1 - c}{2} - \frac{1}{2}q_1$$

which is positive for all $\frac{1-c}{2} - \frac{1}{2}q_1 \geq 0 \Rightarrow q_1 \leq 1 - c$.

Application: Stackelberg Game of Sequential Quantity Competition

Firm 2's BRF:

$$q_2(q_1) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_1 & \text{if } q_1 \leq 1-c \\ 0 & \text{otherwise.} \end{cases}$$

- Intuitively, the follower produces an output of $q_2 = \frac{1-c}{2}$ when the leader is inactive ($q_1 = 0$), but decreases its output in q_1 at a rate of $\frac{1}{2}$.
- When the leader's output is sufficiently large, $q_1 > 1 - c$, the follower responds staying inactive, $q_2 = 0$.
- This BRF coincides with that under the Cournot model of simultaneous quantity competition in chapter 4.

Application: Stackelberg Game of Sequential Quantity Competition

- **First stage, leader – firm 1.**

Firm 1 maximizes its profit as follows

$$\max_{q_1 \geq 0} \pi_1 = [1 - q_1 - q_2(q_1)]q_1 - cq_1$$

which is evaluated at $q_2(q_1)$. Inserting firm 2's BRF, we obtain

$$\max_{q_1 \geq 0} \pi_1 = \left[1 - q_1 - \overbrace{\left(\frac{1-c}{2} - \frac{1}{2}q_1 \right)}^{q_2(q_1)} \right] q_1 - cq_1 = \left[\frac{1+c}{2} - \frac{1}{2}q_1 \right] q_1 - cq_1$$

Differentiating with respect to q_1 :

$$\frac{1+c}{2} - q_1 - c = 0$$

Solving for q_1 yields:

$$q_1^{seq} = \frac{1-c}{2}$$

Application: Stackelberg Game of Sequential Quantity Competition

- The SPE in this Stackelberg game of quantity competition:

$$SPE = \left(q_1^{Seq}, q_2(q_1) \right) = \left(\frac{1-c}{2}, \frac{1-c}{2} - \frac{1}{2}q_1 \right)$$

where firm 2 (the follower) response with its BRF, $q_2(q_1)$, for any output that the leader chooses (both its equilibrium output, $q_1 = q_1^{Seq}$, and any off-the-equilibrium output, $q_1 \neq q_1^{Seq}$).

- In equilibrium, we can claim that the follower observes the leader's output, q_1^* , and inserts it into its BRF, to obtain the follower's equilibrium:

$$q_2^{Seq} = \frac{1-c}{2} - \frac{1}{2} \overbrace{\frac{1-c}{2}}^{q_1^{Seq}} = \frac{1-c}{4}$$

- Importantly, *we do not say* that the SPE of this Stackelberg game with two symmetric firms is $(q_1^{Seq}, q_2^{Seq}) = \left(\frac{1-c}{2}, \frac{1-c}{4} \right)$. This output vector only describes firms' output along the equilibrium path. In contrast, the SPE must specify each firm's output decision, both in- and off-the-equilibrium path.

Equilibrium Output level in the Cournot and Stackelberg Games

- Recall that the equilibrium output in the Cournot model of simultaneous quantity competition was $q_i^{Sim} = \frac{1-c}{3}$ for every firm i (see section 4.2.1), where we assumed the same indirect demand function $p(Q) = 1 - Q$ and marginal cost c for both firms.
- Therefore,
 - the leader produces more units than when firms choose their output simultaneously, $q_1^{Seq} = \frac{1-c}{2} > \frac{1-c}{3} = q_1^{Sim}$,
 - while the follower produces fewer units, $q_2^{Seq} = \frac{1-c}{4} < \frac{1-c}{3} = q_2^{Sim}$.
- Intuitively, the leader exercises its “First-mover-advantage” by increasing its output, relative to that under Cournot competition, anticipating that the follower will respond to this increase in q_1 by decreasing its own output q_2 .

Equilibrium Output level in the Cournot and Stackelberg Games

- Figure 6.17 depicts the equilibrium output level in the Cournot and Stackelberg games, showing that the output pair moves from the 45° – *line* (where both firms produce the same output) to above this line, which indicates that firm 1 produces more units than firm 2.
- Graphically, firm 1 anticipates firm 2's BRF, $q_2(q_1)$, and chooses the point along this line that yields the highest profit.




Figure 6.17. Equilibrium quantities in Cournot vs. Stackelberg games.

Equilibrium Profit level in the Cournot and Stackelberg Games

As expected, the leader's profit when firms compete sequentially,

$$\begin{aligned}\pi_1^{Seq} &= [1 - q_1^{Seq} - q_2^{Seq}] q_1^{Seq} - c q_1^{Seq} \\ &= \left[1 - \frac{1-c}{2} - \frac{1-c}{4}\right] \frac{1-c}{2} - c \frac{1-c}{2} \\ &= \frac{(1-c)^2}{8},\end{aligned}$$

is larger than its profit when they compete simultaneously i.e.,

$$\pi_1^{Seq} = \frac{(1-c)^2}{8} > \pi_1^{Sim} = \frac{(1-c)^2}{9}$$

Equilibrium Profit level in the Cournot and Stackelberg Games

In contrast, the follower's profits when firms compete sequentially,

$$\begin{aligned}\pi_2^{Seq} &= [1 - q_1^{Seq} - q_2^{Seq}] q_2^{Seq} - c q_2^{Seq} \\ &= \left[1 - \frac{1-c}{2} - \frac{1-c}{4}\right] \frac{1-c}{4} - c \frac{1-c}{4} \\ &= \frac{(1-c)^2}{16},\end{aligned}$$

is lower than its profits under simultaneous quantity competition, i.e.,

$$\pi_2^{Seq} = \frac{(1-c)^2}{16} < \pi_2^{Sim} = \frac{(1-c)^2}{9}$$

Sequential Public Good Game

- Let us now return to the public good game we examined in section 4.4, but rather than assuming players simultaneously choose their contributions to the public good, we now consider that
 - Player 1 makes her contribution x_1 in the first period, and after observing this contribution, player 2 responds with her contribution x_2 in the second period.
 - Because this is a sequential-move game of complete information, we solve it by backward induction, focusing on player 2 first, and then moving on to player 1.

Sequential Public Good Game

Player 2, follower. Player 2 takes player 1's contribution, x_1 , as given, and chooses her contribution x_2 to solve the utility maximization problem:

$$\max_{x_2} u_2(x_1, x_2) = (w_2 - x_2)\sqrt{m(x_1 + x_2)}$$

Differentiating with respect to x_2 , we obtain

$$-\sqrt{m(x_1 + x_2)} + \frac{m(w_2 - x_2)}{2\sqrt{m(x_1 + x_2)}} = 0$$

which simplifies to $\frac{m(w_2 - 2x_1 - 3x_2)}{2\sqrt{m(x_1 + x_2)}} = 0$ which holds when $3x_2 = w_2 - 2x_1$

Sequential Public Good Game

Solving for x_2 , we find that player 2's BRF is:

$$x_2(x_1) = \begin{cases} \frac{w_2}{3} - \frac{2}{3}x_1 & \text{if } x_1 < \frac{w_2}{2} \\ 0 & \text{otherwise.} \end{cases}$$

This BRF coincides with that found in section 4.4:

- Originating at $x_2 = \frac{w_2}{3}$ when player 1 does not donate to the public good.
- But decreasing at a rate of $2/3$ for every dollar that player 1 contributes.

Sequential Public Good Game

Player 1, leader. In the first period, player 1 anticipates player 2's BRF in the subsequent period, $x_2(x_1)$, and inserts it into her utility maximization problem as follows:

$$\begin{aligned}\max_{x_1 \geq 0} u_1(x_1, x_2) &= (w_1 - x_1) \sqrt{m[x_1 + x_2(x_1)]} \\ &= (w_1 - x_1) \sqrt{m \left[x_1 + \underbrace{\left(\frac{w_2}{3} - \frac{2}{3} x_1 \right)}_{x_2(x_1)} \right]}\end{aligned}$$

which simplifies to

$$\max_{x_1 \geq 0} (w_1 - x_1) \sqrt{m \frac{w_2 + x_1}{3}}$$

Sequential Public Good Game

Player 1, leader. In the first period, player 1 solves

$$\max_{x_1 \geq 0} (w_1 - x_1) \sqrt{m \frac{w_2 + x_1}{3}}$$

Differentiating with respect to x_1 , we obtain

$$\frac{m(w_1 - x_1)}{2\sqrt{m(w_2 + x_1)}} - \frac{\sqrt{m(w_2 + x_1)}}{\sqrt{3}} = 0$$

$$\Rightarrow \frac{m(w_1 - 2w_2 - 3x_1)}{2\sqrt{3}\sqrt{m(w_2 + x_1)}} = 0$$

Then solving for x_1 , we find player 1's equilibrium contribution in this sequential-move version of the public good game

$$x_1^{Seq} = \frac{w_1 - 2w_2}{3}.$$

Sequential Public Good Game

Therefore, the SPE of the game is:

$$SPNE = \left(x_1^{Seq}, x_2(x_1) \right) = \left(\frac{w_1 - 2w_2}{3}, \frac{w_2}{3} - \frac{2}{3}x_1 \right)$$

As a consequence, players contributions evaluated in equilibrium are:

- $x_1^{Seq} = \frac{w_1 - 2w_2}{3}$ for player 1, and
- $x_2^* = x_2(x_1^{Seq})$ for player 2, that is

$$x_2(x_1^{Seq}) = \frac{w_2}{3} - \frac{2}{3} \left(\frac{w_1 - 2w_2}{3} \right) = \frac{7w_2 - 2w_1}{9}$$

Sequential and Simultaneous Public Good Game Comparison

- We find that the leader contributes less to the charity than when players submit donations simultaneously, i.e.,

$$x_1^{Seq} = \frac{w_1 - 2w_2}{3} < \frac{3w_1 - 2w_2}{5} = x_1^{Sim}$$

because this inequality simplifies to $-w_2 < w_1$, which always holds given that wealth levels are positive by assumption.

- The follower, however, contributed more in the sequential- than simultaneous-move game since:

$$x_2^{Seq} = \frac{7w_2 - 2w_1}{9} > \frac{3w_2 - 2w_1}{5} = x_2^{Sim}$$

simplifies to $-w_1 < w_2$, which also holds for all wealth levels.

- Intuitively, the leader anticipates that her decreased contributions will be responded by the follower with an increase in her own donation.
- In other words, the leader exploits her first-move advantage, which in this case means decreasing her contribution to free-ride the follower's donation.

Application: Ultimatum Bargaining Game

- Bargaining is prevalent in many economic situations where two or more parties negotiate how to divide a certain surplus, such as the federal budget or a salary increase.
- We first examine the shortest bargaining setting, the so-called “ultimatum bargaining game” depicted in Figure 6.18a:
 - Player 1 makes a division of the surplus to player 2, d , where $d \in [0,1]$ can be interpreted as a share of the total surplus.
 - Observing the offer d , player 2 only has the ability to accept it or reject it.
- The arc at the top of the game represents the continuum of offers that player 1 can make to player 2, since $d \in [0,1]$.




Figure 6.18a. Ultimatum bargaining game.

Application: Ultimatum Bargaining Game

- If player 2 rejects the offer, both players earn a payoff of zero.
- If player 2 accepts the offer:
 - She receives d
 - Player 1 earns the remaining surplus, $1 - d$.
- *Example.*
 - If player 1 offers 20 percent of the surplus to player 2, $d = 0.2$, and player 2 accepts it, player 1 keeps the remaining 80 percent of the surplus, $1 - d = 0.8$.
- This bargaining game is, then, equivalent to a “take-it-or-leave-it” offer, or an ultimatum, from player 1, explaining the game’s name.

Application: Ultimatum Bargaining Game

- The smallest subgame that we can identify in Figure 6.18b is that initiated after the responder (player 2) observes the proposer's offer, d .
- The second smallest subgame is the game as a whole.




Figure 6.18b. Ultimatum bargaining game - Smallest subgame.

Application: Ultimatum Bargaining Game

- Applying backward induction
 - **Smallest subgame.**
 - Upon receiving an offer d , the responder accepts it if d satisfies $d \geq 0$.
 - **First Stage.**
 - Proposer anticipates the responder's decision rule, $d \geq 0$, and makes an offer that maximizes her payoff conditional on that being accepted, i.e., proposer solves
$$\max_{d \geq 0} 1 - d$$
which considers the constraint $d \geq 0$ to induce the responder to accept the offer.
 - Differentiating with respect to d , yields -1 (corner solution).
 - Intuitively, the proposer seeks to reduce d as much as possible (Figure with d on the horizontal axis).
- Therefore, the proposer reduces d all the way to $d^* = 0$, which still satisfies acceptance, making the responder indifferent between accepting and rejecting the offer.

Application: Ultimatum Bargaining Game

- For simplicity, we assume that the responder accepts offers that make her indifferent.
 - If, instead, she rejects this type of offers, the proposer could offer her an extremely small division of the surplus, $d \rightarrow 0$, still yielding similar equilibrium results as above.
- Therefore, the SPE of the ultimatum game is:


$$SPE = \{d^* = 0, d \geq 0\}$$

which indicates that the proposer makes an offer $d^* = 0$, and the responder accepts any offer d that satisfies $d \geq 0$.

- Equilibrium path:
 - The proposer offers $d^* = 0$
 - The responder accepts that offer, yielding equilibrium payoff $(1 - d, d) = (1, 0)$ which implies that the proposer keeps all the surplus.
- Recall that this “equilibrium path” doesn’t describe how the responder reacts to an off-the-equilibrium offer from the proposer, that is, $d \neq 0$.

Application: Two-period alternating offers bargaining game

- Figure 6.19 extends the above bargaining:
 - allowing the responder (player 2) to make a counteroffer if she rejects the division that player 1 offers, d^1 .




Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

Time structure:

- Player 1 makes an offer d^1 to player 2, who accepts or rejects it.
- If player 2 accepts d^1 , the game is over, and player 2 earns d^1 .
 - Player 1 earns the remaining surplus, $1 - d^1$.
- If player 2 rejects, however, she can make an offer d^2 to player 1 (which we can interpret as a “counteroffer”).
- Observing d^2 , player 1 responds accepting or rejecting it.




Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

Time structure:

- As in previous stages, both players earn zero if the game ends without a division being accepted.
 - Otherwise, player 1 (the responder in the second stage) earns d^2 , which has a discounted value of $\delta_1 d^2$ in today's terms.
- Discount factor $\delta_1 \in [0,1]$ indicates the relative importance that player 1 assigns to future payoffs.




Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

Time structure:

- Player 2 earns, in this case, the remaining surplus $1 - d^2$, whose discounted value in today's terms is $\delta_2(1 - d^2)$.
- In this context, we can find four subgames, those initiated after:
 1. Player 1 receives offer d^2
 2. Player 2 rejects offer d^1
 3. Player 2 receives offer d^1 , and
 4. The game as a whole



Matrix 6.19. Two-period alternating offers bargaining game

Application: Two-period alternating offers bargaining game

By backward induction, we start with the smallest subgame (subgame (1)).

- **Subgame (1).** At that point of the game, player 1 accepts any offer d^2 from player 2 if and only if $d^2 \geq 0$.
- **Subgame (2).** In subgame (2), player 2 anticipates this best response by player 1, making an offer d^2 that solves

$$\max_{d^2 \geq 0} 1 - d^2$$

which is analogous to the proposer's problem in the ultimatum bargaining game of section 6.4.3.

- Intuitively, when players reject offers and reach their final period of interaction, they face a strategic setting like the ultimatum bargaining game:
 - Not reaching an agreement at that point will lead to a zero payoff for everyone, as in the ultimatum bargaining game.
- The proposer (player 2) offers the lowest division that guarantees acceptance from player 1, $d^2 = 0$, earning a payoff of $1 - d^2 = 1 - 0 = 1$.

Application: Two-period alternating offers bargaining game

- **Subgame (3).** Player 2 chooses whether to:
 - Accept the offer from player 1, earning d^1 , or
 - Reject it and become the proposer tomorrow.
 - In this setting, she anticipates that her equilibrium payoff in subgame (2) will be 1.
 - Discounted value of 1 today is $\delta_2 \times 1 = \delta_2$.
- Therefore, player 2 accepts offer d^1 if and only if $d^1 \geq \delta_2$.

Application: Two-period alternating offers bargaining game

- **Game as a whole.** Player 1 maximizes

$$\max_{d^1 \geq \delta_2} 1 - d^1$$

which yields an equilibrium offer $d^1 = \delta_2$. Therefore, the SPE of the two-period alternating offers bargaining game is

$$SPE = \left\{ \underbrace{(d^1 = \delta_2, d^2 \geq 0)}_{\text{Player 1}}, \underbrace{(d^2 = 0, d^1 \geq \delta_2)}_{\text{Player 2}} \right\}$$

- Player 1 offers a division $d^1 = \delta_2$ in the first period, and accepts any offer $d^1 \geq 0$ in the second period.
- Player 2 offers a division $d^2 = 0$ in the first period, and accepts any offer $d^1 \geq \delta_2$ in the second period.

Application: Two-period alternating offers bargaining game

- **Game as a whole.**

- This SPE implies that, while players have two periods to negotiate how to divide the surplus, they reach an agreement in the first period (instantaneously!):
 - Player 1 offers a division $d^1 = \delta_2$, which player 2 accepts since it satisfies $d^2 \geq 0$, and the game is over (see left side of Figure 6.19).
- In equilibrium,
 - Player 1 earns a payoff $1 - d^1 = 1 - \delta_2$.
 - Player 2 earns a payoff equal to the division she accepted from player 1, $d^1 = \delta_2$.
 - In summary, equilibrium payoffs are $(u_1^*, u_2^*) = (1 - \delta_2, \delta_2)$

Application: Two-period alternating offers bargaining game

Comparative Statics of equilibrium payoffs.

- As player 2 becomes more patient (higher δ_2):
 - Her equilibrium payoff increases while that of player 1 decreases.
 - Player 1 understands that player 2 does not discount future payoffs significantly, meaning that she can reject player 1's offer today, becoming the proposer tomorrow.
- If instead $\delta_2 \rightarrow 0$:
 - Player 1 anticipates that player 2 severely discounts future payoffs, having a stronger preference for today's payoffs.
 - In this case, player 1 can exploit player 2's impatience by offering her a lower division.
- Interestingly, player 1's patience, as captured by his discount factor δ_1 , does not affect:
 - The offer she makes to player 2 in equilibrium, $d^2 = 0$, or
 - Players' equilibrium payoffs, $(u_1^*, u_2^*) = (1 - \delta_2, \delta_2)$, which are only affected by player 2's patience, δ_2 .

Some tricks about solving alternating offer bargaining games

- Equilibrium payoffs in the last stage of the game are $(0,1)$ where the proposer in that stage makes a zero offer, which is accepted, helping the proposer keep the whole surplus.
 - This holds true when the game has only one period of possible negotiations (ultimatum bargaining game), two periods, or more.
- Players then anticipate that, if the game reaches the last period, their equilibrium payoffs will be $(0,1)$ if player 2 is the proposer in that period or $(1,0)$ if player 1 is the proposer.
 - (WLOG, assume that player 2 is the proposer in the last period).
- The proposer in the previous-to-last period (player 1), must then offer a division d^1 that makes player 2 indifferent between:
 - Accepting d^1 today or...
 - Rejecting it to become the proposer tomorrow and earn a payoff of 1, with discounted value $\delta_2 \times 1 = \delta_2$ today.
 - This means that player 1 offers exactly $d^1 = \delta_2$, keeping the remaining surplus $1 - \delta_2$ in equilibrium.

Some tricks about solving alternating offer bargaining games

Applying in longer alternating offer bargaining games

- For instance, if players can negotiate for four periods, starting with player 1 making an offer d^1 , we can operate by backward induction as follows:
 1. *Fourth period.* Player 2 is the proposer, so the equilibrium payoff in that subgame is $(0,1)$.
 2. *Third period.* Player 1 makes an offer, so she must make player 2 indifferent between her offer today or the whole surplus tomorrow, worth δ_2 today. Equilibrium payoffs are, then,
 - $(1 - \delta_2, \delta_2)$.
 3. *Second period.* Player 2 is the proposer, so she must make player 1 indifferent between her offer today and payoff $1 - \delta_2$ tomorrow (when she becomes the proposer), with discounted value $\delta_1(1 - \delta_2)$ today. Therefore, equilibrium payoffs are
$$(\delta_1(1 - \delta_2), 1 - \delta_1(1 - \delta_2))$$

where $1 - \delta_1(1 - \delta_2)$ represents the remaining surplus that player 2 earns, after offering $\delta_1(1 - \delta_2)$ to player 1.

Some tricks about solving alternating offer bargaining games

4. *First period.* Player 1 is the proposer, and she must make player 2 indifferent between her current offer and payoff $1 - \delta_1 (1 - \delta_2)$ tomorrow, with discounted value $\delta_2 [1 - \delta_1 (1 - \delta_2)]$ today. As a consequence, equilibrium payoffs are

$$(1 - \delta_2 [1 - \delta_1 (1 - \delta_2)], \delta_2 [1 - \delta_1 (1 - \delta_2)])$$

where player 1's equilibrium payoff denotes the remaining surplus after she offers $\delta_2 [1 - \delta_1 (1 - \delta_2)]$ to player 2.

Some tricks about solving alternating offer bargaining games

- Table 6.1 summarizes our approach to solve this alternating-offer bargaining game, starting in the last period.
- Vertical arrows indicate that a player must be indifferent between her payoff in the period when she is a proposer and her payoff in the previous period when she is the responder.
- Horizontal arrows, however, denote that the proposing player's equilibrium payoff is the remaining surplus, i.e., the share not offered to her rival.
- Graphically, we move up and sideways, and then up and sideways again, until reaching the first period, arrows resemble a ladder, explaining why our students often refer to this approach as the “ladder”.

<i>Player</i>	<i>Period</i>	<i>Player 1's payoff</i>	<i>Player 2's payoff</i>
Player 1	$t=1$	$1-\delta_2[1-\delta_1(1-\delta_2)] \leftarrow$	$\delta_2[1-\delta_1(1-\delta_2)] \uparrow$
Player 2	$t=2$	$\delta_1(1-\delta_2) \rightarrow$	$1-\delta_1(1-\delta_2)$
Player 1	$t=3$	$1-\delta_2 \leftarrow$	$\delta_2 \uparrow$
Player 2	$t=4$	0	1

Table 6.1. An approach to solve alternating-offers bargaining games.

Alternating-offer bargaining game with infinite periods

- Consider an infinite-period bargaining game where player 1 makes offers to player 2, d^1 , in odd numbered periods while player 2 makes to player 1, d^2 , in even numbered periods
- We seek to find “Stationary strategies,” meaning:
 - Player 1:
 - Makes the same offer to player 2 in every period when player 1 is the proposer (every odd-numbered period), and
 - Follows the same decision rule when he becomes the responder (in every even-numbered period).
 - Player 2:
 - Makes the same offer to player 1 in every period when player 2 is the proposer (every even-numbered period), and
 - Follows the same decision rule when he becomes the responder (in every odd-numbered period).

Alternating-offer bargaining game with infinite periods

Odd-numbered periods:

- Player 2 compares the payoff she gets from accepting the offer from player 1, d^1 , against the payoff that she can earn tomorrow when she becomes the proposer.
- Tomorrow, player 2 offers d^2 , keeping the remaining $1 - d^2$ of the surplus, with discounted value $\delta_2(1 - d^2)$ in today's terms.
- Player 1 accepts player 1's offer today, d^1 , if and only if
$$d^1 \geq \delta_2(1 - d^2).$$
- Player 1 anticipates this decision rule and minimizes her offer, d^1 , making player 2 indifferent: $d^1 = \delta_2(1 - d^2)$.

Alternating-offer bargaining game with infinite periods

Even-numbered periods:

- Player 1 compares her payoff from accepting player 2's offer, d^2 , and the payoff she can earn tomorrow when she becomes the proposer.
- Player 1 offers d^1 to player 2 tomorrow, keeping the remaining $1 - d^1$ of the surplus, with discounted value $\delta_1(1 - d^1)$.
- Player 1 accepts d^2 today if and only if $d^2 \geq \delta_1(1 - d^1)$, where player 2 reduces d^2 , ultimately making player 1 indifferent between accepting and rejecting, that is,

$$d^2 = \delta_1(1 - d^1).$$

Alternating-offer bargaining game with infinite periods

Solving the game.

- We have one equation each from odd and even-numbered periods with two unknowns d^1 and d^2 .
- Solve for these two offers in equilibrium, we can insert one indifferent condition into the other yielding

$$d^1 = \delta_2 \left(1 - \underbrace{\delta_1(1 - d^1)}_{d^2} \right)$$

which rearranges to $d^1 = \delta_2 - \delta_1\delta_2 + \delta_1\delta_2 d^1$, and solving for d^1 , yields

$$d^1 = \frac{\delta_2(1 - \delta_1)}{1 - \delta_1\delta_2}$$

Alternating-offer bargaining game with infinite periods

- Inserting this result into indifference condition $d^2 = \delta_1(1 - d^1)$, we obtain

$$d^2 = \delta_1 \left(1 - \frac{\delta_2(1 - \delta_1)}{1 - \delta_1\delta_2} \right) = \frac{\delta_1(1 - \delta_2)}{1 - \delta_1\delta_2}$$

- We can then conclude that, in the first period, player 1 makes an offer d^1 to player 2, who immediately accepts it, and the game is over, yielding equilibrium payoffs

$$(1 - d^1, d^1) = \left(\frac{1 - \delta_2}{1 - \delta_1\delta_2}, \frac{\delta_2(1 - \delta_1)}{1 - \delta_1\delta_2} \right)$$

where $1 - d^1 = 1 - \frac{\delta_2(1 - \delta_1)}{1 - \delta_1\delta_2} = \frac{1 - \delta_2}{1 - \delta_1\delta_2}$.

Alternating-offer bargaining game with infinite periods

- Therefore, player 2's equilibrium payoff, d^1 , increases in her own discount factor δ_2 , since

$$\frac{\partial d^1}{\partial \delta_2} = \frac{1 - \delta_1}{(1 - \delta_1 \delta_2)^2} \geq 0$$

Implying that, as player 2 assigns more weight to her future payoff, she can reject player 1's offer and wait to become the proposer, which forces player 1 to make a more generous offer today.

- In contrast, player 2's equilibrium payoff decreases in her rival's discount factor, δ_1 , because

$$\frac{\partial d^1}{\partial \delta_1} = \frac{\delta_2(\delta_2 - 1)}{(1 - \delta_1 \delta_2)^2} < 0$$

which intuitively means that, as player 1, assigns more weight to her future payoffs, player 2 must offer her a larger surplus share in future periods (when player 2 becomes the proposer). This, in turn, reduces player 2's future payoffs, making her less attracted to reject the offer from player 1 today.

Appendix:

Mixed and Behavioral Strategies

Appendix – Mixed and Behavioral Strategies

- Definition. **Mixed Strategy.** In a sequential-move game, a mixed strategy is a probability distribution over all player i 's pure strategies $s_i \in S_i$.
- In Figure 6.20, the set of pure strategies is $S_1 = \{Ac, Ad, Bc, Bd\}$, and a mixed strategy is a randomization over all these pure strategies.




Figure 6.20. Mixed and behavioral strategies in a game tree.

Appendix – Mixed and Behavioral Strategies

- Intuitively, player i rolls a dice, before the game starts, and the outcome of this roll determines the path of play that she will follow (i.e., the pure strategy that she chooses).
- For instance, a mixed strategy for player 1 could be $\left(0Ac, 0Ad, \frac{1}{2}Bc, \frac{1}{2}Bd\right)$,
- which puts:
 - no probability weight on pure strategies Ac and Ad , but
 - assigns 50 percent probability on Bc and Bd .




Figure 6.20. Mixed and behavioral strategies in a game tree.

Appendix – Mixed and Behavioral Strategies

- Definition. **Behavioral strategy.**
- A behavioral strategy is a mapping $b_i: H_i \rightarrow \Delta A_i(h_i)$, where $b_i(s_i(h_i))$ is the probability that player i selects action $s_i(h_i)$ at information set h_i .
- In the context of Figure 6.20, an example of behavioral strategy for player 1 could specify that:
 - At the first node where she is called to move, she randomizes between the two available actions in this node, A and B , such as $pA + (1 - p)B$ where $p \in [0,1]$.
 - Similarly, in the second node where she is called to move, she randomizes between actions c and d , e.g., $qc + (1 - q)d$, where $q \in [0,1]$.
- In that context, an example of a behavioral strategy for player 1 is $(p, q) = \left(\frac{1}{3}, \frac{1}{4}\right)$, meaning that:
 - She assigns a probability of $\frac{1}{3}$ to action A in the first node where she is called to move, and
 - Probability $\frac{1}{4}$ to action c in the second node where she plays.

Appendix – Mixed and Behavioral Strategies

- **Equivalence.**

- Are behavioral strategies equivalent to a corresponding mixed strategy?
- Yes, we can construct a one-to-one correspondence between behavioral and mixed strategy.

- To see this with the above example, if probabilities p and q in the behavioral strategy satisfy $p = 0$ and $q = \frac{1}{2}$, then the behavioral strategy becomes

$$\left(B, \frac{1}{2}c + \frac{1}{2}d \right)$$

is equivalent to the mixed strategy

$$\left(0Ac, 0Ad, \frac{1}{2}Bc, \frac{1}{2}Bd \right)$$

Since, in one of them, player 1 starts choosing B in pure strategies, and then randomizes between c and d with equal probability.

- Otherwise, the behavioral and mixed strategy do not produce the same outcome.

Appendix – Mixed and Behavioral Strategies

- **“Manuals and Libraries.”**

- Following Luce and Raiffa’s (1957) analogy, a pure strategy $s_i \in S_i$ can be understood as an “instructional manual” in which each page tells player i which action to choose when she is called on to move at information set h_i .
 - Every manual has as many pages as information sets player i can face.
- In this analogy, strategy space S_i is a “library” with all possible instruction manuals.
- A mixed strategy, then, randomly chooses a specific manual s_i from the library S_i (i.e., a specific complete contingent plan), which the player considers for the rest of the game.
- In contrast, a behavioral strategy chooses pages from different manuals (actions in each information set) with positive probability.

Chapter 7: Repeated Games

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- Situations in which players interact in one of the strategic setting and they play the game for several rounds, which are known as “repeated games”
- Interested in identifying if the game’s repetition provides players with more incentives to cooperate

Prisoner’s Dilemma

- We first study the finitely-repeated version of this game
- Extend players’ interaction to an infinitely repeated game
- We then apply the above tools to different settings, such as collusion in oligopoly models, where every firm chooses its output level in each period of interaction, and to stage games with more than one NE

Repeating the Game twice

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.1 The Prisoner's Dilemma Game

- Let us repeat the game twice and find its SPE since players now interact in a sequential-move game:
 - In the first stage, every player i simultaneously and independently chooses whether to *Confess* or *Not Confess*
 - In the second stage, observing the outcome of the first stage, every player i selects again simultaneously and independently whether to *Confess* or *Not Confess*
 - This sequential-game is depicted in Figure 7.1 (next slide).

Twice-repeated Prisoner's Dilemma Game

- The game tree includes information sets in the first and second stage.
- At the beginning of second stage, however, players perfectly observe the strategy profile played in the first stage (no information sets there).




Figure 7.1. Twice-repeated prisoner's dilemma game.

Twice-repeated Prisoner's Dilemma Game

- Payoff pairs in the terminal nodes are just a sum of payoffs in the first and second stage, assuming no discounting.
- For instance, if (C, C) is played in the first stage, but (NC, C) occurs in the second stage:
 - Player 1 earns $2 + 0 = 2$, while
 - Player 2 earns $2 + 8 = 10$.
- A similar argument applies to other payoffs at the bottom of the tree. (Practice!)




Figure 7.1. Twice-repeated prisoner's dilemma game.

Twice-repeated Prisoner's Dilemma Game

- There are five subgames:
 1. One initiated after players choose (C, C) in the first stage.
 2. Another initiated after (NC, C)
 3. Another initiated after (C, NC)
 4. Another initiated after (NC, NC)
 5. And the game as a whole
- Operating by backward induction, we solve each of these subgames, starting with subgames 1-4.
- Since these subgames are simultaneous-move games, we can solve each of them by transforming each to its matrix form.




Figure 7.1. Twice-repeated prisoner's dilemma game.

Twice-repeated Prisoner's Dilemma Game - Subgames for Second Stage

		Player 2	
		Confess	Not Confess
Player 1	Confess	<u>4</u> , <u>4</u>	<u>10</u> ,2
	Not Confess	2, <u>10</u>	6,6

Matrix 7.2a Twice-repeated Prisoner's Dilemma game – Subgame 1

		Player 2	
		Confess	Not Confess
Player 1	Confess	<u>10</u> , <u>2</u>	<u>16</u> ,0
	Not Confess	8,8	12,4

Matrix 7.2b Twice-repeated Prisoner's Dilemma game – Subgame 2

		Player 2	
		Confess	Not Confess
Player 1	Confess	<u>2</u> , <u>10</u>	8,8
	Not Confess	0, <u>16</u>	4,12

Matrix 7.2c Twice-repeated Prisoner's Dilemma game – Subgame 3

		Player 2	
		Confess	Not Confess
Player 1	Confess	<u>6</u> , <u>6</u>	<u>12</u> ,4
	Not Confess	4, <u>12</u>	8,8

Matrix 7.2d Twice-repeated Prisoner's Dilemma game – Subgame 4

Twice-repeated Prisoner's Dilemma Game – First Stage

• Summary

- Every player $i = \{1,2\}$ chooses to play *Confess* in the first stage, and
- Players *Confess* in the second stage regardless of the strategy profile played in the first stage




Figure 7.2. Twice-repeated prisoner's dilemma game - First stage.

		Player 2	
		Confess	Not Confess
Player 1	Confess	<u>4,4</u>	<u>10,2</u>
	Not Confess	<u>2,10</u>	6,6

Matrix 7.2e Twice-repeated Prisoner's Dilemma game – Game as whole

Summary

- Intuitively, players anticipate that they will both play *Confess* in the last stage independently of how cooperative or uncooperative they were in the previous stage
- Informally, every player thinks:
 - “It does not matter what I do today, tomorrow we will all confess.
 - I can then treat tomorrow’s game as independent from today’s play, as if I played two separate Prisoner’s Dilemma games.”
- In summary, repeating the game twice did not help us sustain cooperation in the SPE.
- Can cooperation emerge if we repeat the game for more periods, generally, for $T \geq 2$ periods?

Repeating the game $T \geq 2$ times

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2 + P_1^{T-1}, 2 + P_2^{T-1}}$	$\underline{8 + P_1^{T-1}, 0 + P_2^{T-1}}$
	Not Confess	$\underline{0 + P_1^{T-1}, 8 + P_2^{T-1}}$	$\underline{4 + P_1^{T-1}, 4 + P_2^{T-1}}$

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

- Matrix 7.3 represents the game that players face when interacting in the last round (period T)
- Notation:
 - P_1^{T-1} denotes the discounted sum of all payoffs that player 1 earned in the previous $T - 1$ periods.
 - P_2^{T-1} is a similar expression for player 2.
- This presentation helps us describe the subgame that players face at period T for any previous history of play.
 - That is, for any strategy profiles occurring in previous periods.
 - *Intuitively*: regardless of how we played the game in previous periods.

Repeating the game $T \geq 2$ times

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2 + P_1^{T-1}, 2 + P_2^{T-1}}$	$\underline{8 + P_1^{T-1}, 0 + P_2^{T-1}}$
	Not Confess	$\underline{0 + P_1^{T-1}, 8 + P_2^{T-1}}$	$\underline{4 + P_1^{T-1}, 4 + P_2^{T-1}}$

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

- To illustrate the generality of this matrix representation, evaluate terms P_1^{T-1} and P_2^{T-1} in the special case where the game is only repeated $T = 2$ times.
- In this context:
 - P_1^{T-1} captures the payoff that player 1 earned in the first period of interaction.
 - And similarly, with P_2^{T-1} for player 2.
- For instance, if (C, C) emerged in the first stage of the game,
 - $P_1^{T-1} = P_2^{T-1} = 2$.
 - While if (C, NC) occurred, then $P_1^{T-1} = 8$ and $P_2^{T-1} = 0$.

Repeating the game $T \geq 2$ times

- If the Prisoner's Dilemma can produce four different strategy profiles in each period, there are 4^{T-1} nodes at the beginning of period T .
- In the twice-repeated Prisoner's Dilemma game, for instance, there are $4^{2-1} = 4$ nodes at the beginning of the second period, as shown in the previous tree.
- More generally, if the unrepeated version of the game has
 - $k > 2$, different strategy profiles (cells) and
 - the game is repeated T times,
- Then, there are:
 - k^{T-1} nodes at the beginning of period T ,
 - only k^{T-2} nodes at the beginning of period $T - 1$,
 - and similarly in previous periods...
 - with just $k^{T-T} = k^0 = 1$ node at the beginning of the game (first period of interaction).

Iterative Method: Period T

		Player 2	
		Confess	Not Confess
Player 1	Confess	<u>$2 + P_1^{T-1}$</u> , <u>$2 + P_2^{T-1}$</u>	<u>$8 + P_1^{T-1}$</u> , $0 + P_2^{T-1}$
	Not Confess	$0 + P_1^{T-1}$, <u>$8 + P_2^{T-1}$</u>	$4 + P_1^{T-1}$, <u>$4 + P_2^{T-1}$</u>

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

Period T :

- Underlining best response payoffs, we can see that (C, C) is a NE in the T -period subgame in Matrix 7.3.
- Player 1's best response:
 - When Player 2 chooses *Confess*, player 1's best response is to *Confess* because $2 + P_1^{T-1} > 0 + P_1^{T-1}$ simplifies to $2 > 0$.
 - When Player 2 selects *Not Confess*, player 1's best response is *Confess* since $8 + P_1^{T-1} > 4 + P_1^{T-1}$ simplifies to $8 > 4$.

Iterative Method: Period T

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2 + P_1^{T-1}}, \underline{2 + P_2^{T-1}}$	$\underline{8 + P_1^{T-1}}, 0 + P_2^{T-1}$
	Not Confess	$0 + P_1^{T-1}, \underline{8 + P_2^{T-1}}$	$4 + P_1^{T-1}, 4 + P_2^{T-1}$

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

Period T :

- Note that:
 - Players' best responses in period T don't depend on players' payoffs in previous periods, as captured by P_1^{T-1} and P_2^{T-1} .
 - As if they were sunk, or they couldn't be changed at this point.
 - In other words, player 1 treats P_1^{T-1} as a constant; while player 2 treats P_2^{T-1} as a constant.
 - This implies that every player finds:
 - *Confess* to be a best response to her rival's strategy *regardless* of the previous history of play.

Iterative Method: Period $T - 1$

- **Period $T - 1$.**

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2 + P_1^{T-2}}, \underline{2 + P_2^{T-2}}$	$\underline{8 + P_1^{T-2}}, 0 + P_2^{T-2}$
	Not Confess	$0 + P_1^{T-2}, \underline{8 + P_2^{T-2}}$	$4 + P_1^{T-2}, 4 + P_2^{T-2}$

Matrix 7.4. Finitely-repeated Prisoner's Dilemma Game – Subgames at period $T - 1$

- (C, C) is again the unique NE.
- The intuition is identical to that in the twice-repeated game:
 - Players anticipate that everyone will play *Confess* in the last round of interaction (period T).
 - Therefore, players have no incentives to cooperate today (in period $T - 1$) since misbehavior will not be disciplined in the following period.

Iterative Method: Period $T - 2$

- **Period $T - 1$.**

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2 + P_1^{T-3}}, \underline{2 + P_2^{T-3}}$	$\underline{8 + P_1^{T-3}}, 0 + P_2^{T-3}$
	Not Confess	$0 + P_1^{T-3}, \underline{8 + P_2^{T-3}}$	$4 + P_1^{T-3}, 4 + P_2^{T-3}$

Matrix 7.4. Finitely-repeated Prisoner's Dilemma Game – Subgames at period $T - 1$

- The same argument follows for period $T - 2$, as illustrated in the above matrix.
- Same argument applies to period $T - 3$, and so on.
- Extending that reasoning to all previous periods, we find that (C, C) is the unique NE:
 - in every subgame and
 - in every period.

Repeating the Game infinitely many times

- We could not support cooperation by repeating the game twice or $T \geq 2$ times...
 - But our above results point us to why uncooperative outcomes emerge:
 - The presence of a terminal period when players know that the game ends
- At that last stage, both players:
 - Behave as in the unrepeated version of the game,
 - Regardless of how the game was played in previous rounds.
- Using backward induction, they extend this behavior to all previous interactions.

Repeating the Game infinitely many times

- Therefore, if we seek to sustain cooperation:
 - Players cannot know with certainty when the game will end.
 - They interact today, at time t , and will keep playing with each other tomorrow, at time $t+1$, with a probability p , where $p \in [0, 1]$.
 - For instance, if $p = 0.7$, the chances of interacting for more than 20 rounds are less than 1 percent.
 - Probability of interacting converges fast to zero, but it is always strictly positive.

Uncooperative outcome

- We first show that (C,C) in every period can still be sustained as one of the SPEs in the infinitely-repeated game.
- Note that if a player chooses C at every period t , anticipating that his rival will choose C as well, he obtains a sequence of 2, with discounted value

$$2 + 2\delta + 2\delta^2 + \dots = 2[1 + \delta + \delta^2 + \dots] = 2 \frac{1}{1 - \delta}$$

Recall that

$$1 + \delta + \delta^2 + \dots = \delta^0 + \delta^1 + \delta^2 + \dots = \sum_{t=0}^{\infty} \delta^t = \frac{1}{1 - \delta}$$

Uncooperative outcome (C, C)

- If, instead, she unilaterally deviates to *Not Confess* for one period, he earns

$$\underbrace{0}_{\text{Not Confess}} + \underbrace{2\delta + 2\delta^2 + \dots}_{\text{Confess}} = 0 + 2\frac{\delta}{1-\delta}$$

Hence, this player earns a higher stream of payoffs choosing *Confess* than *Not Confess* because

$$2\frac{1}{1-\delta} \geq 2\frac{\delta}{1-\delta}$$

simplifies to $1 \geq \delta$ which holds given that $\delta \in (0,1)$ by assumption.

Cooperative Outcome: Grim-Trigger Strategy

- Can we sustain cooperation as the SPE of the infinitely-repeated game?
- For that, we need players to:
 - Cooperate and keep cooperating as long as every player cooperated in previous period, but...
 - Otherwise, threat to move the game to a “grim” outcome where players earn lower payoffs than by cooperating.
- The following Grim-Trigger Strategy (GTS) helps us achieve exactly that:
 1. In period $t = 1$, choose *Not Confess*
 2. In every period $t \geq 2$,
 - a. Keep choosing *Not Confess* if every player chose *Not Confess* in all previous periods, or
 - b. Choose *Confess* thereafter for any other history of play (i.e., if either player chose Confess in any period)

Cooperative Outcome: Grim-Trigger Strategy

- The name of the GTS strategy should be clear at this point:
 - Any deviation triggers a grim punishment thereafter, where every player reverts to the NE of the unrepeated version of the game (often known as “Nash reversion”).
- To show that the GTS is a SPE of the infinitely-repeated game, we need to demonstrate that it is an optimal strategy:
 - for *every* player, and
 - at *every* subgame where they are called to move.
- This means that players find the GTS to be optimal:
 - at *any* point t , and
 - after *any* previous history of play.
- A daunting task? No, we only need to examine two histories of play.

Case (1). No Cheating History

- At any period t (both $t = 1$ and $t \geq 2$), if no previous history of cheating occurs, the GTS prescribes that:
 - Every player keeps cooperating, which yields a payoff of 4 for every player i , entailing a payoff stream of

$$4 + 4\delta + 4\delta^2 + \dots = 4 \frac{1}{1 - \delta}$$

Case (1). No Cheating History

- If, instead, player i chooses *Confess* today (which we interpret as her deviating from the above GTS, or more informally as “cheating”), her current payoff increases from 4 to 8.
 - This is a unilateral deviation as only player i chooses *Confess* while her opponent still plays according to the GTS.
- Therefore, player i 's discounted payoff stream from cheating is

$$\underbrace{8}_{\text{Deviation to Confess}} + \underbrace{2\delta + 2\delta^2 + \dots}_{\text{Punishment thereafter}} = 8 + 2\frac{\delta}{1 - \delta}$$

Case (1). No Cheating History

- **Comparison:**
- Then, after a history with no previous cheating episodes, every player chooses to cooperate (player *Not Confess*) if

$$4 \frac{1}{1 - \delta} \geq 8 + 2 \frac{\delta}{1 - \delta}$$

- Multiplying both sides by $1 - \delta$, yields

$$4 \geq 8(1 - \delta) + 2\delta$$

or $4 \geq 8 - 6\delta$, which solving for δ yields

$$\delta \geq \frac{2}{3}$$

Case (2). Some Cheating History

- At period t , if one or both players cheated in previous periods, then the GTS dictates that:
 - Every player responds with *Confess* thereafter, earning a discounted stream of payoffs
- Player i could, instead, unilaterally deviate to *Not Confess*, while his rival plays *Confess* as part of the punishment in the GTS.
 - You may suspect that such deviation is not profitable:
 - Player i is the only one choosing *Not Confess*, decreasing his payoff in that period, rather than increasing it.
 - And his deviation does not change his rival's behavior in subsequent stages.
 - That suspicion is correct: Player i 's discounted stream of payoffs from this deviation is

$$2 + 2\delta + 2\delta^2 = 2 \frac{1}{1 - \delta}$$

$$\underbrace{0}_{\text{Deviation to Not Confess}} + \underbrace{2\delta + 2\delta^2 + \dots}_{\text{Punishment thereafter}}$$

Case (2). Some Cheating History

- In period t , players observe (NC, C) or (C, NC) being played, implying that a deviation from the fully cooperative outcome (NC, NC) occurred, triggering an infinite punishment of (C, C) thereafter, with a payoff of 2 to every player. This payoff stream simplifies to

$$2(\delta + \delta^2 + \dots) = 2 \frac{\delta}{1 - \delta}$$

- **Comparison:** After a history of cheating, every player i prefers to implement the punishment prescribed by the GTS if

$$2 \frac{1}{1 - \delta} \geq 2 \frac{\delta}{1 - \delta} \Rightarrow 2 \geq 2\delta$$

which holds for $\delta \in (0,1)$

Summary

- Combining our results from cases (1) and (2), we only found one condition restricting the value of δ for us to sustain the GTS as a SPE of the infinitely-repeated Prisoner's Dilemma game:

$$\delta \geq \frac{2}{3}$$

- Intuitively:
 - When players assign a sufficiently high weight to future payoffs...
 - They start cooperating in the first period and keep cooperating in all subsequent periods, yielding outcome (C, C) in every round of interaction.

Summary

- Figure 7.3 illustrates the trade-off that every player experiences when deciding whether to cooperate, playing *Not Confess*, after a history of cooperation.
- If she sticks to the GTS, she earns a payoff of 4 thereafter, as depicted in the flat dashed-line in the middle of the figure.




Figure 7.3. Incentives from cheating in the infinitely-repeated prisoner's dilemma game.

Summary

- If, instead, she cheats:
 - Her current payoff increases from 4 to 8 during one period.
 - But her rival detects her cheating, triggering a punishment with (C, C) thereafter, with associated payoff 2.
- Relative to what she earns by cooperating (4), cheating provides:
 - An instantaneous gain
 - But a future loss due to the punishment.




Figure 7.3. Incentives from cheating in the infinitely-repeated prisoner's dilemma game.

Extensions: Temporary Punishments

- The GTS in the above example assumes an infinite reversion to the NE of the stage game (unrepeated version of the game)
- But, what if:
 - players only revert to this NE during a finite number of periods, N ,
 - moving back to cooperation once every player observes that both players implemented the punishment during N periods?
- Graphically, the right rectangle in the figure would be narrower.
- In summary, shortening the punishment, while keeping its severity, shrinks the parameter values where we can sustain cooperation.




Figure 7.4. Temporary punishments.

Extensions: More Attractive Cheating

- If every player i earns 10 rather than 8 when she chooses *Confess* while her opponent plays *Not Confess*...
 - the height of the instantaneous gain from cheating in Figure 7.5 increase, thus making cheating more attractive.
- Formally, this increases the minimal discount factor sustaining cooperation, $\underline{\delta}$,
 - where the GTS can be sustained as a SPE if $\delta \geq \underline{\delta}$.




Figure 7.5. More attractive cheating.

Extensions: More Attractive Cheating

- Because the range of δ 's sustaining cooperation satisfy $\delta \in [\underline{\delta}, 1)$, an increase in cutoff $\underline{\delta}$ entails that cooperation emerges in a more restricted range of δ 's.
- For compactness, the literature says that:
 - an increase in cutoff $\underline{\delta}$ “hinders cooperation in equilibrium”,
 - while a decrease in $\underline{\delta}$ facilitates such cooperation.




Figure 7.5. More attractive cheating.

Extensions: More Severe Punishments

- If every player earns 0 at the NE of the game, rather than 2:
 - the right rectangle in Figure 7.3 becomes deeper, as depicted in Figure 7.6,
 - indicating a more severe future payoff loss from cheating today.
- Intuitively:
 - cheating becomes less attractive,
 - which decreases the minimal discount factor sustaining cooperation, $\underline{\delta}$,
 - implying that the GTS can be sustained under a larger range of δ 's, i.e., $\delta \in [\underline{\delta}, 1)$




Figure 7.6. More severe punishment.

Extensions: Lag in Detecting Cheating

- However, in some real-life examples, players may detect cheating $k \in \mathbb{Z}_+$ periods after it happened.
- If:
 - $k = 0$, we would still have immediate detection,
 - while $k > 0$ entails a lag in detecting cheating episodes.




Figure 7.7. Lag at detecting cheating.

Extensions: Lag in Detecting Cheating

- This lag widens the left square in Figure 7.7 representing the gain from cheating.
 - It is not instantaneous in this setting.
 - Meaning that player i enjoys her cheating payoff during $1 + k$ periods.
- This makes cheating more attracting and:
 - increases the minimal discount factor sustaining cooperation, $\underline{\delta}$,
 - which hinders cooperation in equilibrium.




Figure 7.7. Lag at detecting cheating.

Extensions: Lag in Starting Punishments

- A similar argument as in point (4) applies if players:
 - Despite detecting cheating immediately after it occurs...
 - Need k periods to revert to the NE of the stage game.
- This happens, for instance, when “cooperation” means producing few units while “cheating” indicates producing a significantly larger number of units.
- If the cheated firm needs several periods to expand its production process, the beginning of the punishment phase is, essentially, delayed.
- This lag expands the payoff gain from cheating in the left square of Figure 7.7 during $1 + k$ periods.

Folk Theorem

- Can we also support other, partially cooperative outcomes, where, for instance, players choose *Confess* during only some periods?
- In particular, we seek to identify the *per-period payoffs* that players earn at different SPEs of this game.
- Definition.
- **Per-period payoff.** If player i 's present value from an infinite payoff stream is defined as $PV_i = \sum_{t=0}^{\infty} \delta^t v_i^t$, her per-period payoff $\bar{v}_i$ is the constant payoff that solves

$$PV_i = \frac{\bar{v}_i}{1 - \delta}$$

or, after solving $\bar{v}_i$, $\bar{v}_i = (1 - \delta)PV_i$.

- Intuitively, when player i receives that constant sum $\bar{v}_i$ in every period, she is indifferent between that (flat) payoff stream and her (potentially variable) stream of payoffs, as they both yield the same present value, PV_i .

Feasible and Individually Rational Payoffs

- We now seek to find all SPEs that can be sustained in the infinitely-repeated version of the game
 - and, as a consequence, predict which per-period payoffs players earn in each SPE.
- Afterwards, we restrict this set of payoffs to those where every player earns a higher payoff than at the NE of the stage game.
- Definition.
- **Feasible Payoffs (FP).** A feasible payoff vector $v = (v_1, v_2, \dots, v_N)$ can be achieved by convex combinations of any two payoff vectors v' and v'' .
- Therefore, in payoffs vector v , where player i earns v_i , is found by a convex combination $v_i = \alpha v'_i + (1 - \alpha)v''_i$ where $\alpha \in [0,1]$ represents the weight on v'_i .
- Intuitively, the FP set captures all possible payoff vectors that players can earn if they play the game in different ways.

Example 7.1. Finding FP in the Prisoner's Dilemma Game

- First, each vertex depicts one of the payoff pairs that players earn by playing one of the pure strategy profiles every period.
- But as Figure 7.8 suggests, FP includes more than just the four vertices.
- Convex combinations of these vertices can yield other, still feasible, per-period payoffs.
- Examples next.

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.5. The Prisoner's Dilemma Game




Figure 7.8. FP set in the prisoner's dilemma game.

Example 7.1. Finding FP in the Prisoner's Dilemma Game

- For instance, if players alternate between (C,C) and (NC,NC), the per-period payoff is $0.5 \cdot 2 + 0.5 \cdot 4 = 3$ to each player.
 - Graphically positioned in the diagonal connecting (2,2) and (4,4).
- If, instead, players alternate between (C,C) and (NC,C), the per-period payoff for:
 - player 1 is $0.5 \cdot 2 + 0.5 \cdot 0 = 1$,
 - while that of player 2 is $0.5 \cdot 2 + 0.5 \cdot 8 = 5$.
- Similarly, for all payoff pairs inside the FP set.

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.5. The Prisoner's Dilemma Game




Figure 7.8. FP set in the prisoner's dilemma game.

Example 7.1. Finding FP in the Prisoner's Dilemma Game

- Recall that the FP set does not mean that players earn any payoff pair in the FP every period.
- Instead, it means that players can play the game in such a way that, even if their payoffs vary across time, their per-period payoff would lie inside the FP.

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.5. The Prisoner's Dilemma Game




Figure 7.8. FP set in the prisoner's dilemma game.

Individually Rational Payoffs (FIR)


- Definition.
 - **Individually Rational Payoffs (IR).** An individually rational (IR) payoff vector $v = (v_1, v_2, \dots, v_N)$ satisfies
$$v_i \geq v_i^{NE} \text{ for every player } i,$$
 - where v_i^{NE} denotes her NE payoff in the unrepeated version of the game.
- Combining FP and IR, we obtain FIR set (feasible individually rational payoffs).
- Example next.

Individually Rational Payoffs (FIR)

- Example 7.2. **Finding FIR payoffs.**
- Consider the Prisoner's Dilemma game of Example 7.1.
- A per-period payoff of player i is individually rational if $v_i \geq 2$, since every player earns 2 in the NE of the stage game.

Feasibly Individually Rational Payoffs (FIR)

- Figure 7.9 depicts:
 - a vertical line representing $v_1 \geq 2$ for player 1, which holds for all payoffs to the right-hand side of 2, and
 - a horizontal line capturing $v_2 \geq 2$ for player 2, which occurs for all payoffs above 2.
- The FIR diamond is the shaded area of the FP set, indicating that both players earn a higher per-period payoff than at the NE of the stage game.



Folk Theorem and Cooperation

- Definition. **Folk Theorem.**
- Every per-period payoff vector $v = (v_1, v_2, \dots, v_N)$ in the FIR set can be sustained as the SPE of the infinitely-repeated game for a sufficiently high discount factor, δ , where $\delta \geq \underline{\delta}$.
- Graphically, the Folk theorem says that any point in the FIR diamond (on the edges or strictly inside) can be supported as a SPE of the infinitely-repeated game as long as players care enough about the future (high δ).
- *Examples:*
 - The uncooperative outcome, where (C, C) emerges in every period, can be sustained for all values of δ , yielding per-period payoffs $(2, 2)$ as depicted on the southwest corner of the FIR diamond.
 - Similarly, the fully cooperative outcome, where (NC, NC) arises in every period, yielding per-period payoffs $(4, 4)$, can be supported if $\delta \geq \frac{2}{3}$.
 - What about partial cooperation, with other per-period payoffs? Next slide.

Example 7.3. Supporting Partial Cooperation

- Consider the following modified GTS where players:
 - Alternate between (NC, C) and (C, NC) over time,
 - starting with (NC, C) in the first period.
- If either or both players deviates, both players revert to the NE of the stage game, (C, C) , forever.
- To determine whether this modified GTS can be sustained as a SPE, we must show that no player benefits by unilaterally deviating (cheating).

	Outcome	Payoffs
Period 1	NC, C	0,8
Period 2	C, NC	8,0
Period 3	NC, C	0,8
Period 4	C, NC	8,0
...		

Table 7.1. Modified GTS inducing partial cooperation

Example 7.3. Supporting Partial Cooperation

- *Odd-numbered periods.* When player 1 cooperates in this GTS, her stream of discounted payoffs starting at an odd-numbered period (e.g. period 1) is

$$\begin{aligned} 0 + \delta 8 + \delta^2 0 + \delta^3 8 + \dots &= 0(1 + \delta^2 + \dots) + 8(\delta + \delta^3 + \dots) \\ &= 8\delta(1 + \delta^2 + \dots) \\ &= \frac{8}{1 - \delta^2} \end{aligned}$$

since $\sum_{t=1}^{\infty} \delta^{2t} = \delta^0 + \delta^2 + \delta^4 + \dots = \frac{1}{1 - \delta^2}$.

Example 7.3. Supporting Partial Cooperation

- If, instead, player 1 deviates to C in an odd-numbered period, her current payoff increases from 0 to 2, yielding a stream of discounted payoffs of

$$\underbrace{2}_{\text{Deviation}} + \underbrace{\delta 2 + \delta^2 2 + \delta^3 2 + \dots}_{\text{Reversion to NE}} = 2(1 + \delta + \delta^2 + \dots) = \frac{2}{1 - \delta}$$

- **Comparison:** Comparing player 1's payoff streams, we find that she sticks to the GTS in every odd-numbered period if

$$\frac{8\delta}{1 - \delta^2} \geq \frac{2}{1 - \delta}$$

and since $1 - \delta^2 = (1 - \delta)(1 + \delta)$, we can rearrange the above equality to obtain $8\delta \geq 2(1 + \delta)$, yielding $\delta \geq \frac{1}{3}$.

Example 7.3. Supporting Partial Cooperation

- *Even-numbered periods.* When player 1 cooperates in this GTS, her stream of discounted payoffs starting at an odd-numbered period (e.g. period 1) is

$$8 + \delta 0 + \delta^2 8 + \delta^3 0 + \dots = 8(1 + \delta^2 + \dots) + 0(\delta + \delta^3 + \dots) = \frac{8}{1 - \delta^2}$$

- If, instead, player 1 unilaterally deviates from C to NC , in an even-numbered period, her current payoff actually decreases (from 8 to 4), yielding a stream of discounted payoffs

$$\underbrace{4}_{\text{Deviation}} + \underbrace{\delta 2 + \delta^2 2 + \delta^3 2 + \dots}_{\text{Reversion to NE}} = 4 + 2\delta(1 + \delta + \delta^2 + \dots) = 4 + \frac{2\delta}{1 - \delta}$$

Example 7.3. Supporting Partial Cooperation

- **Comparison:** Therefore, player 1 sticks to the GTS instead of deviating at even-numbered periods if

$$\frac{8}{1 - \delta^2} \geq 4 + \frac{2\delta}{1 - \delta}$$

and since $1 - \delta^2 = (1 - \delta)(1 + \delta)$, we can rearrange this equality to obtain


$$8 \geq 4(1 - \delta)(1 + \delta) + 2\delta(1 + \delta),$$

which further simplifies to $2\delta^2 - 2\delta + 4 \geq 0$.

Example 7.3. Supporting Partial Cooperation

- **Comparison:**

- This inequality further simplifies to $2\delta^2 - 2\delta + 4 \geq 0$.
- Plotted here for illustration purposes:
 - Showing that $2\delta^2 - 2\delta + 4$ lies above the horizontal axis for all δ .



- This inequality holds for all admissible values of $\delta \in [0, 1]$
- This means that player 1 sticks to the GTS in every even-numbered period regardless of her discount factor.

Folk Theorem

- The Folk theorem provides us with a positive and a negative result.
- Positive result:
 - We can sustain cooperation in games where such cooperation couldn't be supported in their unrepeated version or in their finitely-repeated versions.
 - In other words, players can reach Pareto-improving outcomes.
- Negative result:
 - We can reach *any* per-period payoff in the FIR diamond.
 - That's a large set of SPEs!
 - Limited predictive power.

Application to Collusion in Oligopoly

- Consider the duopoly with two firms competing à la Cournot.
- Firms face the same marginal cost of production c , where $1 > c \geq 0$
- Inverse demand function is $p(Q) = 1 - Q$, where $Q \geq 0$ denotes aggregate output
- We know:
 - $q_i = \frac{1-c}{3}$
 - NE is $(q_i, q_j) = \left(\frac{1-c}{3}, \frac{1-c}{3}\right)$
 - $\pi_i = \frac{(1-c)^2}{9}$
- As in the Prisoner's Dilemma game, we study next if a GTS can be sustained as a SPE of the infinitely-repeated version of this Cournot competition game.

Application to Collusion in Oligopoly: Collusion

$$\max_{q_1, q_2 \geq 0} \pi(q_1, q_2) = \underbrace{[(1 - q_1 - q_2)q_1 - cq_1]}_{\pi_1} + \underbrace{[(1 - q_1 - q_2)q_2 - cq_2]}_{\pi_2}$$

Since $Q = q_1 + q_2$

$$\max_{Q \geq 0} \pi(Q) = (1 - Q)Q - cQ$$

which leads to

$$\begin{aligned} Q^c &= \frac{1 - c}{2} \\ q_i^c &= \frac{Q^c}{2} = \frac{1 - c}{4} \\ p^c &= 1 - Q^c = 1 - \frac{1 - c}{2} = \frac{1 + c}{2} \\ \pi_i^c &= (p^c - c)q_i^c = \frac{(1 - c)^2}{8} \end{aligned}$$

GTS in Collusion

- We are now ready to specify the GTS in this game:
 1. In period $t = 1$, every firm i chooses the collusive output $q_i^C = \frac{1-c}{4}$.
 2. In all periods $t > 1$, every firm i chooses the collusive output $q_i^C = \frac{1-c}{4}$ if both firms produced q_i^C in every previous period. Otherwise, every firm i reverts to the NE of the stage game, choosing $q_i = \frac{1-c}{3}$, thereafter.
- To show that this GTS can be sustained in equilibrium, we need to show that firms have no incentives to deviate from it:
 - at every period t , and
 - after any previous history of play.
- We separately analyse whether firms have incentives to deviate from the GTS:
 - after observing that firms cooperated in previous periods and
 - after observing some uncooperative episodes.

GTS in Collusion

- **After a history of cooperation.** If both firms cooperated in all previous periods, choosing q_i^C , every firm i can collude in period t , as specified by the GTS, or deviate from this collusive output.
- *Profits from collusion.* If firm i colludes, it earns $\pi_i^C = \frac{(1-c)^2}{8}$ in every period, entailing a discounted present value of

$$\pi_i^C + \delta\pi_i^C + \delta^2\pi_i^C + \dots = \frac{(1-c)^2}{8(1-\delta)}$$

Profits from Deviation

If, instead, firm i deviates from the collusive output q_i^C , we must first determine firm i 's most profitable unilateral deviation, which means that firm j still produces the collusive output, q_j^C .

$$\max_{q_i \geq 0} \pi(1 - q_i - q_j^C)q_i - cq_i$$

which leads to

$$q_i(q_j^C) = \frac{1-c}{2} - \frac{1}{2}q_j^C$$

since $q_j^C = \frac{1-c}{4}$,

$$q_i^{Dev} = \frac{3(1-c)}{8}$$

(This can be directly done by inserting $q_j^C = \frac{1-c}{4}$ into firm i 's BRF.)

Therefore, equilibrium price is

$$p^{Dev} = 1 - q_i^{Dev} - q_j^C = \frac{3+5c}{8}$$

Profits from Deviation

And the profits,

$$\pi_i^{Dev} = (p^{Dev} - c)q_i^{Dev} = \frac{9(1 - c)^2}{64}$$

Therefore, if firm i deviates at any period t , it earns deviation profits $\pi_i^{Dev} = \frac{9(1-c)^2}{64}$ in that period, but Cournot profits in every period afterward

$$\begin{aligned}\pi_i^{Dev} + \delta\pi_i + \delta^2\pi_i + \dots &= \pi_i^{Dev} + \delta\pi_i(1 + \delta + \delta^2 + \dots) \\ &= \pi_i^{Dev} + \pi_i \left(\frac{\delta}{1 - \delta} \right) \\ &= \frac{9(1 - c)^2}{64} + \frac{\delta(1 - c)^2}{9(1 - \delta)}\end{aligned}$$

Comparison: We can say that at any period t after a history of collusion, firm i keeps colluding if

$$\pi_i^C \left(\frac{1}{1 - \delta} \right) \geq \pi_i^{Dev} + \pi_i \left(\frac{\delta}{1 - \delta} \right) \Rightarrow \delta \geq \delta \equiv \frac{9}{17}$$

Therefore, collusion can be sustained if firms assign a sufficiently high value to future profits. Otherwise, collusion cannot be sustained.

After a History of No Cooperation

- If one firm deviated from collusive output $q_i^C = \frac{1-c}{4}$ in any previous period, the GTS prescribes that every firm i reverts to the NE of the stage game.
- To confirm that the GTS is a SPE of the infinitely-repeated game, we need to show that every firm i has incentives to, essentially, implement this punishment.
- If, upon observing a deviation, firm i behaves as prescribed by the GTS, producing the Cournot output, it earns a discounted profit stream of

$$\pi_i + \delta\pi_i + \delta^2\pi_i + \dots = \frac{\pi_i}{1-\delta}$$

- If, instead, firm i unilaterally deviates from the Cournot output, producing $q_i \neq \frac{1-c}{3}$, while firm j behaves as prescribed by the GTS, its best response function prescribes firm i to choose:

$$q_i \left(\frac{1-c}{3} \right) = \frac{1-c}{2} - \frac{1}{2} \frac{1-c}{3} = \frac{1-c}{3}$$

- In other words, even if firm i wanted to deviate from the GTS, its best response would be to behave as prescribed by the GTS.

Minimal Discount Factor supporting collusion

- Our analysis above considered that firms compete à la Cournot, sell a homogeneous good, and face the same marginal cost of production, c
- However, we could allow for firms to compete in prices, sell differentiated products, and/or face different production costs.
- Here, we follow a more general approach.
- In particular, after a history of cooperation, every firm i 's present value from cooperating in period t is

$$\pi_i^C + \delta \pi_i^C + \delta^2 \pi_i^C + \dots = \pi_i^C (1 + \delta + \delta^2 + \dots) = \pi_i^C \left(\frac{1}{1 - \delta} \right)$$

Minimal Discount Factor supporting collusion

- If, instead, firm i deviates at period t , it earns π_i^{Dev} during that period but π_i^{NE} thereafter, yielding:

$$\begin{aligned}\pi_i^{Dev} + \delta\pi_i^{NE} + \delta^2\pi_i^{NE} + \dots &= \pi_i^{Dev} + \pi_i^{NE}(1 + \delta + \delta^2 + \dots) \\ &= \pi_i^{Dev} + \pi_i^{NE} \left(\frac{1}{1 - \delta} \right)\end{aligned}$$

Minimal Discount Factor supporting collusion

- Therefore, firm i cooperates if and only if

$$\pi_i^C \left(\frac{1}{1-\delta} \right) \geq \pi_i^{Dev} + \pi_i^{NE} \left(\frac{1}{1-\delta} \right)$$

which we can rearrange as $\pi_i^C \geq (1-\delta)\pi_i^{Dev} + \delta\pi_i^{NE}$.

- Solving for δ , yields a minimal discount factor:

$$\delta \geq \frac{\pi_i^{Dev} - \pi_i^C}{\pi_i^{Dev} - \pi_i^{NE}} \equiv \underline{\delta}$$

Minimal Discount Factor supporting collusion

$$\delta \geq \frac{\pi_i^{Dev} - \pi_i^C}{\pi_i^{Dev} - \pi_i^{NE}} \equiv \underline{\delta}$$

- Minimal discount factor $\underline{\delta}$ increases in $\pi_i^{Dev} - \pi_i^C$.
 - As deviations become more attractive, cooperation can only be sustained for higher discount factors.
- In contrast, $\underline{\delta}$ decreases in $\pi_i^{Dev} - \pi_i^{NE}$.
 - When profit loss from reverting to the NE of the stage game become more severe, firms have less incentives to cheat, expanding the range of discount factors sustaining cooperation.

Minimal Discount Factor supporting collusion

- The above ratio is convenient, as it can be readily applied to different market structures:
 - markets with two or more firms, selling homogeneous products, with symmetric or asymmetric costs, etc.
- We only need to find:
 - the NE in the stage game, and its associated profits, π_i^{NE}
 - the collusive output (or prices) that maximize firms' profits, yielding π_i^C to firm i ; and
 - firm i 's optimal deviation, fixing firm j 's collusive behaviour, so firm i 's profits are π_i^{Dev}
- Inserting profits π_i^{NE} , π_i^C , and π_i^{Dev} in the above expression for $\underline{\delta}$, we obtain the minimal discount factor sustaining collusion.

Other Collusive GTS

- We said that the above GTS with perfect collusion, where firms maximize their joint profits, can be sustained:
 - “as a SPE,”
 - rather than saying “as the unique SPE,” of the infinitely repeated game.
- From the Folk theorem, we know that the other per-period payoffs can be supported as SPEs too.

Other Collusive GTS

- The above GTS can be interpreted as the most northwest payoff pair in the FIR set
 - since firms earn the highest symmetric payoff relative to what they earn in the NE of the Cournot model.
- More generally, we can characterize the line connecting profits in the above GTS and profits in the NE of the stage game as follows:

$$\pi_i = \alpha \pi_i^C + (1 - \alpha) \pi_i^{NE}$$

where weight $\alpha \in [0,1]$ can be understood as how close firms are to the perfect collusion.

- In particular, firm i produces

$$q_i = \alpha q_i^C + (1 - \alpha) q_i^{NE} = \alpha \frac{1 - c}{4} + (1 - \alpha) \frac{1 - c}{3} = \frac{(1 - c)(4 - \alpha)}{12}$$

Implying that this firm's output coincides with the collusive output when $\alpha = 1$, but becomes the NE output when $\alpha = 0$.

Other Collusive GTS

- Consider then, the following GTS
 1. In period $t = 1$, every firm i chooses the output $q_i = \alpha q_i^C + (1 - \alpha)q_i^{NE}$
 2. In all periods $t > 1$:
 - Every firm i chooses output q_i if both firms produced q_i in every previous period.
 - Otherwise, every firm i reverts to the NE of the stage game, choosing $q_i = \frac{1-c}{3}$, thereafter.

What if the stage game has more than one NE?

- So far, we considered games with a unique NE in its unrepeated version.
 - *Examples:* the Prisoner's Dilemma game, Cournot quantity competition, or Bertrand price competition.
- The NE in these games was rather uncooperative.
 - Yet, we showed that cooperation can be supported as a SPE of the infinitely-repeated game.
- A natural question is whether cooperation can be sustained in games with more than one NE in its unrepeated version as in Matrix 7.6.

What if the stage game has more than one NE?

		Player 2		
		A	B	C
Player 1	A	8,8	-2, <u>10</u>	2,3
	B	<u>10</u> , -2	<u>2</u> , <u>2</u>	1,1
	C	1,1	1,1	<u>6</u> , <u>6</u>

Matrix 7.6. A game with two NEs

- The game has two psNEs:
 1. A “bad” NE, (B,B), where every player earns a payoff of 2.
 - Formally, we say that (C,C) Pareto dominates (B,B).
 2. A “good” NE, (C,C), where every player earns a payoff of 6.
 3. Yet, there is an even better outcome, where every player earns higher payoffs:
 - At (A,A) every player earns 8.
 - How can we sustain the cooperative outcome (A,A) being played?

What if the stage game has more than one NE?

		Player 2		
		A	B	C
Player 1	A	8,8	-2, <u>10</u>	2,3
	B	<u>10</u> , -2	<u>2</u> , <u>2</u>	1,1
	C	1,1	1,1	<u>6</u> , <u>6</u>

Matrix 7.6. A game with two NEs

- Consider the following GTS:
 1. In the first period, every player chooses A .
 2. In the second period:
 - Every player chooses C if (A, A) was played in the first period.
 - Otherwise, every player chooses B .

Stick-and-carrot

- Informally, this is known as a “stick-and-carrot” strategy, because in the second period:
 - It prescribes the Pareto dominant NE, (C, C) , if players cooperated in the first period.
 - But prescribes the Pareto dominated NE, (B, B) , if either player was uncooperative.
- Important lesson:
 - While a GTS can specify an outcome that is not a NE of the stage game during the first period, such as (A, A) ,...
 - It cannot prescribe an outcome that is not a NE in the last period.
 - Otherwise, if the GTS specified outcome (A, A) in both periods, every player would have incentives to deviate from A in the second (last) period.
- In summary, when designing GTSs in finitely-repeated games, we can:
 - “Shoot for the stars” during $T-1$ periods, but...
 - must settle for a NE in the last round of interaction.

Stick-and-carrot, Solution

Second period:

- Technically, there are 9 outcomes in each stage, for a total of $9^2 = 81$ terminal nodes
- Out of the 9 second-period subgames, there is only one initiated after outcome (A, A) was played in the first period.
 - The remaining 8 second-period subgames emerge because one or both players did not select A in the first period.
- *Non-cooperative history.* Upon observing any outcome different from (A, A) in the first period:
 - Player i anticipates that player j will play B in the second period.
 - Player i does not have incentives to deviate from the GTS in the second period, since B is a best response to B (see middle column, for instance).
- *Cooperative history.* If, instead, player i observes that outcome (A, A) was played in the first period:
 - She expects that player j will choose C in the second period.
 - Therefore, player i does not want to deviate from the GTS in the second period, since C is a best response to C (see right column, for instance).

Stick-and-carrot, Solution

- **First period:**

- If player i chooses A , as prescribed by the GTS, her payoff stream is

$$8 + \delta 6,$$

since he anticipates that (C, C) will be played in the second period (the Pareto dominant NE, or “carrot”)

- If, instead, player i unilaterally deviates:

- Her best deviation is to B , which yields a payoff of 10 today,
 - but that triggers outcome (B, B) , the Pareto dominated NE or “stick,” in the second period,
 - Payoff stream from deviation is, then:

$$10 + \delta 2.$$

- **Comparison:** player i behaves as prescribed by the GTS in the first period if

$$8 + \delta 6 \geq 10 + \delta 2 \Rightarrow \delta \geq \frac{1}{2}$$

Stick-and-carrot, Solution

- Cooperation can, then, be sustained if players assign a sufficiently high weight to future payoffs.
- Players condition last-period play
 - Choosing the good or bad NE.
 - on whether players were cooperative or uncooperative in previous interactions.
- In stage games with a unique NE, we couldn't condition last-period play.
 - Players will just behave according to the unique NE.
 - This inability to condition future play lead to the “unraveling” result, where players behave as in the stage game during every period of interaction.

Modified GTSs: An eye for an eye

- We now consider modified GTSs.
- “An eye for an eye”:
 - Code of Hammurabi, written around 1750 BC.
- It prescribes that individuals start cooperating, as in standard GTSs.
- However, if an individual is uncooperative:
 - the victim can be uncooperative, inflicting the same damage to the cheater,
 - and then players return to cooperation.
- In an infinitely-repeated game context, this adage means that, if a player cheats:
 - the cheater player gets to cheat during one period while the cheating player cooperates in that same period (so punishment phase only lasts one period),
 - afterwards players return to cooperation

Modified GTSs: An eye for an eye

- To implement this cheating sequence in a GTS, consider the game in Matrix 7.7:
 1. In the first period, every player i chooses C
 2. In subsequent periods:
 1. Every player i chooses C if (C, C) was played in all previous periods.
 2. Otherwise, every player reverts to the NE of the stage game, (B, B) .
- This GTS can be sustained if $\delta \geq \frac{2}{3}$ (see Exercise 7.17).

		Player 2		
		A	B	C
Player 1	A	0,0	4, <u>2</u>	5,-2
	B	<u>2</u> ,4	<u>8</u> , <u>8</u>	<u>14</u> ,3
	C	-2,5	3, <u>14</u>	10,10

Matrix 7.7. An eye for an eye

Modified GTSs: An eye for an eye

- But what about using a GTS that follows the “an eye for an eye” adage?
 - The punishment phase now prescribes that the victim gets to cheat her opponent for a period.
 - Then they return to the cooperative outcome (C, C) thereafter.
- As an illustration, if player 1 cheats, deviating to B while player 2 still cooperates choosing C , table 7.2 summarizes the sequence of outcomes.

	Period $t - 1$	Period t	Period $t + 1$
Outcome	B, C	C, B	C, C
Payoffs	14,3	3,14	10,10
Who cheats?	<i>Player 1</i>	<i>Player 2</i>	No one

Table 7.2. Timing for an eye-for-an-eye GTS

After a History of Cooperation

- If no player has cheated on previous periods, player i 's payoff stream from cooperating is

$$10 + 10\delta + 10\delta^2 + \dots$$

- If, instead, player 1 deviates, we first find that, conditional on player 2 cooperating, player 1's BR is to play B , earning a payoff of 14.
- Therefore, player 1's deviation:
 - Increases her current payoff from 10 to 14,
 - But then triggers a punishment, which reduces her payoff to 3 during one period.
 - Afterwards, players return to cooperation.

- In summary, the stream of payoffs that player 1 earns when deviating is

$$\underbrace{14}_{\text{Cheat}} + \underbrace{3\delta}_{\text{Cheated}} + \underbrace{10\delta^2 + 10\delta^3 + \dots}_{\text{Back to cooperation}}$$

Therefore, player cooperates if and only if

$$10 + 10\delta + 10\delta^2 + \dots \geq 14 + 3\delta + 10\delta^2 + 10\delta^3 + \dots$$

which simplifies to $10 + 10\delta \geq 14 + 3\delta$, or $\delta \geq \frac{4}{7}$.

After a History of Cheating

- When one player cheats, the GTS prescribes that the cheating player cooperates for one period, playing C , while her opponent cheats in that period, choosing B .
- If player 1 cheats in period $t - 1$, the GTS prescribes that (C, B) is played in period t , and (C, C) occurs in all subsequent periods, yielding

$$\underbrace{3}_{\text{Cheated}} + \underbrace{10\delta + 10\delta^2 + 10\delta^3 + \dots}_{\text{Back to cooperation}}$$

- If player 1 does not choose C in period t , being cheated, her best deviation is to B since, conditional on player 2 choosing B in period t (see Table 7.2), player 1's best response is B , leading to a payoff stream

$$\underbrace{8}_{\text{Cheat}} + \underbrace{3\delta}_{\text{Cheated}} + \underbrace{10\delta^2 + 10\delta^3 + \dots}_{\text{Back to cooperation}}$$

After a History of Cheating

- Intuitively, if the cheating player does not follow the GTS, not allowing the victim to cheat during one period, he only postpones the return to cooperation.
- Therefore, player 1 allows player 2 to cheat her after player 1 cheated in the first place, if and only if

$$3 + 10\delta + 10\delta^2 + 10\delta^3 + \dots \geq 8 + 3\delta + 10\delta^2 + 10\delta^3 + \dots$$

which coincide after the third term, $10\delta^2$, and for all subsequent periods. We can then rewrite this inequality as

$$3 + 10\delta \geq 8 + 3\delta \Rightarrow \delta \geq \frac{5}{7}$$

Summary

- Comparing this condition on δ with that found after a history of cooperation, $\delta \geq \frac{4}{7}$, we see that condition $\delta \geq \frac{5}{7}$ is more demanding than $\delta \geq \frac{4}{7}$ since $\frac{5}{7} \geq \frac{4}{7}$.
- In other words, if condition $\delta \geq \frac{5}{7}$ holds, condition $\delta \geq \frac{4}{7}$ must also be satisfied, so $\delta \geq \frac{5}{7}$ is a sufficient condition for cooperation.

Comparing different GTS:

An-eye-for-an-eye GTS vs Standard GTS

- **Equity of payoffs:**

- We find that the punishment phase is more equitable in the former than the latter since the cheater party is compensated in the following period, when she cheats.
- Under the standard GTS, however, the cheated party sees its payoff decrease during the cheating period, and then both players revert to the NE of the stage game, yielding a lower per-period payoff than the cheater's.

- **Minimal discount factors:**

- An-eye-for-an-eye GTS can be sustained if $\delta \geq \frac{5}{7}$,
- Standard GTS can be sustained if $\delta \geq \frac{2}{3}$,
- Cooperation is, then, more difficult to arise in the former than the latter since $\frac{5}{7} \geq \frac{2}{3}$.
- *Intuition:* The punishment phase in this GTS requires the cheating party to be cheated during one period, which is not very attractive for this player, generating strong incentives to deviate from the GTS at precisely that moment.

Short and Nasty Punishments

- Consider the same payoff matrix as before.
- The punishment is now (A, A) , yielding the lowest symmetric payoff in the matrix, $(0,0)$, and players return to cooperation immediately after.
- Consider the GTS as follows:
 1. In the first period, every player i chooses C
 2. In subsequent periods:
 - a. Every player i chooses C if (C, C) was played in all previous periods
 - b. Otherwise, every player chooses A during one period. If (A, A) was played in the last period, every player selects C thereafter

Short and Nasty Punishments

- If player 1 cheats, deviating to B while player 2 still cooperates choosing C
- Table 7.3 summarizes the sequence of outcomes if players behave according to this GTS

	Period $t - 1$	Period t	Period $t + 1$
Outcome	B, C	A, A	C, C
Payoffs	14,3	0,0	10,10

Table 7.3. Timing of the short-and-nasty punishments GTS

Short and Nasty Punishments

- Cooperation can be sustained with this GTS if $\delta \geq \frac{2}{5}$.
 - (Exercise 7.1, as a practice)
- Table 7.4 reports the minimal discount factor $\underline{\delta}$ under different GTS in the game in Matrix 7.7, where $\frac{5}{7} > \frac{2}{3} > \frac{2}{5}$.

	Minimal discount factor, $\underline{\delta}$
Standard GTS, section 7.4.2	$\underline{\delta} = \frac{2}{3}$
An-eye-for-an-eye GTS, section 7.8.1	$\underline{\delta} = \frac{5}{7}$
Short-and-nasty-punishments GTS, section 7.8.2	$\underline{\delta} = \frac{2}{5}$

Table 7.4. Minimal Discount Factors sustaining cooperation in the game in Matrix 7.7 when using different GTS

Imperfect Monitoring

- Players may imperfectly monitor her rival's actions, observing which action she chose with a certain probability.
- Consider Matrix 7.8:
 - (C, C) is the NE of the stage game,
 - but (NC, NC) would increase both players' payoffs from 2 to 4.

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.8. The Prisoner's Dilemma Game

Imperfect Monitoring

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.8. The Prisoner's Dilemma Game

- Under imperfect monitoring, if player j chooses to cooperate, NC , we assume that, for simplicity, player i cannot perfectly detect deviation to C .
- If, instead, player j chooses C , the probability that player i observes a deviation from NC is $p \in [0,1]$.
- Intuitively:
 - when probability $p \rightarrow 0$, deviations are rarely observed,
 - while when $p \rightarrow 1$, deviations are almost always observed, as in a context with perfect monitoring.

Imperfect Monitoring

- Consider the standard GTS in this context:
 1. In the first period, every player chooses NC .
 2. In every subsequent period, every player chooses NC if (NC, NC) was observed in all previous periods. Otherwise, revert to the NE of the stage game, (C, C) , thereafter.
- That is, the cooperative outcome (NC, NC) is now observed but it may not have been played.
 - This suggests that outcome (NC, NC) may not have been played, because at least one person deviated, but cooperation continues if such outcome is observed.
- The opposite also applies:
 - Outcome (NC, NC) may have been played, but cooperation stops if one of both players did not observe such outcome in a previous period.

After a history of observed cooperation

- If both players observe outcome (NC, NC) in all previous periods, player i 's payoff stream from cooperating is

$$4 + 4\delta + 4\delta^2 + \dots = \frac{4}{1 - \delta}$$

- If, instead player i unilaterally deviates to C , her payoff increases to 8 in the current period, but in the next period:

- Player j detects her deviation with probability p , reverting to the NE of the stage game with a payoff of 2 forever, or...
- Player j doesn't detect this deviation with probability $(1 - p)$, which lets player i return to cooperation (choosing NC tomorrow) as if no deviation ever happened.

- In summary, player i 's payoff stream from unilaterally deviating to C is

$$\underbrace{8}_{\text{Current gain from dev.to } C} + \delta \left[p \underbrace{\frac{2}{1 - \delta}}_{\text{Detected}} + (1 - p) \underbrace{\frac{4}{1 - \delta}}_{\text{Undetected}} \right]$$

After a history of observed cooperation


- For player i to cooperate after a history of cooperation, we need that

$$\frac{4}{1-\delta} \geq 8 + \delta \left[p \frac{2}{1-\delta} + (1-p) \frac{4}{1-\delta} \right]$$

And solving for δ , yields

$$\delta \geq \frac{2}{2+p} \equiv \underline{\delta}(p)$$


- Figure 7.9 depicts the minimal discount factor $\underline{\delta}(p)$.



After a history of observed cooperation

$$\Rightarrow \delta \geq \frac{2}{2+p} \equiv \underline{\delta}(p)$$

- As deviations are more likely detected (higher p):
 - the expected punishment increase,
 - expanding the range of δ 's sustaining cooperation.
- In contrast, when deviation cannot be detected ($p = 0$):
 - the minimal discount factor becomes $\underline{\delta}(0) = 1$,
 - implying that collusion cannot be supported in equilibrium.



After a history of observed deviation

- If player observes that (NC, NC) was not played in a previous period, this means that player j did not cooperate.
- If player i behaves as prescribed by the GTS, she should revert to the NE of the stage game, (C, C) , thereafter, earning a payoff stream

$$2 + 2\delta + 2\delta^2 + \dots = \frac{2}{1 - \delta}$$

After a history of observed deviation

- If, instead, player i deviates, choosing NC while player j chooses C (recall that this is a unilateral deviation from the GTS, upon observing no cooperation in previous periods) player i earns 0 in this period.
- Players then play (C, C) in all subsequent periods yielding

$$0 + 2\delta + 2\delta^2 + \dots = 2\delta(1 + \delta + \delta^2 + \dots) = \frac{2\delta}{1 - \delta}$$

- Therefore, upon observing a deviation, player i prefers to behave as prescribed by the GTS than deviating from it since

$$\frac{2}{1 - \delta} \geq \frac{2\delta}{1 - \delta} \Rightarrow \delta \leq 1, \text{ which holds by assumption since } \delta \in [0, 1).$$

- Overall, the only condition that we need to support this GTS as SPE is $\delta \geq \underline{\delta}(p)$, as found above.

Chapter 8: Bayesian Nash Equilibrium

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- More often than not, players interact in games where at least one of them is uninformed about some relevant information.
- In auction, all players are uninformed about some piece of information.
 - Every bidder privately observes her valuation for the object on sale...
 - but does not observe the valuation that other bidders assign to the object.
- A similar example applies to an industry where firms compete in prices:
 - each firm privately observes its marginal production costs but not observes the exact cost of its rivals.
- However, players observe probability distribution over this parameters.
 - This can come from months of research, or because the player hired a consulting company to provide estimates.

Introduction – “types”

- We will refer to the private information a player observes as her “type.”
 - Every player observes her type (e.g., production cost) but...
 - doesn’t observe her rivals’ types.
- For the game to qualify as incomplete information, we must have that:
 - At least one player does not observe the type of at least one of her rivals.

Introduction – Best responses

- We seek to adapt the NE solution concept to a context of incomplete information.
- We start by defining a player's strategy:
 - The action that she chooses as a function of her type.
- We then use this definition of strategy, to identify a player's best response in this context.
 - A player's BR should also be a function of her type.
- Because the player cannot observe her rivals' types:
 - She will need to find her best response *in expectation*, but conditional on her type.
- Following the same steps as in Chapter 3, we then use BR to describe a NE in a context of incomplete information.
 - We will obtain Bayesian Nash Equilibrium (BNE).
 - We will present two approaches to find BNEs.

Background:

Players' types and their associated probability

Discrete types.

- In a game of incomplete information, every player i observes her type, θ_i , where $\theta_i \in \Theta_i$.
 - This may represent, for instance, a high or low production costs, $\theta_i = H$ or $\theta_i = L$ implying that $\Theta_i = \{H, L\}$.
- Player i , however, does not observe her rivals' types:
 - θ_j in a two-player game.
 - $\theta_{-i} = (\theta_1, \theta_2, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_N)$ in an N -player game.
- Players, however, know the probability distribution over types.
- *Example:*
 - Firm i knows that its rival's type is either:
 - $\theta_j = H$ with probability q , where $q \in [0, 1]$,
 - or $\theta_j = L$ with probability $1 - q$ (and this information is common knowledge).

Background:

Players' types and their associated probability

Discrete types.

- More generally, in a setting where player i can have K different types, we write that her type space is

$$\Theta_i = \{\theta_i^1, \theta_i^2, \dots, \theta_i^K\}$$

- The probability that her type is θ_i^1 can be expressed as

$$\Pr(\theta_i = \theta_i^1) = p_i^1,$$

- and similarly for the probability that her type is θ_i^K can be written as

$$\Pr(\theta_i = \theta_i^K) = p_i^K.$$

- Since, in addition, probabilities of each type, $(p_i^1, p_i^2, \dots, p_i^K)$ satisfy $p_i^k \in [0,1]$ for every $k = \{1, 2, \dots, K\}$ and $\sum_{k=1}^K p_i^k = 1$, we can omit the probability of the last type, writing it, instead, as

$$p_i^K = 1 - \sum_{k \neq K} p_i^k.$$

- In a context with three types, for instance, $\Theta_i = \{\theta_i^1, \theta_i^2, \theta_i^3\}$, the associated probabilities can be expressed as

$$(p_i^1, p_i^2, 1 - p_i^1 - p_i^2).$$

Background:

Players' types and their associated probability

Continuous Types.

- Our notation can be adapted to a setting where types are continuous.
- A player i 's type in this setting, θ_i , is drawn from a continuous cumulative probability distribution, that is,
$$F(x) = \Pr(\theta_i \leq x)$$
- Intuitively, $F(x)$ measures the probability that player i 's type, θ_i , lies weakly below x .
- This representation also helps us find a density function, $f(x)$, associated above $F(x)$, if one exists, by computing its first-order derivative:
$$f(x) = F'(x).$$
- Recall that density $f(x)$ describes the probability that “player i 's type, θ_i , is exactly x ,” that is, $f(x) = \Pr\{\theta_i = x\}$, which can be quite useful in some games.

Background:

Players' types and their associated probability

Continuous Types.

- Uniform distribution.
 - If player i 's types are uniformly distributed, $F(x) = x$.
 - The density function is $f(x) = 1$, meaning that all types are equally likely to occur.
- Exponential distribution.
 - If types are exponentially distributed, $F(x) = 1 - \exp(-\lambda x)$, its density function is $f(x) = \lambda \exp(-\lambda x)$
 - implying that parameter λ represents how quickly the density function decreases as we increase x , and is often known as the “decay rate.”
 - Figure.
 - Intuitively, a higher λ means that $f(x)$ puts most probability weight on low values of x .

Strategies under Incomplete Information

- If players operate under incomplete information, we must have that:
 - at least one player does not observe the types of at least one of her opponents.
- If player i does not observe any private information:
 - The player only has access to the “public information” in the game that everyone else also observes.
- In contrast, if player i observes some piece of private information:
 - She can condition her strategy on her type,
 - implying that her strategy in this context is a function of θ_i , which we express as $s_i(\theta_i)$.
- In some games, we may have:
 - Some perfectly informed players, who observe everyone’s types because of their experience in the industry or because they get to act before everyone else;
 - Some uninformed players who observe their own types but not their rival’s.

Representing Asymmetric Information as Incomplete Information

		Player 2	
		L	R
Player 1	U	$x, 17$	5, 10
	D	10, 0	10, 17

Matrix 8.1. Simultaneous move game where player 1 privately observes x

- Player 1 observes the realization of random variable x before playing the game.
- Player 2 only knows that its realization is either $x = 20$ or $x = 12$ with equal probabilities. This information is common knowledge.
- Matrix 8.1, therefore, suggests that player 2 faces *imperfect* information because she does not observe the realization of x , while player 1 does.

Representing Asymmetric Information as Incomplete Information

- Player 2 knows, however, that she faces either of the two games depicted in Figures 8.1a and 8.1b.
 - We circled the only payoff that differs across both figures, corresponding to (U, L) .




Figure 8.1a. Simultaneous move game where $x = 20$




Figure 8.1b. Simultaneous move game where $x = 12$

Representing Asymmetric Information as Incomplete Information

- Alternatively, we can represent the above setting as a game of *incomplete* information.
- See figure.
- Player 2:
 - Instead of not observing the realization of random variable x ,
 - Doesn't observe the move from a fictitious player ("nature"), who determines whether:
 - $x = 20$ or $x = 12$ at the beginning of the game.
 - We are just connecting the two trees in the previous slide, with a move of nature.
 - This trick is due to Harsanyi (1967).




Figure 8.2. Combining both games of figure 8.1 to allow for incomplete information.

Best Response under Incomplete Information

Definition.

Best response under incomplete information. Player i regards $s_i(\theta_i)$ as a best response to strategy profile $s_{-i}(\theta_{-i})$, if

$$EU_i(s_i(\theta_i), s_{-i}(\theta_{-i})) \geq EU_i(s'_i(\theta_i), s_{-i}(\theta_{-i}))$$

for every available strategy $s'_i(\theta_i) \in S_i$ and every type $\theta_i \in \Theta_i$.

- This means that strategy $s_i(\theta_i)$ yields a weakly higher *expected* payoff than any other strategy $s'_i(\theta_i)$ against $s_{-i}(\theta_{-i})$, and this holds for all player i 's types, θ_i .
- For instance, if $\Theta_i = \{H, L\}$, and player i 's type is H, player i maximizes her expected payoff responding with $s_i(H)$ against her rival strategy; and similarly, when i 's type is L, she maximizes her expected payoff responding with $s_i(L)$.

Best Response under Incomplete Information

Relative to the definition of best response in contexts of complete information (chapter 3), this definition differs in two dimensions.

1. *Expected Utility.* First, player i seeks to maximize her expected, instead of certain, utility level.
 - That's because she faces the uncertainty from not observing some private information.
2. *Best Response as a function of types.* In addition, player i finds a strategy $s_i(\theta_i)$, as opposed to the strategy s_i , in the best responses of chapter 3.
 - Intuitively, her choice may now depend on her privately observed type, meaning that her strategy may differ for at least two of her types.
 - Player i must find an optimal strategy for each of her type, potentially different.

Bayesian Nash Equilibrium

Definition.

Bayesian Nash Equilibrium (BNE). A strategy profile $(s_i^*(\theta_i), s_{-i}^*(\theta_{-i}))$ is a Bayesian Nash Equilibrium if every player chooses a best response (under incomplete information) given her rivals' strategies.

- Therefore, in a two-player game, a strategy profile is a BNE if it is a *mutual* best response, thus being analogous to the definition of NE.
- As a result, no player has unilateral incentives to deviate.
- This definition assumes, of course, that players select best responses in the sense defined in section 8.2, where players seek to maximize their expected payoffs given the information they observe.

Ex-ante and ex-post stability

- As under complete information:
 - In a BNE every player must be playing mutual best responses to each other's strategies,
 - Thus making the strategy profile stable.
- This stability is, in this context, understood from an ex-ante perspective,
 - That is, given the information that every player observes when she is called to move.
- As opposed to ex-post:
 - which assumes that at the end of the game, every player gets to observe her rivals' types.
- Players could say
 - "if I had to play the game again, given the (little) information I had, I would not have incentives to deviate. I would play in the same way I just did."

Ex-ante and ex-post stability

- Ex-post stability would, however, imply that, after playing according to a BNE, players would say:
 - “If I had to play the game again, but given the (rich) information I have now, once the game is over and I can observe everyone’s types, I would still play in the same way.”
- Ex-post stability is, therefore, more demanding than ex-ante stability:
 - It requires that players have no incentives to deviate from their equilibrium strategies for all realizations of type profiles
 - (i.e., regardless of her rival’s types).

Tool 8.1. Finding BNEs using the Bayesian Normal Form Representation (two-player game)

1. Write down each player's strategy set, S_i and S_j . Recall that privately informed players condition their strategies on their types. Similarly, players who observed other players actions before being called to move can condition their strategy on the node or information set where they are called to move.
2. Building the Bayesian Normal Form representation of the game:
 - a. Depict a matrix with as many rows as strategies in S_i and as many columns as strategies in S_j , leaving all cells in this matrix empty.
 - b. Find the expected utility that player i earns in each cell.
3. Find each player's best response against her opponent strategy and underline her best response payoffs.
4. Identify which cell/s have all players' payoff underlined, thus being mutual best responses. This cell/s are the BNEs of the game.

Example 8.1. Finding BNEs

		Player 2	
		L	R
Player 1	U	x,17	5,10
	D	10,0	10,17

Matrix 8.2. Simultaneous move game where player 1 privately observes x

- **First step.** For simplicity, we often start with the uninformed player (player 2 in this setting), whose strategy set is $S_2 = \{L, R\}$ since she cannot condition her strategy on player 1's type.
- In contrast, player 1 observes her type (the realization of x) and can condition her strategy on x , that is, $s_1(x)$. Therefore, player 1's strategy set is $S_1 = \{U^{20}U^{12}, U^{20}D^{12}, D^{20}U^{12}, D^{20}D^{12}\}$
- Strategy $U^{20}U^{12}$, for instance, prescribes that player 1 chooses U both after observing that $x = 20$ and that $x = 12$. In $U^{20}D^{12}$ ($D^{20}U^{12}$), however, player 1 chooses U only after observing that $x = 20$ (*that* $x = 12$, *respectively*). Finally, in $D^{20}D^{12}$, player 1 selects D regardless of her type.

Example 8.1. Finding BNEs

- **Second step.** From the first step, player 1 has four available strategies while player 2 only has two, yielding the Bayesian normal form representation of the game in Matrix 8.3.

		Player 2	
		L	R
Player 1	$U^{20}U^{12}$		
	$U^{20}D^{12}$		
	$D^{20}U^{12}$		
	$D^{20}D^{12}$		

Matrix 8.3. Player 1 privately observes x — Bayesian normal form representation

- where

$$EU_1(U^{20}U^{12}, L) = \underbrace{\frac{1}{2}20}_{\text{if } x=20} + \underbrace{\frac{1}{2}12}_{\text{if } x=12} = 16$$

$$EU_2(U^{20}U^{12}, L) = \underbrace{\frac{1}{2}17}_{\text{if } x=20} + \underbrace{\frac{1}{2}17}_{\text{if } x=12} = 17$$

$$EU_1(U^{20}D^{12}, L) = \frac{1}{2}20 + \frac{1}{2}10 = 15$$

$$EU_2(U^{20}U^{12}, L) = \frac{1}{2}17 + \frac{1}{2}0 = 8.5$$

Example 8.1. Finding BNEs

- Operating similarly for the remaining strategy profiles, we find the expected payoffs in 8.3, obtaining matrix 8.4.

		Player 2	
		L	R
Player 1	$U^{20}U^{12}$	16,17	5,10
	$U^{20}D^{12}$	15,8.5	7.5,13.5
	$D^{20}U^{12}$	11,8.5	7.5,13.5
	$D^{20}D^{12}$	10,0	10,17

Matrix 8.4. Player 1 privately observes x — Finding expected payoffs

Example 8.1. Finding BNEs

- **Third step.** Underlining best response payoffs, we obtain Matrix 8.5. The two cells where both players' (expected) payoffs are underlined indicate the two BNEs of this game: $(U^{20}U^{12}, L)$ and $(D^{20}D^{12}, R)$.

		Player 2	
		L	R
Player 1	$U^{20}U^{12}$	<u>16,17</u>	5,10
	$U^{20}D^{12}$	15,8.5	7.5, <u>13.5</u>
	$D^{20}U^{12}$	11,8.5	7.5, <u>13.5</u>
	$D^{20}D^{12}$	10,0	<u>10,17</u>

Matrix 8.5. Player 1 privately observes x — Best Response payoffs

Tool 8.2. Finding BNEs by focusing on the informed player first

1. Focus on the privately informed player i .
 - For each of her $k \geq 2$ types, find her best response function to her rival's strategies, e.g., $s_i(s_j|\theta_i)$.
 - You must find k best response functions, one for each type, e.g., $s_i(s_j|H)$ and $s_i(s_j|L)$.
2. If all players are privately informed about their types, simultaneously solve for $s_i(\theta_i)$ in best response function $s_i(s_j|\theta_i)$, finding k equilibrium strategies, one for each type, that is, $s_i^*(\theta_i) = (s_i^*(\theta_i^1), \dots, s_i^*(\theta_i^k))$.
 - Therefore, $(s_i^*(\theta_i^1), s_i^*(\theta_i^2))$ is the BNE of the game.
3. If one player is uninformed (does not privately observe some information), analyze the uninformed player j .
 - a. Find which strategy gives player j the highest expected utility to her rival's strategies. This strategy is her best response function $s_j(s_i)$, which is not conditional on her rival's types.
 - b. Simultaneously solve for $s_i(\theta_i)$ and s_j in best response functions $s_i(s_j|\theta_i)$, one for each player i 's types, and $s_j(s_i)$.

Example 8.2. Cournot competition with asymmetrically informed firms

- Consider a duopoly market where firms compete à la Cournot and face inverse demand function $p(Q) = 1 - Q$
- Firms interact in an incomplete information context where firm 2's marginal costs are either:
 - $c_H = \frac{1}{2}$ or $c_L = 0$, occurring with probability p and $1 - p$, respectively, where $p \in (0,1)$.
- Firm 2 privately observes its marginal cost but firm 1 does not observe firm 2's costs.
- Because firm 1 has operated in the industry for a long period, all firms observe that this firm's costs are $c_H = \frac{1}{2}$.

Example 8.2. Cournot competition with asymmetrically informed firms

- *Privately informed firm, high costs.*

$$\max_{q_2^H \geq 0} \pi_2^H(q_2^H) = (1 - q_2^H - q_1)q_2^H - \frac{1}{2}q_2^H$$

Differentiating with respect to q_2^H , yields

$$1 - 2q_2^H - q_1 - \frac{1}{2} = 0$$

and solving for q_2^H , we obtain firm 2's best response function when its costs are high,

$$q_2^H(q_1) = \frac{1}{4} - \frac{1}{2}q_1$$

which originates at $\frac{1}{4}$, decreases in its rival's output at rate of $\frac{1}{2}$, and becomes zero for all $q_1 > \frac{1}{2}$.

Example 8.2. Cournot competition with asymmetrically informed firms

- *Privately informed firm, low costs.*

$$\max_{q_2^L \geq 0} \pi_2^L(q_2^L) = (1 - q_2^L - q_1)q_2^L$$

Differentiating with respect to q_2^L , yields

$$1 - 2q_2^L - q_1 = 0$$

and solving for q_2^L , we obtain firm 2's best response function when its costs are low,

$$q_2^L(q_1) = \frac{1}{2} - \frac{1}{2}q_1$$

which originates at $\frac{1}{2}$, decreases in its rival's output at rate of $\frac{1}{2}$, and becomes zero for all $q_1 > 1$.

Equilibrium Output under Incomplete Information

- Figure 8.3a compares firm 2's best response functions when its costs are high and low.
- We find that
$$q_2^L(q_1) > q_2^H(q_1),$$
- indicating that:
 - for a given output by firm 1, firm 2 produces more units when its own costs are low than high.




Figure 8.3a. Firm 2 - High or Low costs.

Figure 8.3b. Equilibrium output under incomplete information.

Example 8.2. Cournot competition with asymmetrically informed firms

- *Uninformed firm.*

$$\max_{q_1 \geq 0} \pi_1(q_1) = p \underbrace{[(1 - q_2^H - q_1)q_1]}_{\text{Firm 2's costs are high}} + (1 - p) \underbrace{[(1 - q_2^L - q_1)q_1]}_{\text{Firm 2's costs are low}} - \frac{1}{2} q_1$$

where the last term indicates firm 1's high costs, which are certain, and observed by all firms.

Differentiating with respect to q_1 , yields:

$$p(1 - q_2^H - 2q_1) + (1 - p)(1 - q_2^L - 2q_1) - \frac{1}{2} = 0$$

which leads to firm 1's best response function:

$$q_1(q_2^H, q_2^L) = \frac{1}{4} - \frac{pq_2^H + (1 - p)q_2^L}{2}$$

which originates at $\frac{1}{4}$, and decreases at rate of $\frac{1}{2}$ when firm 2's expected output increases as captured by the term $pq_2^H + (1 - p)q_2^L$.

Equilibrium Output under Incomplete Information

- Figure 8.3b depicts firm 2's expected output

$$E(q_2) \equiv pq_2^H + (1 - p)q_2^L$$

in the dashed line between $q_2^H(q_1)$ and $q_2^L(q_1)$.

- It also plots firm 1's best response as a function of $E(q_2)$.
- Only one BRF for the uninformed firm!




Figure 8.3a. Firm 2 - High or Low costs.

Figure 8.3b. Equilibrium output under incomplete information.

Example 8.2. Cournot competition with asymmetrically informed firms

- *Finding equilibrium output in the BNE.* A common approach is to insert best response functions $q_2^H(q_1)$ and $q_2^L(q_1)$ into $q_1(q_2^H, q_2^L)$, as follows,

$$q_1 = \frac{1}{4} - \frac{p \overbrace{\left(\frac{1}{4} - \frac{1}{2} q_1\right)}^{q_2^H(q_1)} + (1-p) \overbrace{\left(\frac{1}{2} - \frac{1}{2} q_1\right)}^{q_2^L(q_1)}}{2}$$

Rearranging, this expression simplifies to

$$q_1 = \frac{1}{4} + \frac{2q_1 + p - 2}{8}$$

which, solving for q_1 , yields firm 1's equilibrium output

$$q_1^* = \frac{p}{6}$$

Inserting q_1^* into firm 2's best response functions, we obtain

$$q_2^H(q_1^*) = \frac{3-p}{12}; \quad q_2^L(q_1^*) = \frac{6-p}{12}$$

Example 8.2. Cournot competition with asymmetrically informed firms

- Therefore, the BNE of this game is given by the triplet of output levels

$$(q_1^*, q_2^{H*}, q_2^{L*}) = \left(\frac{p}{6}, \frac{3-p}{12}, \frac{6-p}{12} \right)$$

- Figure 8.4 depicts these three output levels, showing that:
 - firm 1's equilibrium output, q_1^* , is increasing in the probability that its rival has a high cost, p , but...
 - firm 2's equilibrium output is decreasing in this probability, which holds both its costs are both high and low.




Figure 8.4. Cournot competition with asymmetrically informed firms – Equilibrium output

Example 8.2. Cournot competition with asymmetrically informed firms

- The figure also embodies complete information as special cases.
- When $p \rightarrow 0$, firm 2 has (almost certainly) low costs, thus being very competitive.
 - The uninformed firm 1 remains inactive, while firm 2 produces $\frac{1}{4}$ when its cost is high and $\frac{1}{2}$ when its cost is low.
- When $p \rightarrow 1$, firm 2 has (almost certainly) high costs, thus being very competitive.
 - The uninformed firm 1 produces $\frac{1}{6}$, as much as firm 2's output.




Figure 8.4. Cournot competition with asymmetrically informed firms – Equilibrium output

Evaluating BNE as a solution concept

1. Existence? Yes.

- When we apply BNE to any game, we find that at least one equilibrium exists.
- This result is equivalent to the existence of NEs in complete information games, but extended to an incomplete information context.
- We may need to rely on mixed strategies. See exercise 8.17 and auctions in chapter 9.

2. Uniqueness? No.

- Games of incomplete information may have more than one BNE, as illustrated in Example 1 where we found two BNEs.

3. Robust to small payoff perturbations? Yes.

- We find that BNE is also robust to small payoff perturbations.
- This is due to the fact that, if a strategy s_i yields a higher expected payoff than another strategy s_{-i} , it must still generate a higher expected payoff than s_{-i} after we apply a small payoff perturbation.

4. Socially optimal? No.

- BNE does not necessarily yield socially optimal outcomes
- Strategy profile $(D^{20}D^{12}, R)$ is a BNE, with equilibrium payoffs $(10, 17)$, but it is not socially optimal because other strategy profiles, such as $(U^{20}U^{12}, L)$, provide payoffs $(16, 17)$, thus improving player 1's payoff without lowering player 2's.

What if both players are privately informed?

Example 8.3. Cournot competition with symmetrically uninformed firms.

- Assume now that every firm $i = \{1,2\}$ observes its marginal cost.
 - Its marginal cost is either $c_H = \frac{1}{2}$ or $c_L = 0$, occurring with probability p and $1 - p$, respectively.
 - Every firm i doesn't observe its rival's (firm j 's) marginal cost.
- This means that every firm will have two BRFs, one upon observing its cost is high and another after observing its cost is low.
- High costs.* If firm i has high marginal costs, $c_H = \frac{1}{2}$, it solves the following problem:

$$\max_{q_i \geq 0} \pi_i^H(q_i) = p \underbrace{[(1 - q_j^H - q_i)q_i]}_{\text{Firm } j' \text{'s costs are high}} + (1 - p) \underbrace{[(1 - q_j^L - q_i)q_i]}_{\text{Firm } j' \text{'s costs are low}} - \frac{1}{2}q_i$$

Differentiating with respect to q_i and finding BR function:

$$q_i^H(q_j^H, q_j^L) = \frac{1}{4} - \frac{pq_j^H + (1 - p)q_j^L}{2}$$

What if both players are privately informed?

Low costs. If firm i has, instead, low marginal costs, $c_L = 0$, it solves:

$$\max_{q_i \geq 0} \pi_i^L(q_i) = p \underbrace{\left[(1 - q_j^H - q_i)q_i \right]}_{\text{Firm } j' \text{'s costs are high}} + (1 - p) \underbrace{\left[(1 - q_j^L - q_i)q_i \right]}_{\text{Firm } j' \text{'s costs are low}}$$

Firm j 's costs are high

Firm j 's costs are low

Differentiating with respect to q_i , yields

$$p(1 - q_j^H - 2q_i) + (1 - p)(1 - q_j^L - 2q_i) = 0$$

which leads to firm i 's best response function:

$$q_i^L(q_j^H, q_j^L) = \frac{1}{2} - \frac{pq_j^H + (1 - p)q_j^L}{2}$$

What if both players are privately informed?

Comparing best response functions:

$$q_i^L(q_j^H, q_j^L) = \frac{1}{2} - \frac{pq_j^H + (1-p)q_j^L}{2}$$

Relative to $q_i^H(q_j^H, q_j^L)$, the best response function $q_i^L(q_j^H, q_j^L)$:

- originates at a higher vertical intercept, $\frac{1}{2}$,
- shares the same slope.
- Intuitively, firm i produces a larger output, for a given expected output from its rival, when its own costs are low than high.

What if both players are privately informed?

- *Finding equilibrium output in the BNE.* In a symmetric BNE, firms produce the same output level when their costs coincide, that is, $q_i^H = q_j^H = q^H$ and $q_i^L = q_j^L = q^L$.
- Inserting this property in best response functions $q_i^H(q_j^H, q_j^L)$ and $q_i^L(q_j^H, q_j^L)$, yields

$$q^H = \frac{1}{4} - \frac{pq^H + (1-p)q^L}{2} \quad \text{and} \quad q^L = \frac{1}{2} - \frac{pq^H + (1-p)q^L}{2}$$

Solving two equations and two unknowns, we obtain

$$q^H = \frac{1+p}{12} \quad \text{and} \quad q^L = \frac{4+p}{12}$$

which are both increasing in the probability that its rival's costs are high, p .

Comparison across information settings

- How equilibrium output levels vary from those we found in example 8.2, $(q_1^*, q_2^{H*}, q_2^{L*})$, where only one player was privately informed about its costs (firm 2), while firm 1's costs were common knowledge.
- **Firm 1:** it produces more units when its rival does not observe its costs (example 8.3) than when its rival does (example 8.2), because

$$q^H - q_1^* = \frac{1+p}{12} - \frac{p}{6} = \frac{1-p}{12} > 0, \text{ and}$$
$$q^L - q_1^* = \frac{4+p}{12} - \frac{p}{6} = \frac{4-p}{12} > 0.$$

since $0 < p < 1$ by assumption.

- **Firm 2:** In contrast, firm 2's output satisfies

$$q^H - q_2^{H*} = \frac{1+p}{12} - \frac{3-p}{12} = -\frac{1-p}{6} < 0, \text{ and}$$
$$q^L - q_2^{L*} = \frac{4+p}{12} - \frac{6-p}{12} = -\frac{1-p}{6} < 0.$$

Indicating that firm 2 produces more output when it observes its rival's cost (example 8.2) than when it does not (example 8.3).

Chapter 9: Auction Theory

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- Auctions have been historically reported as early as Babylon around 500 BC and in the Roman Empire in 193 AD.
 - They remain popular nowadays through online auctions such as eBay, eBid, QuiBids, AuctionMaxx, DealDash, or LiveAuctioneers, etc.
- Understanding optimal bidding behavior in different auction formats, then, is not only an interesting application of BNEs but also a useful tool for online shopping.
- We will consider different auction formats:
 - First-price auction (FPA). The bidder submitting the highest bid wins the auction. She pays the bid she submitted.
 - Second-price auction (SPA). The bidder submitting the highest bid wins the auction. But does not pay her bid; she pays, instead, the second-highest bid.
 - Common value auctions.

Auctions as allocation mechanisms

Two basic criteria to characterize auctions: the assignment and payment rules.

a. **Assignment Rule**

- Informally, the assignment (or allocation) rule in an auction just answers the question “who gets the object?”
 - In the first-price auction, for instance, the individuals submitting the highest bid receives the object for sale.
 - A similar assignment rule applies to other auction formats, such as:
 - the second-price, third-price, k^{th} -price, or all-pay auctions (APA), where the object still goes to the bidder submitting the highest bid.
 - In settings where the seller offers $m \geq 2$ units of the same good (multi-unit auctions), the assignment rule in first-price auctions, for instance, determines that:
 - the first unit of the object goes to the individual submitting the highest bid,
 - the second unit to the individual submitting the second-highest bid,
 - and similarly for the remaining units.

Auctions as allocation mechanisms

a. Assignment Rule

- In lottery auctions, however, bidder i may not receive the object even when she submits the highest bid. In particular, bidder i 's probability of winning the auction is

$$\Pr(\text{win}) = \frac{b_i}{b_i + B_{-i}}$$

where $B_{-i} = \sum_{j \neq i} b_j$ denotes the aggregate bids of player i 's rivals. Therefore, bidder i 's probability of winning:

- i. Increases in her own bid, b_i ;
 - ii. Satisfies $\Pr(\text{win}) < 1$ if at least one of her rivals submits a positive bid, $B_{-i} > 0$; and
 - iii. Decreases in her rivals' bids.
- *Example:* bidder i submits $b_i = \$100$, $N = 10$, and $b_j = \$5$ for every $j \neq i$, then $\Pr(\text{win}) = \frac{100}{100 + (10 \times 5)} = \frac{2}{3}$.

Auctions as allocation mechanisms

b. Payment Rule

- The criterion answers the question “how much each bidder pays,” which allows for the winner and losers to pay a monetary amount to the seller.
- In the first-price auction, the winner is the only bidder paying to the seller, and she must pay the bid she submitted.
- In the second-price auction, the winner is again the only player to the seller, but in this case she pays the second-highest bid (not the bid that she submitted).
- A similar argument applies to k^{th} -price auctions, where only the winning bidder pays to the seller, specifically the k^{th} -highest bid.
- In the all-pay auction, all bidders (the winner and the $N - 1$ losers) pay the bid that each of them submitted.

Setting

- All auctions we consider in this chapter will share these ingredients:
- A seller offers an object to $N \geq 2$ bidders.
- Every bidder i privately observes her valuation for the object, $v_i \in [0,1]$, but does not observe bidder j 's valuation, v_j , for every bidder $j \neq i$.
- It is common knowledge, however, that valuations v_i and v_j are independent and identically distributed (i.i.d.)
- v_j is drawn from a cumulative distribution function $F_j(v) = \Pr\{v \leq v_j\}$, with positive density in all its support, i.e., $f_j(v) = F_j'(v) > 0$ for all $v \in [0,1]$.
- After observing her valuation v_i , every bidder i simultaneously and independently submits her bid $b_i \geq 0$ for the object.
- The seller observes the profile of submitted bids, $b = (b_1, b_2, \dots, b_N)$, and:
 - According to the assignment and payment rules of the auction, the seller declares:
 - a winning bidder,
 - the price that the winning bidder must pay for the object,
 - and potentially the prices that losing bidders must pay.

Second-price auctions (SPA)

- Assignment rule in SPA prescribes that the object goes to the individual who submitted the highest bid.
- Payment rule says that the winner must pay the second-highest bid while all other bidders pay zero.
- *Tie-breaking rule*: If two or more bidders submit the highest bid, then the object is randomly assigned among them with equal probabilities.

Second-price auctions (SPA)

- We seek to show that submitting a bid equal to her valuation,

$$b_i(v_i) = v_i$$

is a weakly dominant strategy for every player, i.e., it is a weakly dominant BNE of this game.

- This means that:
 - Regardless of the valuation bidder i assigns to the object, and
 - Independently of her opponents' bids,
- submitting a bid equal to her valuation, $b_i(v_i) = v_i$, yields an expected profit equal to or higher than deviating to:
 - A bid lower than her valuation, $b_i(v_i) < v_i$, or
 - A bid higher than her valuation, $b_i(v_i) > v_i$.

Case 1: Bid equal to his valuation

1.1 Bidder i wins.

- Define the highest competing bid among all bidder i 's rivals, h_i , as follows

$$h_i = \max\{b_1, b_2, \dots, b_{i-1}, b_{i+1}, \dots, b_N\}$$

- If the highest competing bid, h_i , lies below bidder i 's bid, $h_i < b_i$:
 - bidder i wins the auction,
 - earning a net payoff of $v_i - h_i$ because in a SPA, the winning bidder does not pay the bid she submitted, but rather the second highest bid, h_i .

1.2 Bidder i loses.

- If, instead, the highest competing bid lies above her bid, $h_i > b_i$, then:
- bidder i loses the auction, earning zero payoff.

1.3 Tie in bids.

- If a tie occurs, where $b_i = h_i$, the object is randomly assigned, and bidder i 's expected payoff becomes $\frac{1}{2}(v_i - h_i)$, where $\frac{1}{2}$ is the probability that bidder i receives the object.
- However, because $v_i = h_i$, in this case, the bidder earns a zero expected payoff.

Case 2: Downward deviations, bidding below her valuation

2.1 If the highest competing bid h_i lies below her bid (i.e., $h_i < b_i < v_i$):

- bidder i still wins the auction,
- earning a payoff $v_i - h_i$, as when she does not shade her bid.
- This indicates that bidder i 's payoff does not increase relative to case 1.1, thus not having incentives to shade her bid.

2.2 If the highest competing bid, h_i , lies between b_i and v_i (i.e., $b_i < h_i \leq v_i$):

- bidder i loses,
- earning a payoff of zero.
- As a consequence, bidder i 's payoff is the same (zero) as when she does not shade her bid, as shown in case 1.2.

2.3 If the highest competing bid h_i is higher than v_i (i.e., $b_i < v_i < h_i$):

- bidder i loses the auction,
- yielding the same outcome as when she submits a bid, $b_i = v_i$.
- In this case, bidder i 's payoff exactly coincides when she bids according to her valuation (case 1.3) and when she shades her bid.

Case 2: Downward deviations, bidding below her valuation

- Overall, when bidder i shades her bid:
 - she earns the same or lower payoff than when she submits a bid that coincides with her valuation for the object.
- In other words, she does not have strict incentives to shade her bid:
 - her payoff would not strictly improve from doing so,
 - regardless of the exact position of the highest competing bid.

Case 3: Upward deviations, bidding above her valuation

3.1 If the highest competing bid h_i lies below bidder i 's valuation, v_i :

- she wins,
- earning a payoff of $v_i - h_i$.
- Her payoff in this case coincides with that when submitting a bid equal to her valuation, $b_i = v_i$, as in case 1.1, implying that she has no strict incentives to bid above her valuation.

3.2 If the highest competing bid h_i lies between v_i and b_i ($v_i < h_i < b_i$):

- bidder i wins the object,
- but earns a negative payoff because $v_i - h_i < 0$.
- If, instead, bidder i submits a bid equal to her valuation, $b_i = v_i$, she would have lost the object, earning a zero payoff.
- In other words, bidder i would be better off submitting a bid $b_i = v_i$, and losing, than submitting a bid $b_i > v_i$, and winning but earning a negative payoff.

3.3 If the highest competing bid h_i lies above b_i ($v_i < b_i < h_i$):

- bidder i loses the auction,
- earning a zero payoff.
- When she submits a bid equal to her valuation, $b_i = v_i$, in case 1.3, a tie occurs, but her expected payoff $\frac{1}{2}(v_i - h_i)$, is zero given that $v_i = h_i$ in that case. Therefore, bidder i has no strict incentives to deviate from bidding $b_i = v_i$.

Case 3: Upward deviations, bidding above her valuation

- Overall, bidder i can earn the same payoff as in cases 1.1-1.3, or a lower payoff,
 - but cannot strictly improve her payoff.
- In short, bidding according to her valuation $b_i = v_i$, is a weakly dominant strategy for every bidder i in the SPA.

Discussion

- **Bidding BNE vs. Dominant Strategies.**

- Every bidder i finds that bidding according to her valuation is $\tilde{N}$
 - not only a BNE of the SPA,
 - it is the BNE where every player uses weakly dominant strategies.
- If bidder i finds that her bidding function $b_i^* = b_i(v_i)$ is a BNE, this means that

$$EU_i(b_i^*, b_{-i}^* | v_i) \geq EU_i(b_i, b_{-i}^* | v_i)$$

for every bid $b_i \neq b_i^*$ and every valuation v_i .

- Intuitively, equilibrium bid b_i^* provides bidder i with a higher expected payoff than any other bid $b_i \neq b_i^*$, conditional on her rivals selecting equilibrium bids b_{-i}^* .
- But we showed something stronger than that! (Next slide.)

Discussion

- **Bidding BNE vs. Dominant Strategies.**

- When we say that bidder i finds that her equilibrium bid b_i^* is a *weakly dominant strategy*, we mean that

$$EU_i(b_i^*, b_{-i} | v_i) \geq EU_i(b_i, b_{-i} | v_i)$$

for:

- every bid $b_i \neq b_i^*$
- every valuation v_i , and
- every bidding profile her rivals use b_{-i} .
- This inequality entails that bidder i 's expected payoff from submitting bid b_i^* is higher than from any other bid $b_i \neq b_i^*$, *regardless* of the specific bidding profile that her rivals use, b_{-i} , that is:
 - both when they submit equilibrium bids, $b_{-i} = b_{-i}^*$,
 - and when they do not, $b_{-i} \neq b_{-i}^*$.
- This is a strong property in the bidding strategy in SPAs, saying that bidder i , when submitting $b_i(v_i) = v_i$, can essentially, ignore her opponents' bids:
 - both when they submit equilibrium bids and when they do not.

Discussion

- **No Bid shading.**

- Intuitively, by shading her bid, $b_i(v_i) < v_i$, bidder i :
 - lowers the chance that she wins the auction, but...
 - does not lower the price that she pays upon winning.
 - No tradeoff!
- In the FPA, in contrast, bid shading gives rise to a trade-off:
 - A lower chance of winning the auction,
 - But paying a lower price upon winning the object.

First-price auctions (FPA)

- Assignment rule coincides with that in the SPA:
 - the winner is the bidder submitting the highest bid.
- Payment rule in the FPA, however, differs:
 - the winning bidder must pay the highest bid.
- This seemingly small difference between both auction formats give rise to bid shading in the FPA
 - a result that we could not sustain in equilibrium when bidders face a SPA.

First-price auctions (FPA)

Step 1. Writing bidder i 's maximization problem.

- Bidder i 's maximization problem is as follows:

$$\max_{b_i \geq 0} \quad Pr(win) \times \underbrace{(v_i - b_i)}_{\substack{\text{net payoff that} \\ \text{bidder } i \text{ earns} \\ \text{when she wins}}} - Pr(lose) \times \underbrace{0}_{\substack{\text{payoff} \\ \text{from losing}}}$$

- At this point, we need to write the probability of winning, $Pr(win)$, as a function of bidder i 's bid, b_i .
- To do this, note that every bidder i uses a symmetric bidding function $b_i: [0,1] \rightarrow \mathbb{R}_+$, a function mapping her valuation $v_i \in [0,1]$ into a positive dollar amount (her bid).

First-price auctions (FPA)

- If bidding functions are symmetric across players and monotonic, bidder i wins when her bid satisfies $b_j \leq b_i$,
 - which must indicate that her valuation satisfies $v_j \leq v_i$.
 - This ranking between valuations v_i and v_j occurs if $\Pr(v_j \leq v_i) = F(v_i)$.
- Therefore, when bidder i faces $N - 1$ rivals, her probability of winning the auction is:
 - the probability that her valuation exceeds that of all other $N - 1$ rivals.
- Since valuations are i.i.d., we can write this probability as the product

$$\begin{aligned} & \Pr(v_j \leq v_i) \times \Pr(v_k \leq v_i) \times \dots \times \Pr(v_l \leq v_i) \\ &= \underbrace{F(v_i) \times F(v_i) \times \dots \times F(v_i)}_{N-1 \text{ times}} = F(v_i)^{N-1} \end{aligned}$$

where bidders $j \neq k \neq l$ represent i 's rivals.

First-price auctions (FPA)

As a result, we can express the above expected utility maximization problem as follows:

$$\max_{b_i \geq 0} \underbrace{F(v_i)^{N-1}}_{Pr(win)} (v_i - b_i)$$

First-price auctions (FPA)

- Using the above bidding function, we can write
 - $b_i(v_i) = x_i$, where $x_i \in \mathbb{R}_+$ is the bidder i 's bid when her valuation is v_i .
 - Or $v_i = b_i^{-1}(x_i)$ by applying the inverse of $b_i(\cdot)$ on both sides.
- So, the program becomes

$$\max_{x_i \geq 0} F \left(b_i^{-1}(x_i) \right)^{N-1} (v_i - x_i)$$

First-price auctions (FPA)

Step 2. Finding equilibrium bids.

- Differentiating with respect to x_i , yields

$$-\left[F\left(b_i^{-1}(x_i)\right)^{N-1}\right] + (N-1)F\left(b_i^{-1}(x_i)\right)^{N-2}f\left(b_i^{-1}(x_i)\right)\frac{\partial b_i^{-1}(x_i)}{\partial x_i}(v_i - x_i) = 0$$

- Since $b_i^{-1}(x_i) = v_i$, we can use the inverse function theorem to obtain $\frac{\partial b_i^{-1}(x_i)}{\partial x_i} = \frac{1}{b'(b_i^{-1}(x_i))}$. So simplifying and rearranging,

$$F(v_i)^{N-1}b'(v_i) + (N-1)F(v_i)^{N-2}f(v_i)x_i = (N-1)F(v_i)^{N-2}f(v_i)v_i$$

- The left-hand side is $\frac{\partial[F(v_i)^{N-1}b_i(v_i)]}{\partial v_i}$, which let us write the above expression more compactly as

$$\frac{\partial[F(v_i)^{N-1}b_i(v_i)]}{\partial v_i} = (N-1)F(v_i)^{N-2}f(v_i)v_i$$

First-price auctions (FPA)

Step 2. Finding equilibrium bids.

- Integrating both sides of $\frac{\partial [F(v_i)^{N-1} b_i(v_i)]}{\partial v_i} = (N-1)F(v_i)^{N-2} f(v_i) v_i$, yields

$$F(v_i)^{N-1} b_i(v_i) = \int_0^{v_i} (N-1) F(v_i)^{N-2} f(v_i) v_i dv_i$$

- We could now solve for $b_i(v_i)$ and obtain an equilibrium bidding function

$$b_i(v_i) = \frac{N-1}{F(v_i)^{N-1}} \int_0^{v_i} F(v_i)^{N-2} f(v_i) v_i dv_i$$

- But this presentation does not help us see the role of bid shading.
 - For that, we need to represent the bid as a function of bidder i 's valuation, v_i , minus one term capturing bid shading.
 - For that, we need to apply integration by parts on the RHS. (Next slide.)

First-price auctions (FPA)

Step 3. Applying integration by parts.

- Recall that

$$g(x)h(x) = \int g'(x)h(x)dx + \int g(x)h'(x)dx$$

- Reordering this expression, we find

$$\int g'(x)h(x)dx = g(x)h(x) - \int g(x)h'(x)dx$$

- Applying this in the above expression (RHS), yields:

$$\int_0^{v_i} \underbrace{[(N-1)F(x)^{N-2}f(x)]}_{g'(x)} \underbrace{x}_{h(x)} dx = \underbrace{F(v_i)^{N-1}}_{g(x)} \underbrace{v_i}_{h(x)} - \int_0^{v_i} \underbrace{F(x)^{N-1}}_{g(x)} \underbrace{1}_{h'(x)} dx$$

First-price auctions (FPA)

Step 3. Applying integration by parts.

- Inserting this results in the right-hand side of first-order condition, yields

$$F(v_i)^{N-1} b_i(v_i) = F(v_i)^{N-1} v_i - \int_0^{v_i} F(v_i)^{N-1} dv_i$$

- Solving for $b_i(v_i)$, we obtain the equilibrium bid in the FPA:

$$b_i(v_i) = v_i - \underbrace{\frac{\int_0^{v_i} F(v_i)^{N-1} dv_i}{F(v_i)^{N-1}}}_{\text{Bid shading}}$$

- Intuitively, bidder i submits a bid:
 - equal to her valuation for the object, v_i ,
 - less an amount captured by the second term, which we refer as her “bid shading.”

First-price auctions (FPA)

Step 4. Checking Monotonicity

- We finally check that the above bidding function $b_i(v_i)$ is monotonically increasing in bidder i 's valuation, v_i . A marginal increase in v_i produces the following effect in bidder i 's equilibrium bidding function:

$$\begin{aligned}\frac{\partial b_i(v_i)}{\partial v_i} &= 1 - \frac{F(v_i)^{N-1}F(v_i)^{N-1} - (N-1)F(v_i)^{N-2} \int_0^{v_i} F(v_i)^{N-1} dv_i}{[F(v_i)^{N-1}]^2} \\ &= \frac{(N-1)F(v_i)^{N-2}f(v_i) \int_0^{v_i} F(v_i)^{N-1} dv_i}{[F(v_i)^{N-1}]^2}\end{aligned}$$

which is positive since $F(v_i) \in [0,1]$, $f(v_i) > 0$ for all v_i and $N \geq 2$ by definition.

But this effect is less than proportional (See next slide).

First-price auctions (FPA)

Step 4. Checking Monotonicity

- Note also that $\frac{\partial b_i(v_i)}{\partial v_i} = \frac{(N-1)F(v_i)^{N-2}f(v_i) \int_0^{v_i} F(v_i)^{N-1} dv_i}{[F(v_i)^{N-1}]^2} < 1$:

$$\begin{aligned} \Rightarrow \frac{(N-1)f(v_i) \int_0^{v_i} F(v_i)^{N-1} dv_i}{F(v_i)^N} &= (N-1) \frac{f(v_i)}{F(v_i)} \times \frac{\int_0^{v_i} F(v_i)^{N-1} dv_i}{F(v_i)^{N-1}} \\ &\Rightarrow (N-1) \times \frac{d \log F(v_i)}{dv_i} \times \frac{\int_0^{v_i} F(v_i)^{N-1} dv_i}{F(v_i)^{N-1}} \end{aligned}$$

- Let $\log F(v_i) = g(v_i)$, then we have:

$$\begin{aligned} \Rightarrow (N-1) \times \frac{dg(v_i)}{dv_i} \times \frac{\int_0^{v_i} e^{(N-1)g(v_i)} dv_i}{e^{(N-1)g(v_i)}} &= (N-1) \frac{\int_0^{v_i} e^{(N-1)g(v_i)} dg(v_i)}{e^{(N-1)g(v_i)}} \\ \Rightarrow \left. \frac{\left(\frac{1}{N-1}\right) e^{(N-1)g(v_i)}}{e^{(N-1)g(v_i)}} \right|_0^{v_i} \times (N-1) &= \frac{e^{(N-1)g(v_i)} - e^{(N-1)g(0)}}{e^{(N-1)g(v_i)}} < \frac{e^{(N-1)g(v_i)}}{e^{(N-1)g(v_i)}} = 1. \end{aligned}$$

Meaning that an increase in bidder i 's valuations leads her to increase her bid, but less than proportionally.

Example 9.1. FPA with uniformly distributed valuations

- Consider, for instance, when individual valuations are uniformly distributed, i.e. $F(v_i) = v_i$.
- In this setting, we obtain $F(v_i)^{N-1} = v_i^{N-1}$ and $\int_0^{v_i} F(v_i)^{N-1} dv_i = \frac{1}{N} v_i^N$, producing a bidding function of

$$b_i(v_i) = v_i - \frac{\frac{1}{N} v_i^N}{v_i^{N-1}} = v_i - \underbrace{\frac{v_i}{N}}_{\text{Bid shading}} = v_i \left(\frac{N-1}{N} \right)$$

- In this context, every bidder shades her bid by $\frac{v_i}{N}$, which increases in the number of competing bidders.
- In addition, the equilibrium bidding function $b_i(v_i) = v_i \left(\frac{N-1}{N} \right)$ is monotonically increase in the valuation that bidder i assigns to the object, v_i , as required.

Example 9.1. FPA with uniformly distributed valuations

- When only 2 bidders compete for the object, $N = 2$, this bidding function simplifies to $b_i(v_i) = \frac{v_i}{2}$ as depicted in Figure 9.1.
- When $N = 3$, equilibrium bids increase to $b_i(v_i) = \frac{2v_i}{3}$.
- When $N = 4$, equilibrium bids increase to $b_i(v_i) = \frac{3v_i}{4}$.




Figure 9.1. Equilibrium bids in the first-price auction with uniformly distributed valuations.

Example 9.1. FPA with uniformly distributed valuations

- Informally, as more bidders participate in the auction, every bidder i submits more aggressive bids since:
 - she faces a higher probability than another bidder j has a higher valuation for the object.
 - And, given symmetric bidding functions, bidder j submits a higher bid than she does, leading to bidder i to lose the auction.
- In the extreme case that $N \rightarrow \infty$, the bidding function converges to $b_i(v_i) = v_i$.





Figure 9.1. Equilibrium bids in the first-price auction with uniformly distributed valuations.

Efficiency in auctions

- An auction is deemed “efficient” when it assigns the object to the individual with the highest valuation.
 - That is, when the assignment rule allocates the object to bidder i if only if her valuation, v_i , satisfies $v_i > v_j$ for all $j \neq i$.
- Otherwise, if bidder j receives the object despite having a lower valuation than bidder i :
 - these two bidders could negotiate at the end of the auction,
 - with bidder i paying a price p that satisfies $v_i > p > v_j$,
 - making both bidders better off.
- The assignment rule that allocates the object to bidder j is Pareto inefficient.
 - We can find an alternative allocation that improves the payoff of at least one individual without making any other individual worse off.
- If bidders use a *symmetric, strictly increasing, bidding function* in equilibrium:
 - the winner of the auction must be the individual with the highest valuation,
 - making the auction efficient.


Efficiency in auctions

- *Symmetric and strictly increasing* bidding function:
 - Because bidder A has a higher valuation than B does, he submits a higher bid and wins the object.




Efficiency in auctions

- *Symmetric* but *Not strictly increasing* bidding function:
 - Bidder A has a higher valuation than B does, but he may submit a lower bid and lose the object.



Efficiency in auctions

- *Asymmetric* but *strictly increasing* bidding functions:
 - Bidder A has a higher valuation than B does, but he may submit a lower bid and lose the object.




Efficiency in auctions

- Two typical features that “break” efficiency in auctions:
 - Budget constraints, either symmetric or asymmetric.
 - Asymmetric risk averse bidders.

Budget constraints

- When bidder i faces a budget constraint w_i ,
 - her bidding function is increasing in v_i for all $b_i < w_i$ (affordable bids).
 - But becomes flat at the height of w_i for all valuations for which $b_i > w_i$ (unaffordable bids)




Budget constraints

- Budget constraints imply that bidding functions are weakly increasing in v_i , not strictly increasing,
 - so we cannot guarantee that the auction is efficient.
 - Recall that we need bidders to use “*symmetric, strictly increasing, bidding functions* in equilibrium”
- Figure to illustrate this point (next slide).

Budget constraints

- Bidder A values the object more than B does...
 - yet he may submit a lower bid, losing the object.



Budget constraints

- Budget constraints imply that bidding functions are weakly increasing in v_i , not strictly increasing,
 - so we cannot guarantee that the auction is efficient.
 - Recall that we need bidders to use “*symmetric, strictly increasing, bidding functions* in equilibrium”
- This occurs even if:
 - only one individual suffers from budget constraints,
 - or if all bidders face the same budget constraint $w_i = w$,
 - making their bidding functions symmetric (but not strictly increasing).

Asymmetric risk averse bidders

- Another context where efficiency is not satisfied is that where players exhibit different degrees of risk aversion.
- Bidders competing in a FPA submit more aggressive bids when they become more risk-averse.
 - If both bidders are symmetric in their risk aversion, the above condition hold (i.e. bidders use a symmetric, strictly increasing, bidding function).
- When bidders are asymmetric in their risk preferences, however, the FPA is not efficient.

Asymmetric risk averse bidders

- Consider FPA with two bidders, A and B , where bidder A exhibits more (less) risk aversion.
- Scenarios:
 - Bidder A can be risk averse, while bidder B is risk neutral or risk-loving,
 - both bidders can be risk-averse but A is more risk averse than B ,
 - or both are risk lovers but A is less so than B .
- In any of these settings, bidder A submits a more aggressive bid than B does, $b_A > b_B$, implying that A wins the auction.
- If bidder A values the object more than B does, $v_A > v_B$, the outcome of the auction is still efficient; but otherwise the outcome is inefficient.
- Generally, then, we cannot guarantee that the object goes to the individual who values the object the most, especially if bidders are relatively asymmetric in their risk aversion, entailing that the FPA is not efficient when players are risk averse.

Seller's expected revenue: Expected revenue in the FPA

Step 1. *Finding each bidder's payment*

- The seller receives a payment from bidder i if she wins the auction. In other words, bidder i 's payment is

$$m(v_i) = \Pr(\text{win}) \times b_i(v_i)$$

- We know that $\Pr(\text{win}) = [F(v_i)]^{N-1}$, so

$$m(v_i) = [F(v_i)]^{N-1} \times b_i(v_i) = G(v_i) \times b_i(v_i), \text{ where } G(v_i) = [F(v_i)]^{N-1}$$

- Recall that the equilibrium bidding function in the FPA is

$$b_i(v_i) = v_i - \frac{\int_0^{v_i} F(x)^{N-1} dx}{F(v_i)^{N-1}} = v_i - \frac{\int_0^{v_i} G(x) dx}{G(v_i)} = \frac{G(v_i)v_i - \int_0^{v_i} G(x) dx}{G(v_i)}$$

Seller's expected revenue: Expected revenue in the FPA

Step 1. *Finding each bidder's payment*

- Applying integration by parts in the numerator, we obtain

$$G(v_i)v_i - \int_0^{v_i} G(x) dx = \int_0^{v_i} xg(x) dx$$

- So we can rewrite the equilibrium bidding function:

$$b_i(v_i) = \frac{\int_0^{v_i} xg(x) dx}{G(v_i)}$$

- Inserting this equilibrium bidding function in bidder i 's expected payment, yields

$$m(v_i) = G(v_i) \times \overbrace{\left(\frac{\int_0^{v_i} xg(x) dx}{G(v_i)} \right)}^{\text{Bid, } b_i(v_i)} = \int_0^{v_i} xg(x) dx$$

Seller's expected revenue: Expected revenue in the FPA

Step 2. *Finding the expected payment.*

- Since the seller cannot observe bidder i 's value for the object, she needs to take expectations over all possible values to find the expected payment from this bidder, $E[m(v_i)]$, as follows:

$$E[m(v_i)] = \int_0^1 m(x) f(x) dx = \int_0^1 \underbrace{\left[\int_0^{v_i} x g(x) dx \right]}_{m(v_i)} f(x) dx$$

Step 3. *Sum across all bidders.*

- The seller sums across all N bidders participating in the auction, which yields the expression of her revenue in the FPA.

$$R^{FPA} = \sum_{i=1}^N E[\pi_i(v_i)] = N \times E[m(v_i)] = N \int_0^1 \left[\int_0^{v_i} x g(x) dx \right] f(x) dx$$

Example 9.2. Expected Revenue in FPA with Uniformly Distributed Valuations

- When valuations are uniformly distributed, $F(v_i) = v_i$, we obtain that:
 - $f(v_i) = F'(v_i) = 1$
 - $G(v_i) = [F(v_i)]^{N-1} = v_i^{N-1}$, which implies that:
 - $g(v_i) = G'(v_i) = (N-1)v_i^{N-2}$,
 - $v_i g(v_i) = v_i (N-1)v_i^{N-2} = (N-1)v_i^{N-1}$,

- Therefore, bidder i 's equilibrium bid is

$$b_i(v_i) = v_i \left(\frac{N-1}{N} \right)$$

- **Step 1.** The expected payment from bidder i is:

$$m(v_i) = \int_0^{v_i} xg(x)dx = \int_0^{v_i} \underbrace{(N-1)x^{N-1}}_{xg(x)} dx = \frac{N-1}{N} [x^N]_0^{v_i} = \frac{N-1}{N} (v_i^N - 0) = \frac{N-1}{N} v_i^N$$

Example 9.2. Expected Revenue in FPA with Uniformly Distributed Valuations

- **Step 2.** Therefore, bidder i 's expected payment is:

$$E[m(v_i)] = \int_0^1 m(v_i) f(v_i) dv_i = \int_0^1 \underbrace{\frac{N-1}{N} v_i^N}_{m(v_i)} \underbrace{1}_{f(v_i)} dv_i = \frac{N-1}{N} \int_0^1 v_i^N dv_i = \frac{N-1}{N} \left[\frac{v_i^{N+1}}{N+1} \right]_0^1 = \frac{N-1}{N(N+1)}$$

- **Step 3.** Finally, the seller sums across all N bidders to obtain the expected revenue from the FPA, R^{FPA} , as follows,

$$R^{FPA} = \sum_{i=1}^N E[m(v_i)] = N \times E[m(v_i)] = N \frac{N-1}{N(N+1)} = \frac{N-1}{N+1}$$

Example 9.2. Expected Revenue in FPA with Uniformly Distributed Valuations

- Therefore,

$$R^{FPA} = \frac{N - 1}{N + 1}$$

- The expected revenue is:
 - increasing in the number of bidders, N ,
 - but at a decreasing rate.
 - Approaches 1 when $N \rightarrow \infty$.
- This result goes in line with that in Example 9.1:
 - as more bidders compete in the auction, they submit more aggressive bids, i.e., $b_i(v_i)$ increases in N ,
 - increasing as a result the expected winning bid that the seller earns.

Expected Revenue in the SPA

- In this auction format, the seller anticipates that every bidder i submits a bid $b_i(v_i) = v_i$.
- The winning bidder pays the second-highest bid.
- The second-highest bid coincides with the second-highest valuation for the object, $v_1^{[2]}$. That is,

$$R^{SPA} = E \left[v_1^{[2]} \right] = \int_0^1 x f^{[2]}(x) dx$$

Expected Revenue in the SPA

- We first identify the cumulative distribution function $F^2(x) = Pr(v^2 \leq x)$ which happens when two events occur:
 1. The valuations of all N bidders are below x , or formally, $v_i \leq x$ for every bidder i . This event happens with probability
$$Pr\{v_1 \leq x\} \times \dots \times Pr\{v_N \leq x\} = \underbrace{F(x) \times \dots \times F(x)}_{N \text{ times}} = [F(x)]^N$$
 2. The valuations of $N - 1$ bidders are below x , $v_i \leq x$, but that of only one bidder j is above x , $v_i > x$. This even can occur in N different ways:
 - $v_1 > x$ for bidder 1 but $v_i \leq x$ for every bidder $i \neq 1$;
 - $v_2 > x$ for bidder 2 but $v_i \leq x$ for every bidder $i \neq 2$;
 - ...
 - $v_N > x$ for bidder N but $v_i \leq x$ for every bidder $i \neq N$.

Expected Revenue in the SPA

- Each of these N cases happens with probability

$$\underbrace{[1 - F(x)]}_{v_i > x \text{ for } i} \times \underbrace{[F(x)]^{N-1}}_{v_j \leq x \text{ for every } j \neq i}$$

where $[1 - F(x)]$ denotes the probability that $v_i > x$ for a given bidder i , while $[F(x)]^{N-1}$ represents the probability that $v_j \leq x$ for all other bidders $j \neq i$.

- Summing over the above N cases, we find that event (2) happens with probability

$$\sum_{i=1}^N [1 - F(x)][F(x)]^{N-1} = N(1 - F(x)) [F(x)]^{N-1}$$

- Summarizing, the cumulative distribution function of the second-highest valuation $F^{[2]}(x)$, is

$$F^{[2]}(x) = [F(x)]^N + N(1 - F(x))[F(x)]^{N-1}$$

Expected Revenue in the SPA

- Rearranging,

$$F^{[2]}(x) = N[F(x)]^{N-1} - (N-1)[F(x)]^N$$

- Differentiating with respect to x :

$$f^{[2]}(x) = N(N-1)F(x)^{N-2}[1-F(x)]f(x)$$

- Inserting density function into seller's expected revenue

$$R^{SPA} = \int x f^{[2]}(x) dx = \int x \underbrace{N(N-1)F(x)^{N-2}[1-F(x)]f(x)}_{f^{[2]}(x)} dx$$

Example 9.3. Expected Revenue in SPA with uniformly distributed valuations

- When valuations are uniformly distributed, $F(x) = x$ and $f(x) = 1$, the seller's expected revenue becomes

$$\begin{aligned} R^{SPA} &= \int_0^1 xN(N-1)x^{N-2}[1-x]dx \\ &= (N-1) \int_0^1 Nx^{N-1}dx - N(N-1) \int_0^1 x^N dx \\ &= (N-1)x^N \Big|_0^1 - \frac{N(N-1)x^{N+1}}{N+1} \Big|_0^1 \\ &= (N-1) - \frac{N(N-1)}{N+1} \\ &= \frac{N-1}{N+1} \end{aligned}$$

which coincides with that in the FPA, R^{SPA} , found in Example 9.2., thus being increasing and concave in the number of bidders, N .

Revenue Equivalence Principle

- When bidders' valuations are uniformly distributed, $R^{FPA} = R^{SPA}$, implying that the seller can expect to earn the same revenue from both auction formats.
- This “revenue equivalence” result extends to several other auction formats:
 - yielding the same expected revenue as the FPA and SPA,
 - and to settings where bidders' valuations are non-uniformly distributed.
- We can identify the two main requirements that two auction formats must satisfy to yield the same expected revenue for the seller:
 1. Same allocation rule in both auction formats, e.g., the bidder submitting the highest bid receives the object, as in the FPA and SPA.
 2. Same expected utility of the bidder who has the lowest valuation for the object, e.g., zero in most auction formats since this bidder loses the auction, not receiving the object.

Revenue Equivalence Principle

- It is straightforward to note that the comparison of FPA and SPA satisfies conditions (1) and (2), thus generating the same expected revenue.
- But the comparison of FPA and APA does not satisfy condition (2):
 - because the bidder with the lowest valuations earns a zero payoff in the FPA but a negative payoff in the APA after paying here bid.
- Therefore, the FPA and APA do not necessarily generate the same expected revenue.
- Similarly, the comparison of the FPA and the lottery auction does not satisfy condition (1),
 - since the lottery auction does not necessarily assign the object to the individual submitting the highest bid in the FPA.
- As a consequence, the FPA and the lottery auction do not yield the same expected revenue for the seller.

Common value auctions and the winner's curse

- Every bidder i shares a common value for the object, v , but privately observes a noisy signal s_i about the object's value drawn from $F(s_i)$, where $s_i \in [0,1]$.
- Based on this signal, every bidder i submits her bid, b_i .
- This setting is known as “common value” auctions
- Bidders participate in a first-price, sealed-bid auction:
 - if bidder i wins, her realized payoff becomes $v - b_i$,
 - and if she loses her payoff is zero.
- Experimentally tested with a jar of nickels (Explain).
- For simplicity, we consider a setting with only two bidders, and that the true value is equal to the average of bidders' signals, so that $v = \frac{s_i + s_j}{2}$.

Bid shading is a must!

- Note that bidder i falls prey of the winner's curse if her bid exceeds the object's true value (which no bidder observes), $b_i > v$,
 - thus earning a negative payoff, $v - b_i < 0$, from winning the auction.
- In particular, this occurs if

$$b_i > \frac{s_i + s_j}{2}$$

- When bidder i 's bid is a function of her privately observed signal, $b_i = \alpha s_i$, where $\alpha \in [0,1]$, this inequality becomes

$$\alpha s_i > \frac{s_i + s_j}{2} \Rightarrow s_i > \frac{s_j}{2\alpha - 1}$$

- When bidder i submits a bid equal to the signal she received, the winner's curse occurs if $s_i > s_j$.

Bid shading is a must!

- Intuition:
 - If every bidder submits a bid that coincides with her privately observed signal, the bidder who received the highest signal ends up submitting the highest bid,
 - She wins the auction
 - But suffers from the *winner's curse*.
- In other words, the fact that she won the auction means that she received an overestimated signal of the object.
- If instead, bidder i submits a bid equal to $\frac{3}{4}$ of the signal she received, $\alpha = \frac{3}{4}$, the winner's curse only emerges if $s_i > 2s_j$;
 - that is, when bidder i 's signal is larger than the double of bidder j 's.
- More generally, as bidder i shades her bid more severely (decreasing α),
 - Ratio $\frac{s_j}{2\alpha-1}$ increases,
 - and the winner's curse is less likely to occur.

Equilibrium Bidding in Common Value Auctions

Step 1. *Finding the expected utility.*

- Bidder i 's expected payoff from participating in the auction is

$$Pr(b_i > b_j) \times \{E[v|s_i, b_i > b_j] - b_i\}$$

where $E[v|s_i, b_i > b_j]$ is bidder i 's expected valuation, conditional on her signal s_i , and on knowing that she submitted the highest bid, i.e., $b_i > b_j$.

- Since bidder j uses bidding function $b_j = \alpha s_j$, as this expression becomes

$$Pr(b_i > \alpha s_j) \times \{E[v|s_i, b_i > \alpha s_j] - b_i\}$$

- Solving for s_j in the probability (first term) and in the inequality inside the expectation operator (second term), we obtain that

$$Pr\left(\frac{b_i}{\alpha} > s_j\right) \times \left\{E\left[v|s_i, \frac{b_i}{\alpha} > s_j\right] - b_i\right\}$$

Equilibrium Bidding in Common Value Auctions

- Step 1. *Finding the expected utility.*

- Next, inserting $v = \frac{s_i + s_j}{2}$ in the expectation operator, yields

$$Pr\left(\frac{b_i}{\alpha} > s_j\right) \times \left\{ E\left[\frac{s_i + s_j}{2} \mid \frac{b_i}{\alpha} > s_j\right] - b_i \right\}$$

- Recall that bidder i observes her signal, s_i , so that $E[s_i] = s_i$, but does not know her rival's, s_j , entailing that

$$Pr\left(\frac{b_i}{\alpha} > s_j\right) \times \left\{ \frac{s_i}{2} + \frac{1}{2} E\left[s_j \mid \frac{b_i}{\alpha} > s_j\right] - b_i \right\}$$

- Using the uniform distribution on s_j , we have

$$Pr\left(\frac{b_i}{\alpha} > s_j\right) = \frac{b_i}{\alpha} \quad \text{and} \quad E\left[s_j \mid \frac{b_i}{\alpha} > s_j\right] = \frac{b_i}{2\alpha}$$

because s_j , which is a uniformly distributed random variable, falls into the range $\left[0, \frac{b_i}{\alpha}\right]$, yielding an expected value of $\frac{b_i}{2\alpha}$.

- Inserting these results into bidder i 's expected payoff, we obtain that

$$\frac{b_i}{\alpha} \left[\frac{s_i}{2} + \frac{1}{2} \left(\frac{b_i}{2\alpha} \right) - b_i \right]$$

Equilibrium Bidding in Common Value Auctions

Step 2. *Taking first-order conditions.*

- Every bidder i chooses her bid b_i to maximize her expected utility, solving the following problem

$$\max_{b_i \geq 0} \quad \frac{b_i}{\alpha} \left[\frac{s_i}{2} + \frac{1}{2} \left(\frac{b_i}{2\alpha} \right) - b_i \right]$$

Taking the first-order conditions with respect to b_i , we obtain

$$\frac{1}{\alpha} \left[\frac{s_i}{2} + \frac{1}{2} \left(\frac{b_i}{2\alpha} \right) - b_i \right] + \frac{b_i}{\alpha} \left(\frac{1}{4\alpha} - 1 \right) = 0$$

Simplifying, yields

$$\frac{1 - 4\alpha}{2\alpha^2} b_i = -\frac{s_i}{2\alpha} \quad \Rightarrow \quad b_i = \frac{\alpha}{4\alpha - 1} s_i$$

Equilibrium Bidding in Common Value Auctions

Step 3. *Finding the equilibrium bidding function.*

- Recall that we considered a symmetric bidding function $b_i = \alpha s_i$ for a generic value of α .
- The above expression, $b_i = \frac{\alpha}{4\alpha-1} s_i$, is indeed linear in signal s_i , so we can write

$$\alpha s_i = \frac{\alpha}{4\alpha-1} s_i \quad \text{or}$$
$$\alpha = \frac{\alpha}{4\alpha-1} \Rightarrow 4\alpha - 1 = 1 \Rightarrow \alpha = \frac{1}{2}$$

- In summary, a symmetric BNE has every bidder i using the bidding function

$$b_i(s_i) = \frac{1}{2} s_i$$

That is, she submits a bid equal to half of her privately observed signal, s_i .

Equilibrium Bidding in Common Value Auctions

- **Extension to N bidders:**

$$b_i(s_i) = \frac{(N+2)(N-1)}{2N^2} s_i$$

where $(N+2)(N-1) < 2N^2$ simplifies to $N^2 - N + 2 > 0$, which holds since $N \geq 2$ by definition.

- In addition, $\frac{\partial b_i(s_i)}{\partial n} = \frac{(4-N)}{2N^3} s_i$, which is positive for all $n < 4$, but negative otherwise.

Equilibrium Bidding in Common Value Auctions

- A natural question at this point is whether this bidding function helps bidders avoid the winner's curse.
 - Exercise 9.16 helps you confirm that it does.
 - *Hint:* We only need to evaluate the winning bidder's utility at her equilibrium bid. No need to compute expected utility (we know she won).

Chapter 10: Perfect Bayesian Equilibrium

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- In this chapter, we extend the analysis of incomplete information, allowing for sequential interaction.
- Examples:
 - Incumbent monopolist privately observing its production costs while a potential entrant only having an estimate about the incumbent's costs
 - Incumbent chooses its prices and the potential entrant, observing the price but not knowing the incumbent's cost, decides whether to enter the industry or not
- An important issue in sequential-move games of incomplete information is that player's actions can:
 - Convey information to other players acting later in the game, i.e., informative signals.
 - Conceal information from other players acting later in the game, i.e., uninformative signals.

Sequential-move games of incomplete information - Notation

1. Nature.

- Reveals to player i a piece of information, $\theta_i \in \Theta_i$, which we refer to as her “type”.
- Examples:
 - A firm’s production cost,
 - its product quality, or
 - a worker’s ability on a certain task.
- The realization of θ_i is
 - Privately observed by player i
 - Not observed by her rivals $j \neq i$.

Sequential-move games of incomplete information - Notation

2. First mover (sender). Player i privately observes θ_i and chooses an action $s_i(\theta_i)$. Afterwards, her rivals observe $s_i(\theta_i)$ but don't observe θ_i .

a. *Separating strategy.* If action $s_i(\theta_i)$ is type-dependent,

$$s_i(\theta_i) \neq s_i(\theta'_i), \text{ for every two types } \theta_i \neq \theta'_i,$$

we say that player i uses a “separating” strategy, i.e., different types choose different strategies.

- *Example:* An incumbent firm uses a difference price depending on its cost.

b. *Pooling strategy.* If action $s_i(\theta_i)$ is type-independent,

$$s_i(\theta_i) = s_i(\theta'_i), \text{ for every two types } \theta_i \neq \theta'_i,$$

we say that player i employs a “pooling” strategy, since different types choose the same strategy (i.e., they pool into the same strategy).

- *Example:* Following with the above example, the incumbent would use the same price regardless of its production cost.

Sequential-move games of incomplete information - Notation

2. First mover (sender). Continues...

- c. *Partially separating strategy.* If action $s_i(\theta_i)$ satisfies
- $$s_i(\theta_i) \neq s_i(\theta'_i) \text{ for at least two types } \theta_i \text{ and } \theta'_i,$$
- $$\text{but } s_i(\theta''_i) = s_i(\theta'''_i) \text{ for other types,}$$
- we say that player i uses a “partially separating” strategy.

In our ongoing example, the incumbent chooses:

- the same price p when its production cost is high and medium,
$$p = s_i(H) = s_i(M),$$
- but a different price $p' \neq p$ when its cost is low,
$$p' = s_i(L).$$

Sequential-move games of incomplete information - Notation

3. Second mover (receiver). Player j observes action s_i , but does not know player i 's type, θ_i .

We assume, however, that every player knows the prior probability distribution over types, $\mu(\theta_i)$, for every type $\theta_i \in \Theta_i$.

This probability distribution is well behaved:

- a. $\mu(\theta_i) \in [0,1]$
- b. $\sum_{\theta_i \in \Theta_i} \mu(\theta_i) = 1$ if player i 's type space is discrete, e.g., $\Theta_i = \{L, H\}$; and similarly, $\int_{\theta_i \in \Theta_i} \mu(x) dx = 1$ if her type space is continuous..

For example, when player i has two types, $\Theta_i = \{L, H\}$, the prior probability can be expressed as $\mu(L) = q$ and $\mu(H) = 1 - q$, where $q \in [0,1]$.

When player i 's types are uniformly distributed, then $\Pr\{\theta \leq \theta_i\} = F(\theta_i) = \theta_i$, with corresponding density $f(\theta_i) = 1$.

Sequential-move games of incomplete information - Notation

4. Second mover updates her beliefs.

- Upon observing s_i , player j updates her beliefs about player i 's type being θ_i , which we write as

$$\mu(\theta_i | s_i).$$

- These beliefs are often known as “posterior beliefs” as opposed to the prior probability distribution, $\mu(\theta_i)$, which are referred to as “prior beliefs”.
- For compactness, we will generally use “priors” and “posteriors”, or “prior probability” and “beliefs”.
- Like priors, posteriors satisfy $\mu(\theta_i | s_i) \in [0,1]$ and :
 - $\sum_{\theta_i \in \Theta_i} \mu(\theta_i | s_i) = 1$ when types are discrete, and
 - $\int_{\theta_i \in \Theta_i} \mu_i(x | s_i) dx = 1$ when types are continuous.

Sequential-move games of incomplete information - Notation

4. Second mover updates her beliefs.

- As we discuss below, the PBE solution concept requires players to update their beliefs according to Bayes' rule, as follows,

$$\mu_j(\theta_i | s_i) = \frac{\mu(\theta_i) \Pr(s_i | \theta_i)}{\Pr(s_i)}$$

- This says that the conditional probability that, upon observing action s_i , player i 's type is θ_i , as captured by $\mu_j(\theta_i | s_i)$, is:
 - the (unconditional) probability that player i 's type is θ_i and she chooses action s_i , $\mu(\theta_i) \Pr(s_i | \theta_i)$...
 - given that she chose s_i , $\Pr(s_i)$.
- When $\mu_j(\theta_i | s_i)$ is higher (lower) than the prior probability $\mu(\theta_i)$, the observation of s_i increases (decreases) the chances that player i 's type is θ_i , indicating that action s_i provides information that player j did not initially have in the prior probability distribution $\mu(\theta_i)$.

Sequential-move games of incomplete information - Notation

4. Second mover updates her beliefs.

- In the extreme case that $\mu_j(\theta_i|s_i) = 1$, player j becomes “convinced,” after observing s_i , the player i ’s type is θ_i
- Similarly, if $\mu_j(\theta_i|s_i) = 0$, player j is sure that player i ’s type cannot be θ_i .
- In contrast, if $\mu_j(\theta_i|s_i) = \mu_j(\theta_i)$, player j does not change her beliefs upon observing s_i , meaning that s_i did not provide her with additional information. In this case, we say that the signal was uninformative or that player j could not update her beliefs.

Sequential-move games of incomplete information - Notation

5. Second mover responds.

- After player j updates her beliefs to $\mu_j(\theta_i | s_i)$, she responds choosing an action s_j .
 - Player j may have observed her type, θ_j , before choosing her action s_j .
- Player j 's type can be:
 - Publicly observed by all players, becoming common knowledge.
 - Or privately observed by player j (for instance, every player privately observes her type).

Sequential-move games of incomplete information - Notation

- After concluding the above five steps, most of the sequential-move games with incomplete information that we consider in this chapter will be over, distributing payoffs to each player.
- However, we can continue the game, allowing, for instance, player i to observe player j 's action s_j in step 5, update her beliefs about j 's type θ_j , $\mu_i(\theta_j|s_j)$, and respond with an action a_i , where we use a_i to distinguish it from s_i .
 - If player j does not have privately observed types, then player i could just respond with action a_i , without having to update her beliefs about θ_j .

Example 10.1. Two sequential-move games of incomplete information

- Consider this game tree.
- Only one information set.
- Player 1 privately observes whether:
 - a business opportunity is beneficial for player 2 (in the upper part of the game tree, which happens with probability p), or
 - futile (bottom part of the tree, with probability $1 - p$), where $p \in [0,1]$.




Figure 10.1. Sequential-move game with incomplete information - One information set.

Example 10.1. Two sequential-move games of incomplete information

- In this setting, we say that player 1's types are *Beneficial* or *Futile*.
- If player 1 does not make an investment offer to player 2, N (moving leftward),
 - the game is over and both players earn a payoff of zero.
- If, instead, player 1 makes an investment offer to player 2, O (moving rightward),
 - player 2 receives this offer
 - but does not observe whether it is beneficial (at the upper part of the tree) or not (at the bottom part).




Figure 10.1. Sequential-move game with incomplete information - One information set.

Example 10.1. Two sequential-move games of incomplete information

- However, upon receiving the offer, player 2 must respond:
 - accepting (A)
 - or rejecting it (R).
- The game in Figure 10.1, then, has only one information set where the receiver (player 2) is called to move.




Figure 10.1. Sequential-move game with incomplete information - One information set.

Example 10.1. Two sequential-move games of incomplete information

- This game, however, depicts the case in which a receiver is called to respond after both messages from a sender,
 - yielding two information sets.
- In particular, the game represents Spence's (1973) labor-market signaling game:
 - a worker (sender) privately observed her job productivity (her type), which is either high or low,
 - then decides whether to pursue an advanced education degree (such as Masters program) that she might use as a signal about her productivity level to the employer (receiver).




Figure 10.2. Labor market signaling game.

Example 10.1. Two sequential-move games of incomplete information

- Upon observing whether or not the applicant acquired education...
 - but without observing her true productivity,
 - the employer updates her beliefs about the worker's type,
 - and responds hiring her manager (M) or as a cashier (C).




Figure 10.2. Labor market signaling game.

Example 10.1. Two sequential-move games of incomplete information

- Observing the worker's payoffs, one can infer that the cost of acquiring education is only \$4 for the high-productivity worker.
- Indeed, when she is hired as a M , at the top of the figure:
 - the difference in her payoff when she avoids or acquires education, $10 - 6 = 4$.
- A similar argument applies when she is hired as a C , where the payoff difference is $4 - 0 = 4$.




Figure 10.2. Labor market signaling game.

Example 10.1. Two sequential-move games of incomplete information

- However, the cost of acquiring education increases to \$7 for the low-productivity worker.
- In this case, when she is hired as a M :
 - her payoff difference is $10 - 3 = 7$ and,
 - similarly, when she is hired as a C , the payoff difference is $4 - (-3) = 7$.




Figure 10.2. Labor market signaling game.

Example 10.1. Two sequential-move games of incomplete information

- Regarding the firm's payoffs, in the second component of each payoff pair in the terminal nodes of the figure:
 - they are unaffected by the worker's type, namely, \$4 when the worker is hired as a C , regardless of her type.
 - But are \$10 when the worker is hired as a M , and her productivity happens to be high but \$0 otherwise.




Figure 10.2. Labor market signaling game.

Example 10.1. Two sequential-move games of incomplete information

- The firm's payoffs, are, however, unaffected by whether the worker acquires education or not.
- In other words, education is *not productivity-enhancing*.
- This is, of course, a simplification!
 - Helps us focus on education working as a signal to potential employers...
 - facilitating information transmission from workers to firms,
 - even if it does not change the worker's job productivity.




Figure 10.2. Labor market signaling game.

BNE prescribing sequentially irrational behavior

- Why not apply BNE to solve these games?
 - We can apply the BNE solution concept to find equilibrium behavior when players interact sequentially and operate under incomplete information.
- BNE may identify *sequentially irrational behavior*:
 - That is, upon reaching a node or information set, a player acts in a way that does not maximize her expected payoff.
- BNE exhibits similar issues as the NE solution concept did:
 - We can apply it to both simultaneous- and sequential-move games,
 - but when applying it to the latter we may find too many equilibria, some of them prescribing that players are sequentially irrational.
 - See Example 10.2

Example 10.2. Applying BNE to sequential-move games of incomplete information

		Player 2	
		A	R
Player 1	$O^B O^F$	$\underline{4 - p}, \underline{4p}$	$-3 + p, -3 + 5p$
	$O^B N^F$	$3p, \underline{4p}$	$-2p, 2p$
	$N^B O^F$	$4(1 - p), \underline{0}$	$-3(1 - p), -3(1 - p)$
	$N^B N^F$	$0, \underline{0}$	$\underline{0}, \underline{0}$

Matrix 10.1. Bayesian normal-form representation of the game tree in Figure 10.1

- Consider the game in Figure 10.1 again (only one info. set).
- To find the BNE, we first need to represent its Bayesian normal-form.
- We use this representation in Matrix 10.1, where each cell includes expected payoffs for each player.

Example 10.2. Applying BNE to sequential-move games of incomplete information

		Player 2	
		<i>A</i>	<i>R</i>
Player 1	$O^B O^F$	$\underline{4 - p}, \underline{4p}$	$-3 + p, -3 + 5p$
	$O^B N^F$	$3p, \underline{4p}$	$-2p, 2p$
	$N^B O^F$	$4(1 - p), \underline{0}$	$-3(1 - p), -3(1 - p)$
	$N^B N^F$	$0, \underline{0}$	$\underline{0}, \underline{0}$

Matrix 10.1. Bayesian normal-form representation of the game tree in Figure 10.1

- For instance, at $(O^B O^F, A)$:
 - Player 1 makes an offer to player 2 regardless of whether the investment is beneficial or not, and
 - Player 2 accepts the offer.




Figure 10.4a. Pooling strategy profile (O^B, O^F) - Responses.

Example 10.2. Applying BNE to sequential-move games of incomplete information

		Player 2	
		A	R
Player 1	$O^B O^F$	<u>$4 - p$</u> , <u>$4p$</u>	$-3 + p$, $-3 + 5p$
	$O^B N^F$	$3p$, <u>$4p$</u>	$-2p$, $2p$
	$N^B O^F$	$4(1 - p)$, <u>0</u>	$-3(1 - p)$, $-3(1 - p)$
	$N^B N^F$	0 , <u>0</u>	<u>0</u> , <u>0</u>

Matrix 10.1. Bayesian normal-form representation of the game tree in Figure 10.1

- Then, $(O^B O^F, A)$ yields expected payoffs of:
 - $EU_1 = p3 + (1 - p)4 = 4 - p$ for player 1 and
 - $EU_2 = p4 + (1 - p)0 = 4p$ for player 2.
- Underlining expected payoffs, we find two BNEs in this game:
 - $(O^B O^F, A)$ and $(N^B N^F, R)$.




Figure 10.4a. Pooling strategy profile (O^B, O^F) - Responses.

Example 10.2. Applying BNE to sequential-move games of incomplete information

		Player 2	
		A	R
Player 1	$O^B O^F$	$\underline{4 - p}, \underline{4p}$	$-3 + p, -3 + 5p$
	$O^B N^F$	$3p, \underline{4p}$	$-2p, 2p$
	$N^B O^F$	$4(1 - p), \underline{0}$	$-3(1 - p), -3(1 - p)$
	$N^B N^F$	$0, \underline{0}$	$\underline{0}, \underline{0}$

Matrix 10.1. Bayesian normal-form representation of the game tree in Figure 10.1

- Underlining expected payoffs, we find two BNEs in this game:
 - $(O^B O^F, A)$ and $(N^B N^F, R)$.

Example 10.2. Applying BNE to sequential-move games of incomplete information

- **First BNE:** Intuitively, this BNE says that:
 - player 1 makes an offer regardless of whether the investment is beneficial for player 2 or not,
 - and player responds 2 accepting it.
- This occurs because, in expectation, accepting the offer yields a higher payoff than rejecting it,

$$4p > -3 + 5p,$$

which simplifies to $3 > p$; a condition that holds for all probabilities $p \in [0,1]$.

Example 10.2. Applying BNE to sequential-move games of incomplete information

- **Second BNE:** This BNE, however, predicts that:
 - player 1 never makes an offer (regardless of her type),
 - yet, player 2 would respond rejecting an offer if she ever receives one.
- This behavior is sequentially irrational:
 - if player 2 receives an offer,
 - at the only information set where she is called to move
 - (see right hand side of Figure 10.1),
 - she earns a higher expected payoff accepting than rejecting it,
 - This holds *regardless* of the belief, μ , that she sustains about whether the investment is beneficial or not.

Example 10.2. Applying BNE to sequential-move games of incomplete information

- Indeed, player 2's expected payoffs from accepting and rejecting are

$$EU_A = 4\mu + 0(1 - \mu), \text{ and} \\ EU_R = 2\mu + (-3)(1 - \mu) = -3 + 5\mu$$

which satisfy $EU_A > EU_R$ since $4\mu > -3 + 5\mu$ simplifies to $3 > \mu$, which holds for all $\mu \in [0,1]$.

- In other words, player 2 should respond accepting the offer...
 - regardless of her belief of the investment being beneficial, μ ,
 - which contradicts the second BNE, $(N^B N^F, R)$, where player 2 rejects the offer with certainty.
- As a consequence, we claim that the second BNE is *sequentially irrational*:
 - Once an offer is received, it is sequentially rational to accept it.

Perfect Bayesian Equilibrium - Definition

- Then, we need a new solution concept to solve sequential-move games of incomplete information.
- To guarantee players choose sequentially rational strategies in this context.
- Just as we deployed a new solution concept, SPE, in sequential-move games of complete information.
- Before we introduce our new solution concept (Perfect Bayesian Equilibrium, PBE), let's clarify some notation.

Perfect Bayesian Equilibrium - Definition

Simplifying notation, Strategies.

- As in our discussion about equilibrium strategies in BNE:
 - Let $s_i^*(\theta_i)$ denote player i 's equilibrium strategy when her type is θ_i (which could be privately or publicly observed).

- Similarly, let

$$s_{-i}^*(\theta_{-i}) \equiv (s_1^*(\theta_1), \dots, s_{i-1}^*(\theta_{i-1}), s_{i+1}^*(\theta_{i+1}), \dots, s_N^*(\theta_N))$$

represent the strategy profile of player i 's rivals, given their types,

- Then, an equilibrium strategy profile can be compactly expressed as

$$s^* \equiv (s_i^*(\theta_i), s_{-i}^*(\theta_{-i})).$$

- In most applications, where we only consider two players (the first and second mover), this strategy profile simplifies to

$$s^* = (s_i^*(\theta_i), s_j^*(\theta_j)) \text{ where } j \neq i.$$

Perfect Bayesian Equilibrium - Definition

Simplifying notation, Beliefs.

- We can also provide a compact representation of players' beliefs.
- Let μ_j denote the list of player j 's updated beliefs $\mu_j(\theta_i | s_i)$:
 - for every type $\theta_i \in \Theta_i$ that player i may have,
 - and for every action $s_i \in S_i$ that she may take.
- *Example:* in the simplest scenario where player i 's types and actions are binary, $\Theta_i = \{H, L\}$ and $S_i = \{s_i^A, s_i^B\}$, player j 's "system of beliefs" has four components, as follows,

$$\mu_j = \left(\mu_j(L | s_i^A), \mu_j(H | s_i^A), \mu_j(L | s_i^B), \mu_j(H | s_i^B) \right)$$

which can be actually represented by the first and third components alone given that the second term can be expressed as a function of the first term, $\mu_j(H | s_i^A) = 1 - \mu_j(L | s_i^A)$, and, similarly, the fourth term is function of the third, $\mu_j(H | s_i^B) = 1 - \mu_j(L | s_i^B)$.

Perfect Bayesian Equilibrium - Definition

- **Simplifying notation, Beliefs (cont'd).**
- More generally, player j 's system of beliefs has
$$(card(\Theta_i) - 1) \times card(S_i)$$
- components, where “card” denotes the cardinality of the strategy set (i.e., the number of different strategies for that player).
- For example, if player i has three types, $\Theta_i = \{L, M, H\}$, and three possible actions, $S_i = \{s_i^A, s_i^B, s_i^C\}$, player j 's system of beliefs has $2 \times 3 = 6$ components.
- We can apply a similar definition to the system of beliefs by each of j 's rivals, that is,

$$\mu_{-j} \equiv (\mu_1, \dots, \mu_{j-1}, \mu_{j+1}, \dots, \mu_N)$$

so a system of beliefs can be compactly expressed as

$$\mu \equiv (\mu_j, \mu_{-j}) \text{ for all players.}$$

Perfect Bayesian Equilibrium (PBE)

- **Definition.** A strategy profile s^* and a system of beliefs μ over all information sets is a Perfect Bayesian Equilibrium (PBE) if:
 - a. Every player i 's strategies specify optimal actions at each information set where she is called to move,
 - given the strategies of the other players,
 - and given player i 's system of beliefs μ_i ; and
 - b. Beliefs μ are consistent with Bayes' rule whenever possible.

Perfect Bayesian Equilibrium (PBE)

- Intuitively, condition (a) means that:
 - every player i chooses best responses to her rivals' strategies in an incomplete information environment,
 - that is, given her beliefs about her rivals' types at that point of the game tree.
- Condition (a) can, then, be interpreted as
 - the application of best responses under incomplete information that we already considered in BNEs.
- Condition (b), however, states that every player's beliefs must be "consistent with Bayes' rule whenever possible."

Perfect Bayesian Equilibrium (PBE)

- To understand this requirement:
 - First, note that applying Bayes' rule is only possible along the equilibrium path, that is, at information sets that are reached in equilibrium.
- In contrast, when a player is called to move at an information set that is not reached in equilibrium,
 - she cannot use Bayes' rule to update her beliefs of her rivals' types.
- If she did, she would obtain an indeterminate result,

$$\mu_j(\theta_i|s_i) = \frac{0}{0}.$$

- (More about this in Example 10.3.)
- We refer to these beliefs as “off-the-equilibrium” beliefs.
 - Because we cannot use Bayes' rule to update them, we cannot specify any arbitrary value to them, leaving them unrestricted as follows $\mu_j(\theta_i|s_i) \in [0,1]$.

Example 10.3. Applying Bayes' rule

- For simplicity, we consider the labor-market signaling game of Figure 10.2, as that it allows for more concrete results.
- Recall that, in that context, the sender (job applicant) has:
 - one of two types describing her productivity level, $\Theta = \{L, H\}$, and
 - chooses whether to acquire education, E , or not, NE .
- For compactness, we use:
 - $E^H(NE^H)$ to represent that the high-productivity worker acquires (does not acquire) education;
 - a similar notation applies for the low-productivity worker, where $E^L(NE^L)$, respectively).

Example 10.3. Applying Bayes' rule

1. Separating strategy, (E^H, NE^L) .

- Consider that the job applicant requires education only when her productivity is high.




Figure 10.5. Separating strategy profile (E^H, NE^L) .

Example 10.3. Applying Bayes' rule




Figure 10.5. Separating strategy profile (E^H, NE^L) .

1. Separating strategy, (E^H, NE^L) .

- Applying Bayes' rule, we have that, upon observing that the job applicant acquires education, E , the employer's updated belief about the worker's productivity being H is

$$\begin{aligned}\mu(H|E) &= \frac{\mu(H)Pr(E|H)}{Pr(E)} = \frac{\mu(H)Pr(E|H)}{\mu(H)Pr(E|H) + \mu(L)Pr(E|L)} = \frac{q \times 1}{(q \times 1) + ((1 - q) \times 0)} = \frac{q}{q} \\ &= 1\end{aligned}$$

Since:

- prior probabilities are $\mu(H) = q$ and $\mu(L) = 1 - q$ by assumption, and
- in this separating strategy profile, $Pr(E|H) = 1$ and $Pr(E|L) = 0$ because the H -type chooses E with certainty while the L -type never does.

Example 10.3. Applying Bayes' rule

1. *Separating strategy, (E^H, NE^L) .*

- Intuitively, $\mu(H|E) = 1$ says that:
 - the employer, upon observing that the job applicant acquired education,
 - becomes convinced of facing an H -type since only H -types choose E .
- In other words, E can only originate from the H -type.
 - It is the only player type who chooses education.




Figure 10.5. Separating strategy profile (E^H, NE^L) .

Example 10.3. Applying Bayes' rule




Figure 10.5. Separating strategy profile (E^H, NE^L) .

1. Separating strategy, (E^H, NE^L) .


- We can also evaluate the employer's posterior beliefs upon observing NE , that is,

$$\mu(H|NE) = \frac{\mu(H)Pr(NE|H)}{Pr(NE)} = \frac{\mu(H)Pr(NE|H)}{\mu(H)Pr(NE|H) + \mu(L)Pr(NE|L)} = \frac{q \times 0}{(q \times 0) + ((1 - q) \times 1)} = \frac{0}{1 - q} = 0$$

which means that NE can only originate from the L -type of sender.

- Observing NE helps the employer infer that she must face an L -type worker.
- In short, $\mu(H|NE) = 0$ because $\mu(H|E) = 1$.

Example 10.3. Applying Bayes' rule



2. *Separating Strategy, (NE^H, E^L)*


- The opposite separating strategy profile, where only the low-productivity worker acquires education, yields:
 - The opposite results, of course!
- Upon observing that the job applicant chooses NE , the employer's updated beliefs are

$$\mu(H|E) = \frac{\mu(H)Pr(NE|H)}{Pr(NE)} = \frac{\mu(H)Pr(NE|H)}{\mu(H)Pr(NE|H) + \mu(L)Pr(NE|L)} = \frac{q \times 1}{(q \times 1) + ((1 - q) \times 0)} = \frac{q}{q} = 1$$

meaning that:

- NE cannot originate from the L -type of worker,
- stemming instead from the H -type worker,
- that is, $\mu(H|NE) = 1$ and $\mu(H|E) = 0$.

Example 10.3. Applying Bayes' rule



3. Pooling strategy, (E^H, E^L) .

- In a pooling strategy profile in which the job applicant chooses E regardless of her type, the receiver's beliefs are unaffected after applying Bayes' rule because


$$\mu(H|E) = \frac{\mu(H)Pr(E|H)}{Pr(E)} = \frac{\mu(H)Pr(E|H)}{\mu(H)Pr(E|H) + \mu(L)Pr(E|L)} = \frac{q \times 1}{(q \times 1) + ((1 - q) \times 1)} = \frac{q}{q + (1 - q)} = q$$

and since $\mu(H) = q$, we obtain that the employer's posterior and prior beliefs coincide, that is,

$$\mu(H|E) = \mu(H) = q \quad \text{and} \quad \mu(L|E) = \mu(L) = 1 - q$$

- In other words, because all types of sender (job applicant in this example) choose the same message (E in this strategy profile):
 - the receiver (employer) cannot infer or extract any additional information she did not originally had.
- This is a useful property to remember when checking whether a pooling strategy profile can be sustained as a PBE.
 - No need to compute Bayes' rule.

Example 10.3. Applying Bayes' rule



3. Pooling strategy, (E^H, E^L) .

- What if, instead, the employer observes the job applicant choosing NE .
- In this strategy profile, this should not happen in equilibrium.
 - We refer to this type of message as “off the equilibrium”
 - to emphasize it should not occur in the strategy profile we analyze.
- If in this setting, the receiver (employer) tries to use Bayes' rule to update her beliefs, she finds an indetermination because

$$\mu(H|NE) = \frac{\mu(H)Pr(NE|H)}{Pr(NE)} = \frac{\mu(H)Pr(NE|H)}{\mu(H)Pr(NE|H) + \mu(L)Pr(NE|L)} = \frac{q \times 0}{(q \times 0) + ((1 - q) \times 0)} = \frac{0}{0}$$

- In this case, the off-the-equilibrium beliefs (since NE is an off-the-equilibrium strategy) can take any arbitrary number.
 - That is, $\mu(H|NE) \in [0,1]$.
 - A similar argument applies to $\mu(L|NE)$ since, again, NE should have not been observed in this strategy profile.

Example 10.3. Applying Bayes' rule




Figure 10.6. Pooling strategy profile (NE^H, NE^L) .

4. Pooling strategy, (NE^H, NE^L) .

- A similar reasoning applies to the pooling strategy profile where the job applicant does not acquire education, NE , regardless of her type, that is,


$$\mu(H|NE) = \frac{\mu(H)Pr(NE|H)}{Pr(NE)} = \frac{\mu(H)Pr(NE|H)}{\mu(H)Pr(NE|H) + \mu(L)Pr(NE|L)} = \frac{q \times 1}{(q \times 1) + ((1 - q) \times 1)} = \frac{q}{q + (1 - q)} = q$$

meaning that posteriors and priors coincide:

- $\mu(H|NE) = \mu(H) = q$,
- which implies that $\mu(L|NE) = \mu(L) = 1 - q$.
- Following the same argument as in the pooling strategy profile (E^H, E^L) , we can confirm that Bayes' rule does not help us pin down off-the-equilibrium beliefs in this case either since

$$\mu(H|E) = \frac{\mu(H)Pr(E|H)}{Pr(E)} = \frac{\mu(H)Pr(E|H)}{\mu(H)Pr(E|H) + \mu(L)Pr(E|L)} = \frac{q \times 0}{(q \times 0) + ((1 - q) \times 0)} = \frac{0}{0}$$
- We say that off-the-equilibrium beliefs can take any arbitrary number, $\mu(H|E) \in [0, 1]$.

Example 10.3. Applying Bayes' rule




5. *Semi-separating strategy.*

- If the sender (job applicant) uses a semi-separating strategy such as:
 - E^H with probability 1 and
 - E^L with probability $\sigma \in [0,1]$.
- Then player j 's updated beliefs are

$$\begin{aligned}\mu(H|E) &= \frac{\mu(H)Pr(E|H)}{Pr(E)} = \frac{\mu(H)Pr(E|H)}{\mu(H)Pr(E|H) + \mu(L)Pr(E|L)} = \frac{q \times 1}{(q \times 1) + ((1 - q) \times \sigma)} \\ &= \frac{q}{q + (1 - q)\sigma}\end{aligned}$$

and $\mu(H|E) = 1 - \mu(L|E)$.

Example 10.3. Applying Bayes' rule



5. *Semi-separating strategy.*

- When the L -type never plays E , $\sigma = 0$, the above expression simplifies to:
 - $\mu(H|E) = 1$, as in the separating strategy profile (E^H, NE^L) .
- In contrast, when she plays E with certainty, $\sigma = 1$, the employer's posterior belief becomes:
 - $\mu(H|E) = q$, as in the pooling strategy profile (E^H, E^L) .
- More generally:
 - $\mu(H|E) = \frac{q}{q+(1-q)\sigma}$ decreases in the probability with which the L -type plays E , as captured by σ .
 - Starting at $\mu(H|E) = 1$ when $\sigma = 0$ and reaching $\mu(H|E) = q$ when $\sigma = 1$.
 - Figure.
- Informally, the employer is more convinced of facing an H -type worker when the L -type rarely acquires education than when she often does.

Tool 10.1 Finding PBEs in signaling games

1. Specify a strategy profile for the sender.
 - Such as (E^H, NE^L) in the labor-market signaling game.
 - This is just our “candidate” for a PBE.
2. Update the beliefs of the receiver using Bayes’ rule, whenever possible.
 - This means that you can update beliefs in equilibrium.
 - Not off-the-equilibrium, where beliefs are left unrestricted.
3. Given the receiver’s updated beliefs in Step 2, find her optimal response.
4. Given the receiver’s response in Step 3, find the sender’s optimal message.
5. If the sender’s optimal message found in Step 4:
 - a. Coincide with the message initially postulated in Step 1, we just confirmed that the strategy profile listed in Step 1, along with the receiver’s updated beliefs in Step 2 and response in Step 3, can be sustained as a PBE.
 - b. Does not coincide with that initially postulated in Step 1, we can claim that the strategy profile in Step 1 cannot be supported as a PBE.

Finding PBEs in games with one information set

Separating strategy profile (O^B, N^F)

1. Specifying a strategy profile.

- Step 1 is given by specifying a “candidate” of strategy profile for the sender that we seek to test as a PBE, (O^B, N^F) .
- Figure 10.3 highlights the tree branch corresponding to:
 - O^B , at the top right, and
 - N^F , at the bottom left.




Figure 10.3. Separating strategy profile (O^B, N^F) .

Finding PBEs in games with one information set

Separating strategy profile (O^B, N^F)

2. Bayes' rule.

- Step 2 is straightforward in this case.
 - As described in Example 10.3, $\mu = 1$,
 - If the receiver receives an offer in this separating strategy profile, he infers that the investment must be beneficial.
- Graphically, if he is called to move at the only information set of the game tree...
 - he puts full probability weight on being at the top node of this information set.
- Informally, the receiver focuses his attention on the top right corner of the tree.

Finding PBEs in games with one information set

Separating strategy profile (O^B, N^F)

3. Optimal Response

- Given our result from Step 2, $\mu = 1$, the receiver responds accepting the offer, A , since $4 > 2$.
 - (Note: we only compare his payoffs at the top right corner of the tree, as he puts full probability weight on the top node.)
- To keep track of our results, we shade the branch corresponding to A .
- Note that A is shaded both at the top and bottom nodes:
 - since the receiver cannot condition his response to the sender's true type, which he does not observe.




Figure 10.3a. Separating strategy profile (O^B, N^F) - Responses.

Finding PBEs in games with one information set

Separating strategy profile (O^B, N^F)

4. Optimal messages

- From our results in Step 3, we now need to identify the sender's optimal message, which needs to be separately done for each of the sender's types.
 - a. When the investment is beneficial (at the top of the figure), choosing O^B , as prescribed by this strategy profile, yields 3, since the sender anticipates that the offer will be accepted by the receiver, as found in step 3.
 - b. Graphically, we only need to follow the shaded branches.
 - c. If, instead, the sender deviates towards N^B (top left in the figure), her payoff decreases to 0.
 - d. Therefore, the sender chooses O^B when the investment is beneficial.




Figure 10.3a. Separating strategy profile (O^B, N^F) - Responses.

Finding PBEs in games with one information set

Separating strategy profile (O^B, N^F)

4. Optimal messages

- From our results in Step 3, we now need to identify the sender's optimal message, which needs to be separately done for each of the sender's types.
- b. When the investment is futile (at the bottom of the figure), choosing N^F , as prescribed by this strategy profile, yields 0. Deviating towards O^F (bottom right hand of the figure) increases her payoff to 4, since she anticipates the offer is accepted.




Figure 10.3a. Separating strategy profile (O^B, N^F) - Responses.

Finding PBEs in games with one information set

Separating strategy profile (O^B, N^F)

5. Summary

- From step 4b, we found that at least one of the sender types has incentives to deviate from (O^B, N^F) ,
 - namely, when the business opportunity is futile, N^F cannot be supported as optimal for the sender.
- In summary, the separating strategy profile (O^B, N^F) cannot be sustained as PBE.
- Had we found that all sender types had incentives to behave as prescribed in Step 1
 - making an offer when the test is beneficial, O^B , but not making it otherwise, N^F
 - we would be able to claim that (O^B, N^F) can be sustained as PBE.




Figure 10.3a. Separating strategy profile (O^B, N^F) - Responses.

Finding PBEs in games with one information set

Pooling strategy profile (O^B, O^F)

1. Specifying a strategy profile.

- We start specifying the strategy profile “candidate” that we test as a PBE, (O^B, O^F) .
- Figure 10.4 shades the branches corresponding to:
 - O^B , at the top right, and
 - O^F at the bottom right.




Figure 10.4. Pooling strategy profile (O^B, O^F) .

Finding PBEs in games with one information set

Pooling strategy profile (O^B, O^F)

2. Bayes' rule.

- In Step 2, we update the receiver's beliefs.
- As described in section 10.4, posterior and prior beliefs coincide in this strategy profile,
 - entailing that $\mu = p$.
- Intuitively, upon receiving an offer,
 - the receiver cannot infer any information from this offer,
 - as all sender types make offers,
 - being left as uninformed as he was at the beginning of the game.




Figure 10.4. Pooling strategy profile (O^B, O^F) .

Finding PBEs in games with one information set

Pooling strategy profile (O^B, O^F)

3. Optimal Response

- Given our result from Step 2, $\mu = p$, the receiver responds accepting the offer, A , since his expected payoff satisfies

$EU_A = 4p + 0(1 - p) > 2p + (-3)(1 - p) = EU_R$,
which simplifies to

$$EU_A = 4p > -3 + 5p = EU_R$$

which ultimately reduces to $3 > p$, which holds for all $p \in [0, 1]$.

- In other words, player 2 responds accepting the offer regardless of the probability that the investment is beneficial, p .
- To keep track of our results, we shade the branch corresponding to A in Figure 10.4a.




Figure 10.4a. Pooling strategy profile (O^B, O^F) - Responses.

Finding PBEs in games with one information set

Pooling strategy profile (O^B, O^F)

4. Optimal messages

- From our results in Step 3, we now need to identify the sender's optimal message, separately analyzing each sender's types.
 - a. When the investment is beneficial (at the top of figure 10.4a), choosing O^B , as prescribed by this strategy profile, yields 3, whereas deviating to N^B (top left in the figure), decreases her payoff to 0. Therefore, the sender chooses O^B .
 - b. When the investment is not beneficial (at the bottom of the figure), choosing O^F , as prescribed by this strategy profile, yields 4. Deviating towards N^F would decrease her payoff to 0. Then, sender chooses O^F .




Figure 10.4a. Pooling strategy profile (O^B, O^F) - Responses.

Finding PBEs in games with one information set

Pooling strategy profile (O^B, O^F)


5. Summary

- From Step 4b, we found that all sender types prefer to behave as prescribed in Step 1, (O^B, O^F) , implying that:
- this pooling strategy profile can be supported as a PBE,
 - with the receiver holding beliefs $\mu = p$ and
 - responding accepting the offer.

Finding PBEs in games with two information sets

Separating strategy profile (E^H, NE^L)

1. *Specifying a strategy profile.*
 - We first specify the separating strategy profile that we seek to test as a PBE, (E^H, NE^L) .
 - Figure 10.5 shades the branches corresponding to E^H , at the top right, and to NE^L at the bottom left.

Figure 10.5. Separating strategy profile (E^H, NE^L) .

Finding PBEs in games with two information sets

Separating strategy profile (E^H, NE^L)

2. *Bayes' rule.*

- We can now update the firm's beliefs. Beliefs in this strategy profile satisfy

$$\mu = \frac{\frac{1}{3} \alpha^H}{\frac{1}{3} \alpha^H + \frac{2}{3} \alpha^L} = \frac{\frac{1}{3} 1}{\frac{1}{3} 1 + \frac{2}{3} 0} = 1$$

where α^H (α^L) denotes the probability that a high-productivity (low-productivity) worker acquires education.

- Intuitively, upon observing education, the firm believes it must only originate from a high-productivity worker, which corresponds to the top right node of the game tree.




Figure 10.5. Separating strategy profile (E^H, NE^L) .

Finding PBEs in games with two information sets

Separating strategy profile (E^H, NE^L)

2. Bayes' rule.

- Upon observing no education, however, the firm's beliefs are

$$\gamma = \frac{\frac{1}{3}(1 - \alpha^H)}{\frac{1}{3}(1 - \alpha^H) + \frac{2}{3}(1 - \alpha^L)} = \frac{\frac{1}{3}0}{\frac{1}{3}0 + \frac{2}{3}1} = 0$$

or, alternatively, the probability of being at the bottom left node, $1 - \gamma$, is 100 percent.

- Hence, if the firm observes no education, it assigns full probability of facing a low-productivity worker.




Figure 10.5. Separating strategy profile (E^H, NE^L) .

Finding PBEs in games with two information sets

Separating strategy profile (E^H, NE^L)

3. Optimal Response

- Given our results from Step 2, we now analyze the firm's responses upon observing each of the two possible messages (education or no education).
 - Upon observing education, the firm responds hiring the worker as a manager, M , since $10 > 4$ at the top right side of the game tree. (Recall that the firm believes after observing education, that it deals with a high-productivity worker.)




Figure 10.5a. Separating strategy profile E^H, NE^L - Responses.

Finding PBEs in games with two information sets

Separating strategy profile (E^H, NE^L)

3. Optimal Response

- b. If, instead, the firm observes no education, it responds hiring the worker as cashier, C' , because $4 > 0$ at the bottom left corner of the tree.

To keep track of our results, Figure 10.5a shades the branches corresponding to M in the right side of the tree, and those corresponding to C' in the left side.




Figure 10.5a. Separating strategy profile E^H, NE^L - Responses.

Finding PBEs in games with two information sets

Separating strategy profile (E^H, NE^L)

4. *Optimal Messages.* From our results in Step 3, we now identify the worker's optimal message, separately analyzing each type.

a. *High productivity.*

- At the top of the game tree, the high-productivity worker chooses E^H , moving rightward, instead of deviating to NE^H on the left side of the tree, since $6 > 4$.
- Intuitively, this worker type anticipates that education will be recognized as a signal of her high productivity, inducing the firm to respond hiring her as a manager, as indicated by the shaded branches originating from the central node at the top of the figure and moving rightward.




Figure 10.5a. Separating strategy profile E^H, NE^L - Responses.

Finding PBEs in games with two information set

Separating strategy profile (E^H, NE^L)

4. Optimal Messages.

a. High productivity.

- If this worker deviates to NE^H , she would save the education costs, but is identified as a low-productivity type, hired as a cashier, and earning only 4.
- As a consequence, the cost of acquiring education (relatively low for this worker type) is offset by the wage gain that she experiences when she is hired as a manager.




Figure 10.5a. Separating strategy profile E^H, NE^L - Responses.

Finding PBEs in games with two information set

Separating strategy profile (E^H, NE^L)

4. Optimal Messages.

b. Low productivity.

- At the bottom of the game tree, the low-productivity worker chooses NE^L , moving leftward, which yields 4, instead of deviating to acquire education, on the right side of the tree, as that would only yield 3.
- Intuitively, acquiring education helps her “fool” the firm into believing that she is a high-productivity worker and hiring her as a manager.




Figure 10.5a. Separating strategy profile E^H, NE^L - Responses.

Finding PBEs in games with two information set

Separating strategy profile (E^H, NE^L)

4. Optimal Messages.

b. Low productivity.

- Her wage gain, however, does not offset her cost of acquiring education, which is larger than that of the high-type worker.
- As a result, the low productivity worker does not have incentives to mimic the high-productivity type acquiring education.




Figure 10.5a. Separating strategy profile E^H, NE^L - Responses.

Finding PBEs in games with two information set

Separating strategy profile (E^H, NE^L)

5. Summary.

- From Steps 4a-4b, we found that all sender types prefer to behave as prescribed in Step 1, (E^H, NE^L) ,
- implying that this separating strategy profile can be supported as a PBE:
- with the firm holding beliefs $\mu = 1$ and $\gamma = 0$, and
- responding with (M, C') i.e., hiring the worker as a manager upon observing education but as a cashier otherwise.




Figure 10.5a. Separating strategy profile E^H, NE^L - Responses.

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

1. Specifying a strategy profile.

- We start by specifying the pooling strategy profile, (NE^H, NE^L) in which no worker type acquires education




Figure 10.6. Pooling strategy profile (NE^H, NE^L) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

2. *Bayes' rule.*

- Upon observing no education, the firm's beliefs are

$$\gamma = \frac{\frac{1}{3}(1 - \alpha^H)}{\frac{1}{3}(1 - \alpha^H) + \frac{2}{3}(1 - \alpha^L)} = \frac{\frac{1}{3} \cdot 1}{\frac{1}{3} \cdot 1 + \frac{2}{3} \cdot 1} = \frac{1}{3}$$

Implying that posterior beliefs (γ) coincide with prior beliefs ($1/3$).

- In other words, observing that the worker did not acquire education,
- provides no information about her type to the firm,
- since in this strategy profile, all worker types do not acquire education,
- i.e., $1 - \alpha^H = 1 - \alpha^L = 1$.




Figure 10.6. Pooling strategy profile (NE^H, NE^L) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

2. *Bayes' rule.*

- If, instead, the firm observes that the worker has education, which should not occur in this strategy profile, its updated beliefs are

$$\mu = \frac{\frac{1}{3}\alpha^H}{\frac{1}{3}\alpha^H + \frac{2}{3}\alpha^L} = \frac{\frac{1}{3}0}{\frac{1}{3}0 + \frac{2}{3}0} = \frac{0}{0}$$

thus being undefined.

- Graphically, the information set on the right side of the game tree (when the firm observes education) happens off-the-equilibrium path and the firm's beliefs in this information set is known as "off-the-equilibrium beliefs."
- These beliefs cannot be pinned down using Bayes' rule, as we only found an indetermination.
- For generality, these beliefs are left unrestricted, so that they can take any admissible value, i.e., $\mu \in [0,1]$.




Figure 10.6. Pooling strategy profile (NE^H, NE^L) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

3. Optimal Responses.

- Given our results from Step 2, we now analyze the firm's responses upon observing each of the two possible messages (education or no education).
 - Upon observing an educated applicant, the firm responds hiring the worker as a manager, M' , or as a cashier, C' , based on its expected profit on the left side of the tree, as follows

$$E\pi_{Firm}(M') = 10\frac{1}{3} + 0\frac{2}{3} = \frac{10}{3} \simeq 3.33 \quad \text{and}$$

$$E\pi_{Firm}(C') = 4\frac{1}{3} + 4\frac{2}{3} = 4.$$

Therefore, the firm responds hiring the worker as a cashier, C' , when she did not acquire education, since $4 > 3.33$.




Figure 10.6. Pooling strategy profile (NE^H, NE^L) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

3. *Optimal Responses.*

b. If, instead, the firm observes education, it responds hiring the worker as a manager, M , if its expected profit at the right side of the game tree satisfies

$$E\pi_{Firm}(M) = 10\mu + 0(1 - \mu) = 10\mu \quad \text{and}$$

$$E\pi_{Firm}(C) = 4\mu + 4(1 - \mu) = 4$$

and, comparing these expected profits, we obtain that $10\mu > 4$ holds if $\mu > \frac{2}{5}$.

- Then, the firm responds to an educated applicant hiring her as a manager, M , when its off-the-equilibrium beliefs satisfy $\mu > \frac{2}{5}$ (when it assigns a sufficiently high probability weight on facing a high-productivity worker),
- but responds hiring the applicant as a cashier, C , otherwise.




Figure 10.6. Pooling strategy profile (NE^H, NE^L) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

3. *Optimal Responses.*

- we then need to divide our following step, where we examine the worker's decisions, into two cases:
 1. $\mu > \frac{2}{5}$, where the firm responds hiring the worker as a manager, M , after observing education;
 2. $\mu \leq \frac{2}{5}$, where the firm responds hiring her as a cashier, C , after observing education.




Figure 10.6. Pooling strategy profile (NE^H, NE^L) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

4. *Optimal Messages.* From our results in Step 3, we now identify the worker's optimal message, separately analyzing cases (1)-(2)

Case 1: $\mu > \frac{2}{5}$, illustrated in Figure 10.6a, where the firm responds with C' after observing no education (left side of the tree) and with M , after observing education (right side), given that $\mu > \frac{2}{5}$)




Figure 10.6a. Pooling strategy profile (NE^H, NE^L) - Responses (C', M) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

4. Optimal Messages. Case 1: $\mu > \frac{2}{5}$

a. High Productivity.

- At the top of the game tree, if the high-productivity worker chooses NE^H , moving leftward, she earns 4, as she anticipates that the firm responds with C' .
- If, instead, the worker deviates to E^H , she earns 6, implying that she does not have incentives to behave as prescribed by this strategy profile.
- At this point, we do not need to check if the low-productivity worker has incentives to choose NE^L , since we already found that a worker's type (high productivity) would deviate.




Figure 10.6a. Pooling strategy profile (NE^H, NE^L) - Responses (C', M) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

4. *Optimal Messages.* Case 1: $\mu > \frac{2}{5}$

b. *Low Productivity.*

- At the bottom of the game tree, if the low-productivity worker chooses NE^L , moving leftward, she earns 4, which exceeds her payoff from deviating to E^L , on the right side, 3.
- Intuitively, the low-productivity worker anticipates that she will be recognized as such when she does not acquire education, and hired as a cashier.
- Even if she could fool the firm into believing that it deals with a high-productivity worker when she acquires education, and be hired as a manager...
 - The cost of investing in education is too high for this type of worker to undergo such deviation.




Figure 10.6a. Pooling strategy profile (NE^H, NE^L) - Responses (C', M) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

4. *Optimal Messages.* Case 2: $\mu \leq \frac{2}{5}$, illustrated in Figure 10.6b, where the firm responds with C' after observing no education (left side of the tree)
5. and with C , after observing education (right side, given that $\mu \leq \frac{2}{5}$)




Figure 10.6b. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

4. Optimal Messages. Case 2: $\mu \leq \frac{2}{5}$

a. High Productivity.

- At the top of the game tree, if the high-productivity worker chooses NE^H , moving leftward, she earns 4, as she anticipates that the firm responds with C' .
- If, instead, the worker deviates to E^H , she only earns 0, implying that she does not have incentives to deviate.
- Intuitively, the high productivity worker is hired as a cashier regardless of whether she acquires education,
 - entailing that it is optimal for her to not undergo the investment in education as it does not affect the firm's response.




Figure 10.6b. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

4. *Optimal Messages.* Case 2: $\mu \leq \frac{2}{5}$

b. *Low Productivity.*

- At the bottom of the game tree, if the low-productivity worker chooses NE^L , moving leftward, she earns 4, which exceeds her payoff from deviating to E^L , on the right side, -3.
- In this case, this type of worker faces similar incentives as the high-type above, as her educational level does not affect the firm's response,
 - leading her to not invest in education.




Figure 10.6b. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Finding PBEs in games with two information sets

Pooling strategy profile (NE^H, NE^L)

5. Summary.

- When $\mu > \frac{2}{5}$ holds (Case 1), we found that one sender type (the high productivity worker) prefers to deviate from the pooling strategy profile, implying that (NE^H, NE^L) cannot be supported as PBE (see step 4a).
- In contrast, when $\mu \leq \frac{2}{5}$ holds, both worker types have incentives to behave as prescribed by pooling strategy profile (NE^H, NE^L) , implying that it is sustained as a PBE with the firm responding with (C, C') and with equilibrium beliefs $\gamma = 1/3$ and off-the-equilibrium beliefs $\mu \leq \frac{2}{5}$.

Insensible off-the-equilibrium beliefs

- The pooling strategy profile (NE^H, NE^L) , where no worker type acquires education, requires that $\mu \leq \frac{2}{5}$
- Informally, this is like saying that in a world where no worker type goes to college, if a firm observes a worker with a college degree (surprise!), it is likely to face a low-productivity worker.
- You may feel that this off-the-equilibrium belief is a bit insensible
- Recall that, from our above discussion in Case 1 (Step 4b):
 - the low-productivity worker would not like to deviate, from NE^L to E^L , even if it could fool the firm into believing that such a message stems from the high-productivity worker, and respond hiring her as a manager.
- In other words, the cost of acquiring education for the low-productivity worker is so high that it prevents her from deviating from NE^L .

Evaluating PBE as a solution concept

1. Existence? Yes.

- When we apply PBE to any game, we find that at least one equilibrium exists.
- Intuitively, this result is equivalent to the existence of SPE in sequential-move games of complete information,
 - but extended to an incomplete information setting.
- For this finding to hold, however, we may need to allow for mixed strategies; a
 - lthough all applications in this chapter produce a pure-strategy PBE.

Evaluating PBE as a solution concept

2. Uniqueness? No.

- This point is illustrated by the labor market signaling game in section 10.7 where we found two PBEs:
 1. The separating strategy profile where only the high-type worker acquired education (E^H, NE^L) , and
 2. The pooling strategy profile where no types of worker does, (NE^H, NE^L)
- Other games, however, such as that in section 10.6, have a unique PBE:
 - (O^B, N^F) , where player 1 only makes an investment offer, when such investment is beneficial for player 2.
- Therefore, we cannot guarantee that the PBE solution concept provides a unique equilibrium prediction in *all* games,
 - entailing that uniqueness does not hold for PBE.

Evaluating PBE as a solution concept

3. Robust to small payoff perturbations? Yes.

- PBE yields the same equilibrium predictions if we were to change the payoff of one of the players by a small amount (e.g., 0.001 or, generally, any ε that approaches zero).
- This occurs because, if a strategy s_i is sequentially optimal for player i in our original game, it must still be optimal after we apply a small payoff perturbation.

Evaluating PBE as a solution concept

4. Socially optimal? No.

- As described in the labor market signaling game (section 10.7), the presence of incomplete information gives rise to inefficiencies.
- In particular, in the separating PBE (E^H, NE^L) , the high-type worker invests in a costly education (which does not improve her productivity) just to signal her type to the firm, and thus, be hired as a manager.
- In a complete information setting, instead, the firm would observe the worker's types, hiring the high (low) type as a manager (cashier), leaving the worker with no incentives to acquire education to convey their types to the firm.
- In that setting, their payoffs would be $(10,10)$ when the worker's productivity is high, and $(4,4)$ when it is low.

Evaluating PBE as a solution concept

4. Socially optimal? No.

- In contrast, in the separating PBE, equilibrium payoffs are $(6,10)$ and $(4,4)$. As a consequence, if players behaved as under complete information, the high-type worker would improve her payoff from 6 to 10 (savings in education acquisition) while the payoffs of the low-type worker and the firm would remain unaffected.
- A similar argument applies to the pooling PBE (NE^H, NE^L) , where equilibrium payoffs are $(4,4)$ regardless of the worker's type.
- In this context, the high-type worker and firm would improve their payoffs if they could behave as under complete information (increasing from 4 to 10 for both of them); while the low-type worker and firm would see their payoffs unchanged.

Semi-separating PBE

- What if, after checking for all the separating and pooling strategy profiles, as candidates for PBEs (our “usual suspects”)...
 - we find that none of them can be supported as PBE?
- Does it mean that the signaling game has no PBE?
 - No, it just means that we need to allow at least one sender type to randomize her messages.
- This is analogous as our inability to find psNEs in a simultaneous-move games of complete information,
 - where we could identify msNE where at least one player randomizes.

Semi-separating PBE

- Figure 10.7 depicts a simple poker game where, first, player 1 (sender) privately observes whether she has a high or low hand; then player 1 chooses to bet (B) or resign (R).
- If she resigns, the game is over and player 2 earns the initial pot.
- However, if she bets, player 2 (receiver) must respond calling or folding without observing player 1's hand but knowing that high and low hands are equally likely.




Figure 10.7. Poker game.

Semi-separating PBE

- In this setting, none of the separating strategy profiles, (B^H, R^L) and (R^H, B^L) , or the pooling strategy profiles (B^H, B^L) or (R^H, R^L) can be sustained as a PBE.
- Unlike in previous signaling games, this occurs because player 1, when holding a high hand, would like to induce player 2 to respond calling it; but when holding a low hand, she would prefer player 2 to respond folding (informally, to not call her bluff).




Figure 10.7. Poker game.

Semi-separating PBE

- In other words, each sender type would like the receiver (player 2) to *incorrectly infer* her type:
 - i. thinking that player 1's hand is low, calling it, when it is actually high (at the top right side of the tree);
 - ii. Thinking that her hand is high, folding, when it is in fact low (bottom right side).
- To prevent her type from being recognized, player 1 can create some “noise” in the signal, by betting with a positive probability.
- This randomization reduces player 2's ability to infer her type:
 - betting can now originate from the high- or low-hand players, with potentially different probabilities,
 - having a similar effect as in msNEs, where players mix to keep their rivals guessing about what their next moves are.




Figure 10.7. Poker game.

Semi-separating PBE

- Figure 10.7a reproduces 10.7, but highlighting the branch corresponding to betting for the high-hand player 1, B^H , who finds bet to be a strictly dominant strategy, i.e., B^H yields a payoff of 2, if player 2 responds calling, or 1, if she responds folding, both of them exceeding her payoff from resigning, R^H , where she earns 0 with certainty.
- Therefore, the high-hand player 1 play B^H in pure strategies.
- In contrast, the low-hand player 1 assigns probability p on B^L and $1 - p$ on R^L , where $p \in (0,1)$.




Figure 10.7a. Poker game - Semiseparating strategy profile.

Semi-separating PBE

- Finally, we labelled player 2's probability of calling as q , where $q \in (0,1)$, and that we seek to identify.
- Note that if, instead, we allowed for $q = 0$ ($q = 1$), and player 2 folded (called) in pure strategies, the player 1 would like to bet (resign) with certainty when her hand is low, implying that $p = 1$ ($p = 0$), implying that a pooling (separating) could be supported as a PBE, which we know cannot hold.
- Therefore, player 2 must be randomizing with (non-degenerate) strategies, such that $q \in (0,1)$.




Figure 10.7a. Poker game - Semiseparating strategy profile.

Semi-separating PBE: Identifying PBEs




Figure 10.7a. Poker game - Semiseparating strategy profile.

1. Specifying a strategy profile.

- We first specify the semi-separating strategy profile that we seek to test as a PBE:
 - player 1 chooses B^H with certainty, B^L with probability p , and
 - player 2 responds calling with probability $q \in (0,1)$.

2. Bayes' Rule. We can now update player 2's beliefs. Upon observing player 1 betting, beliefs satisfy

$$\mu = \frac{\frac{1}{2} \cdot 1}{\frac{1}{2} \cdot 1 + \frac{1}{2} p} = \frac{\frac{1}{2}}{\frac{1}{2} + \frac{1}{2} p} = \frac{1}{1 + p}$$

since player 1 bets with certainty when her hand is high, but bets with probability p otherwise.

Semi-separating PBE: Identifying PBEs




Figure 10.7a. Poker game - Semiseparating strategy profile.

3. *Optimal responses.* Given our results from Step 2, we now analyze player 2's response.

- Recall that player 2 must be mixing, $q \in (0,1)$. Otherwise, as discussed above, the low-hand player 1 would have incentives to use pure strategies, which are not PBEs of this game.
- In addition, if player 2 is mixing, she must be indifferent between C and F , entailing that

$$\mu(-1) + (1 - \mu)2 = \mu 0 + (1 - \mu)0$$

Where:

- the left side represents player 2's expected utility of calling (which yields a payoff of -1 when player 1's hand is high, but 2 otherwise);
- while the right side denotes her expected utility from folding (which is 0 regardless of player 1's hand. Solving for μ , yields $\mu = \frac{2}{3}$.

Semi-separating PBE: Identifying PBEs




Figure 10.7a. Poker game - Semiseparating strategy profile.

3. *Optimal responses.*

- But, with which probability q does player 2 call?
 - We know that, as in msNEs, player 2 must be mixing with probability q that makes the low-hand player 1 indifferent between betting and resigning.
- This implies that $EU_1(B^L) = EU_1(R^L)$ or

$$q(-1) + (1 - q)1 = 0$$
- since the low-hand player earns -1 when player 2 responds calling but 1 when player 2 folds.
- Solving for q , we find that $q = \frac{1}{2}$, implying that player 2 calls with 50 percent probability.

Semi-separating PBE: Identifying




Figure 10.7a. Poker game - Semiseparating strategy profile.

4. Optimal messages.

- From our results in Step 3, we now identify player 1's optimal message.
- The high-hand player 1 finds betting, B^H , to be strictly dominant,
 - so we focus on the low-hand player 1, who calls with probability p .
- From Bayes' rule, we know that $\frac{2}{3} = \frac{1}{1+p}$, where $\mu = \frac{2}{3}$ from Step 3.
 - Solving for p , we obtain that $p = \frac{1}{2}$, so the low-hand player 1 bets with 50 percent probability.

Semi-separating PBE: Identifying PBEs




Figure 10.7a. Poker game - Semiseparating strategy profile.

5. Summary.

- From the above steps, we found a PBE where:
 - the high-hand player 1 bets, B^H , with certainty;
 - the low-hand player 1 bets, B^L , with probability $p = \frac{1}{2}$; and
 - player 2 responds calling with probability $q = \frac{1}{2}$, sustaining beliefs $\mu = \frac{2}{3}$.

Extensions: What if...

- For presentation purposes, the labor market signaling game only had:
 - two types, two messages, and two responses.
- But, what if...
 - The sender had three or more possible types.
 - The sender had three or more possible messages.
 - The receiver had three or more possible responses.
- We next analyze each of these extensions at a time.
 - First, finding how to depict them in the game tree.
 - And, then, how it affects our application of Tool 10.1 to find PBEs.

Extensions: What if the receiver has more than two available responses?

- We allow for the firm to have 3 available responses:
 - hiring the worker as the company director, manager, or cashier as depicted in Figure 10.8.
- Relative to figure 10.2, only the number of branches stemming from nodes connected by each information sets increases,
 - from two to three branches originating from each node (director, manager, or cashier hire).




Figure 10.8. Labor-market signaling game with three responses.

How is Tool 10.1 affected by more responses?

- Our analysis in section 10.7 would be mostly unaffected.
- In particular, the list of strategy profiles to test in Step 1 remains unchanged,
 $(E^H, NE^L), (NE^H, E^L), (E^H, E^L)$ and (NE^H, NE^L) .
- The firm's updated beliefs in Step 2 are also unaffected,
 - because the number of sender types and the number of available messages are unchanged,
 - therefore, not modifying the firm's information sets.
- However, the firm's optimal response (Step 3) is affected,
 - as the firm would choose the response that yields the highest payoff, by comparing its three possible responses.
 - Exercise 10.9 asks you to test which strategy profiles can be sustained as PBEs.
- This argument extends to other signaling games where we allow for $k \geq 2$ available responses, in which the firm would choose the one that yields the highest payoff.

What if the sender has more than two available messages?

- Alternatively, the labor market signaling game could be extended by allowing the worker to choose between more than two available messages, such as:
 - acquiring an advanced graduate degree (A),
 - an undergraduate degree (E), or
 - no college education (NE).
- Figure 10.9 illustrates this setting
- Relative to Figure 10.2, each worker type (either high or low productivity) chooses now among three, instead of two, possible branches (A , E , or NE).




Figure 10.9. Labor-market signaling game with three messages

What if the sender has more than two available messages?

- Each branch, in turn, gives rise to a different information set:
 - one after the firm observes A ,
 - another after observing E , and
 - another after observing NE .
- In each information set, however, there are still two nodes:
 - upon observing a given message, the firm faces the same type of uncertainty as in Figure 10.5,
 - namely, not knowing whether the message originates from the high- or low-type.




Figure 10.9. Labor-market signaling game with three messages

What if the sender has more than two available messages?

- While Figure 10.9 is relatively similar to Figure 10.2 (except for the branches where the worker chooses A), its tree representation may be difficult to interpret.
- For clarity, games with more than two messages are often graphically represented with a game tree like that in Figure 10.9a.
- Figure 10.9a can be understood as the “vertical” representation of the game, as nature chooses first, followed by each worker type, and the firm after observing each possible message.




Figure 10.9a. Labor-market signaling game with three messages - Alternative representation

What if the sender has more than two available messages?

- Regarding payoffs, note that the high-productivity worker incurs an additional cost of 1 when choosing A^H , relative to when he chooses E^H ;
 - whereas the low-productivity worker's additional cost is 3.
- The firm's payoffs are, of course, unaffected since education (either E or A) is not productivity-enhancing.




Figure 10.9a. Labor-market signaling game with three messages - Alternative representation

How is Tool 10.1 affected by more available messages?

- We now examine how the five-step tool to identify PBE is affected when the sender has more available messages.
- We start at Step 1, where we consider a specific strategy profile as a candidate for PBE.
- In this setting, we have the following strategy profiles to check as PBEs:
 - six separating,
 $(A^H, E^L), (A^H, N^L), (E^H, N^L), (E^H, A^L), (N^H, A^L), (N^H, E^L),$
 - three pooling,
 $(A^H, A^L), (E^H, E^L),$ and $(N^H, N^L).$
- More generally:
 - If there are x sender types and each sender has y available messages, then there are a total of y^x different strategy profiles.
 - In our example, $x = 2$ types and $y = 3$ messages, entailing $3^2 = 9$ different profiles (6 separating and 3 pooling).

How is Tool 10.1 affected by more available messages?

1. Specifying a strategy profile.

- Consider the strategy profile (A^H, E^L) , depicted in Figure 10.9b.
- Intuitively, both worker types acquire education,
 - but the high type invests in an advanced graduate degree while the low type completes an undergraduate degree.




Figure 10.9b. Labor-market signaling game with three messages - Strategy profile (A^H, E^L) .

How is Tool 10.1 affected by more available messages?

2. *Bayes' rule.* In this context, the firm's updated beliefs are

$$\mu(H|j) = \frac{\frac{1}{3}\alpha_j^H}{\frac{1}{3}\alpha_j^H + \frac{2}{3}\alpha_j^L}$$

where $j = \{A, E, NE\}$ denotes the message that the firm observes, and $\alpha_j^H (\alpha_j^L)$ represents the probability that the high (low) type sends message j .

- In the strategy profile (A^H, E^L) , this belief becomes

- $\mu(H|A) = \frac{\frac{1}{3}}{\frac{1}{3} + \frac{2}{3} \cdot 0} = 1$ after observing A ,
- $\mu(H|E) = \frac{\frac{1}{3} \cdot 0}{\frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1} = 0$ after observing E , and
- $\mu(H|NE) = \frac{\frac{1}{3} \cdot 0}{\frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 0} = \frac{0}{0}$ after observing NE (off-the-equilibrium).

How is Tool 10.1 affected by more available messages?

2. *Bayes' rule.*

- In summary, upon observing $A(E)$, the firm is convinced of facing a high (low) productivity worker, as only this worker type uses this message;
 - but upon observing NE , which is off-the-equilibrium path, the firm cannot update its beliefs according to Bayes' rule.
 - We must leave them unrestricted as $\mu(H|NE) \in [0,1]$.
- In most applications, however, researchers assume that,
 - if firm's beliefs are $\mu(H|E) = 0$ after E ,
 - they must be $\mu(H|NE) = 0$ after NE .
 - Intuitively, if an undergraduate degree signals that the worker's productivity is low, observing that the worker did not even complete college (a lower signal) must provide a similar information about the worker's type.
- From a practical approach, this assumption simplifies our analysis in the subsequent steps of Tool 10.1, and we consider it here.

How is Tool 10.1 affected by more available messages?

3. *Optimal Responses.*

- Upon observing, A , the firm responds hiring the worker as a manager, M'' , since $10 > 4$ at the left of the game tree, $\beta = 1$.
- Upon observing, E , the firm responds hiring the worker as a cashier, C , since $4 > 0$ at the right side of the game tree, $\mu = 0$.
- Upon observing, E , the firm responds hiring the worker as a cashier, C' , since $4 > 0$ at the right side of the game tree, $\gamma = 0$.




Figure 10.9c. Labor-market signaling game with three messages – Optimal messages

How is Tool 10.1 affected by more available messages?

3. *Optimal Messages.* From our results in Step 3, we now identify the worker's optimal message, separately analyzing each type.

a. *High productivity.*

- At the right side of the game tree, if the high-productivity worker chooses A^H , moving upward, she earns 5, as she anticipates the firm responding with M'' .
- If instead, she deviates to E^H on the left side of the tree, she would earn 0 (as she is hired as a cashier, C).
- A similar argument applies, if she does not acquire education, NE^H , at the left side of the tree, where she is also hired as a cashier, C' , and her payoff is 4.
- Therefore, she does not have incentives to deviate from A^H .




Figure 10.9c. Labor-market signaling game with three messages – Optimal messages

How is Tool 10.1 affected by more available messages?

3. Optimal Messages.

b. Low productivity.

- At the left side of the game tree:
 - if the low-productivity worker chooses E^L , on the center, she earns -3.
 - Then, she has incentives to deviate to NE^L , on the right side of the tree, which yields 4.
 - (This is her best deviation, as deviating to A^L only entails a payoff of 0.)




Figure 10.9c. Labor-market signaling game with three messages – Optimal messages

How is Tool 10.1 affected by more available messages?


5. *Summary.*

5. From Steps 4a-4b, we found that:

- one sender type (the low-productivity worker)
- deviates from the strategy profile prescribed in Step 1, (A^H, E^L) ,
- implying that this profile cannot be supported as a PBE.

What if the sender has more than two types?

- Finally, Figure 10.10 extends the signaling game of Figure 10.2 to allow for the worker to have more than two types, such as high, medium, or low productivity.
- Relative to the game tree in Figure 10.2:
 - a new branch stems from the initial node where nature draws the worker's type (at the top of the figure), corresponding to medium productivity.



What if the sender has more than two types?

- For clarity, Figure 10.10a depicts the “vertical” version of the game tree in Figure 10.10.
- This representation helps illustrate that:
 - after the initial move of nature,
 - every type of worker privately observes her type, and
 - chooses whether to acquire education, E^k , or not, NE^k , where $k = \{H, M, L\}$.




Figure 10.10a. Labor-market signaling game with three worker types – vertical version

What if the sender has more than two types?

- Two information nodes still.
- But each information set connects *three* nodes.
 - One for each sender type.
- This means that, upon observing a worker with education:
 - the firm does not know whether her productivity is high, medium, or low;
 - as depicted by the dashed line with the label “Firm after E”.
- A similar argument applies when the firm observes a worker with no education,
 - in the information set with the label “Firm after NE”
 - at the bottom of Figure 10.10a.




Figure 10.10a. Labor-market signaling game with three worker types – vertical version

What if the sender has more than two types?

- Regarding players' payoffs:
 - First, note that the medium-productivity worker incurs a cost of 6 when acquiring education,
 - which is in-between that of the low-productivity worker (7) and that of the high-productivity worker (4).
- The firm's profits are always 4 when hiring the applicant as a cashier,
 - regardless of her type (high, medium, or low) and regardless of her education.
- When hiring her as a manager, however, the firm's profit is:
 - 10 when the applicant's productivity is high,
 - 5 when it is medium, and
 - 0 otherwise.




Figure 10.10a. Labor-market signaling game with three worker types – vertical version

How is Tool 10.1 affected by more than two types?

- In this setting, we have different profiles to test:

$$(E^H, E^M, NE^L), (E^H, NE^M, E^L), (E^H, NE^M, NE^L), (NE^H, E^M, E^L), \\ (NE^H, E^M, NE^L), (NE^H, NE^M, E^L), \\ (E^H, E^M, E^L) \text{ and } (NE^H, NE^M, NE^L)$$

- The last two strategy profiles are clearly pooling,
 - because all worker types choose the same message (all acquire education or none do),
- The first six strategy profiles are neither fully separating:
 - as two worker types pool into sending the same message,
- Not fully pooling either:
 - as one worker type chooses a different message than the other two types.
- We may refer to them as “partially separating” profiles.


How is Tool 10.1 affected by more than two types?

- Following our discussion in section 10.10.1:
 - if there are $x = 3$ sender types and each sender has $y = 3$ available messages,
 - then there are a total of $2^3 = 8$ different strategy profiles, as identified in our list above.
- For illustration purposes, we focus on (E^H, E^M, NE^L) .
 - Let's test if it can be sustained as a PBE.

How is Tool 10.1 affected by more than two types?

1. Specifying a strategy profile.

- Consider strategy profile (E^H, E^M, NE^L) , as in Figure 10.10b, where:
 - both the high- and medium-productivity worker acquire education
 - while the low-productivity worker does not.



How is Tool 10.1 affected by more than two types?

2. *Bayes' Rule.* In this setting, the firm's updated beliefs are:

$$\mu(H|E) = \frac{\frac{1}{3}\alpha^H}{\frac{1}{3}\alpha^H + \frac{1}{3}\alpha^M + \frac{1}{3}\alpha^L}$$

where Figure 10.10b assumes, for simplicity, that all types are equally likely.

- In strategy profile (E^H, E^M, NE^L) , for instance, $\alpha^H = \alpha^M = 1$ but $\alpha^L = 0$, yielding an updated belief

$$\mu(H|E) = \frac{\frac{1}{3}}{\frac{1}{3} + \frac{1}{3} + 0} = \frac{1}{2}.$$

- Intuitively, the firm knows that education must originate from either the high or medium types.
 - Since both are equally likely, the probability of facing a high-productivity worker is 50 percent.

How is Tool 10.1 affected by more than two types?

2. *Bayes' Rule.* In this setting, the firm's updated beliefs are:

$$\mu(H|E) = \frac{\frac{1}{3}\alpha^H}{\frac{1}{3}\alpha^H + \frac{1}{3}\alpha^M + \frac{1}{3}\alpha^L}$$

where Figure 10.10b assumes, for simplicity, that all types are equally likely.

- Similarly,

$$\mu(M|E) = \frac{\frac{1}{3}}{\frac{1}{3} + \frac{1}{3} + 0} = \frac{1}{2}, \quad \text{but} \quad \mu(L|E) = \frac{0}{\frac{1}{3} + \frac{1}{3} + 0} = 0$$

since the low-productivity worker does not acquire education in this strategy profile.

- Upon not observing education, however, we obtain more concentrated beliefs, that is, $\mu(H|NE) = \mu(M|NE) = 0$ whereas $\mu(L|NE) = 1$, as the firm is convinced of facing a low-productivity worker.

How is Tool 10.1 affected by more than two types?


3. *Optimal Responses.* Given our results from Step 2, we now analyze the firm's responses upon observing each of the three possible messages.
- a. Upon observing E , the firm responds hiring the worker as a manager, M , since

$$\underbrace{\frac{1}{2}10 + \frac{1}{2}5}_{=7.5} > 4$$

at the right side of the game tree.

The firm now does not have concentrated beliefs upon observing E , as it can originate from the high or medium type.

Therefore, it must compute its expected profit from each of its decisions.




How is Tool 10.1 affected by more than two types?

3. *Optimal Responses.*

- b. Finally, upon observing NE , the firm responds hiring the worker as a cashier too, C' , since we imposed that off-the-equilibrium beliefs satisfy $\mu(H|NE) = \mu(M|NE) = 0$.

Figure 10.10c shades the branches corresponding to M and C' .




How is Tool 10.1 affected by more than two types?

4. Optimal Messages.

a. *High productivity.* At the left of the game tree, the high-productivity worker chooses education:

- Her payoff from doing so, 6, exceeds
- Her payoff from not acquiring education, 4.




How is Tool 10.1 affected by more than two types?

4. Optimal Messages.

a. *Medium productivity.* At the center of the tree, the medium-productivity worker is indifferent between:

- choosing education, earning 4, or
- not acquiring education, also earning 4.




How is Tool 10.1 affected by more than two types?

4. Optimal Messages.

c. *Low productivity.* At the right side of the tree, the low-productivity worker does not acquire education:

- Her payoff from No education, 4, exceeds
- Her payoff from Education, 3.



How is Tool 10.1 affected by more than two types?

5. *Summary.*

- From Steps 4a-4b, we found that all sender types prefer to behave as prescribed in Step 1, (E^H, E^M, NE^L) , implying that:
 - this separating strategy profile can be supported as a PBE,
 - with the firm holding beliefs $\mu(H|E) = \mu(M|E) = \frac{1}{2}$ and $\mu(L|E) = 0$ upon observing education,
 - and $\mu(H|NE) = \mu(M|NE) = 0$, $\mu(L|NE) = 1$ after no education,
 - responding with (M, C') , i.e., hiring the worker as a manager upon observing education but as a cashier otherwise.

Other Extensions

- **Productivity-enhancing education.**

- Previous sections assumed, for simplicity, that education does not affect a worker's job productivity.
- We said productivity wasn't productivity enhancing, thus only serving as a signal of the worker's type.
- In many real-life settings, however, education makes the worker more productive.
- We explore this possibility in Exercises 10.2 and 10.3.

Other Extensions

- **More general cost differentials across types.**
 - Our above model assumed that the cost of acquiring education for the high (low) productivity worker was $c_H = 4$ ($c_L = 7$, respectively).
 - Exercise 10.5 allows, for generality, that the low-productivity worker's cost is $c_L > 4$, so the difference $c = c_L - c_H > 0$ represents the cost differential that the low-productivity worker suffers relative to the high productivity worker when acquiring education.
 - That is, a higher cost differential may deter the low-productivity worker from mimicking the high-type's decision to acquire education, ultimately facilitating the emergence of separating PBEs.

Other Extensions

- **More general profit differentials across types.**

- The standard labor-market signaling game considered that hiring a high (low) productivity worker as a manager provides the firm with a profit of $d_H = 10$ ($d_L = 0$, respectively);
 - hiring her as a cashier produces a profit of 4 regardless of the the worker's type.
- In Exercise 10.6, we allow for a more general profit differential between worker types when hired as a manager, where we still assume that $d_H = 10$ but $d_L = 10 - d$, where $d = d_H - d_L$ and $0 \leq d \leq 10$.
- Intuitively, as d increases, the firm hiring the low-productivity worker as a manager entails a larger loss, which the firm seeks to avoid;
 - while when $d = 0$, the firm earns the same profit hiring both worker types as a manager.

Chapter 11: Equilibrium Refinements

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- Chapter 10 highlights that signaling games and the PBE solution concept are an excellent tool to explain a wide array of economic situations
- However, we found that this class of games may yield:
 - A large number of PBEs
 - Some PBEs based on insensible off-the-equilibrium beliefs.
- In this chapter, we present two commonly used tools to “refine” set of PBEs
 - satisfying different consistency requirements in their off-the-equilibrium beliefs

Introduction

- Cho and Kreps' Intuitive Criterion
- D1 Criterion
- Sequential equilibrium

Intuitive Criterion

- The Cho and Kreps' (1982) “Intuitive Criterion”:
 - eliminates all PBEs that are sustained on insensible off-the-equilibrium beliefs,
 - such as the pooling PBE (NE^H, NE^L) in Figure 11.1.
- Recall that this PBE required an off-the-equilibrium belief $\mu \leq \frac{2}{5}$.
 - Informally, if the firm observes, surprisingly, a worker acquiring education, it believes the worker must likely be low productivity.
 - Crazy, right?




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Intuitive Criterion

- Essentially, the Intuitive Criterion seeks to answer a relatively basic question:
 - “If the receiver observed an off-the-equilibrium message, such as education on the right side of Figure 11.1...
 - which sender types could benefit from sending such an off-the-equilibrium message?”




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Intuitive Criterion

- If no sender types can benefit, then the initial PBE we considered *survives* the Intuitive Criterion.
- If some sender types can benefit, the receiver would then update her off-the-equilibrium beliefs and her response to this message.
- In turn, this updated response could induce some sender types to deviate from their equilibrium messages, implying that the initial PBE we considered *violates* the Intuitive Criterion.
- More details in Tool 11.1




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Tool 11.1 Applying the Cho and Kreps' (1982) Intuitive Criterion

Step 1. Consider a specific PBE

- For instance, the pooling PBE (NE^H, NE^L) .
- Recall that (NE^H, NE^L) could be supported as a PBE if the firm responds with (C, C') and its beliefs are:
 - $\gamma = \frac{1}{3}$ upon observing no education (in the equilibrium path), and
 - $\mu \leq \frac{2}{5}$ upon observing education (off-the-equilibrium path)




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Tool 11.1 Applying the Cho and Kreps' (1982) Intuitive Criterion

Step 2. Identify an off-the-equilibrium message for the sender

- That is, find an action that no type of sender chooses in the PBE we are analyzing,
 - e.g., education in the pooling PBE (NE^H, NE^L)
-
- Think about games with two types but k messages.
 - There would be $k-1$ off-the-equilibrium messages.




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Tool 11.1 Applying the Cho and Kreps' (1982) Intuitive Criterion

Step 3. Find which types of senders can obtain a higher utility level by deviating (i.e., when they send off-the-equilibrium messages) than by choosing their equilibrium message.

- The high-productivity worker can benefit when she deviates from NE^H , where she earns a payoff of 4, to E^H , where she can earn a payoff of 6 if the firm responds hiring her as a manager
- The low-productivity worker, however, cannot benefit when she deviates from NE^L , where she earns a payoff of 4, to E^L , where her highest payoff is 3 when the firm hires her as a manager.




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Tool 11.1 Applying the Cho and Kreps' (1982) Intuitive Criterion

Step 4. Restrict the off-the-equilibrium beliefs of the responder using the results of Step 3.

- Intuitively, if education is observed,
 - the firm believes it can only originate from the high-productivity worker
 - since the low-productivity worker cannot benefit from acquiring education.
- Therefore, the firm's beliefs upon observing education are restricted from $\mu \leq \frac{2}{5}$ to $\mu = 1$.




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Tool 11.1 Applying the Cho and Kreps' (1982) Intuitive Criterion

Step 5. Find optimal response given the restricted beliefs found in Step 4.

- In our ongoing example, upon observing education, the firm is convinced of dealing with a high-productivity worker, $\mu = 1$, and
- optimally responds hiring her as a manager, M (earning a profit of 10), which exceeds the profit of hiring the worker as a cashier (4).




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Tool 11.1 Applying the Cho and Kreps' (1982) Intuitive Criterion

Step 6. Given the optimal response found in Step 5, check if there is one of more sender types who can profitably deviate from her equilibrium message.

- If there is, we say that the PBE we considered *violates* the Intuitive Criterion.
- Otherwise, the PBE *survives* the Intuitive Criterion




Figure 11.1. Pooling strategy profile (NE^H, NE^L) - Responses (C', C) .

Pooling Equilibria May Violate the Intuitive Criterion

- In the pooling PBE (NE^H, NE^L) , education is an off-the-equilibrium message.
- Moving on to Step 3, we note that:
 - only the high-productivity worker can profitably deviate towards education, as she could increase her payoff if the firm responds hiring her as a manager.
 - Formally, the highest payoff she earns from deviating to education is 6, which exceeds her equilibrium payoff of 4.
 - The low-productivity worker cannot profitably deviate, as the highest payoff she can earn from deviating is 3, which lies below her equilibrium payoff of 4.

Pooling Equilibria May Violate the Intuitive Criterion

- In other words, even if she fools the firm into believing that her productivity is high, and the firm hires her as a manager, her cost of acquiring education offsets her wage increase.
- In summary, we say that the low-productivity worker does not find education to be *equilibrium dominated*, since

$$u_L(NE^L, C') = 4 > 3 = \max_{a \geq 0} u_L(E^L, a)$$

where $a \in \{M, C\}$ denotes the firm's response.

- In contrast, the high-productivity worker does not find education to be equilibrium dominated.

Pooling Equilibria May Violate the Intuitive Criterion

- From Step 3, the firm can, in Step 4, restrict its off-the-equilibrium belief to $\mu = 1$
 - (if it observed education, it must stem from the high-productivity worker).
- As a consequence, in Step 5, we find that the firm's optimal response to education, given $\mu = 1$, is:
 - to hire the worker as a manager, M , as its profit, 10 exceeds that of hiring her as a cashier, 4.
- Finally, in Step 6, we obtain that:
 - one sender type (the high-productivity worker) has incentives to deviate from NE^H in the pooling PBE towards E^H ,
 - as she anticipates that the firm will respond hiring her as a manager (Step 5).
- In conclusion, we can claim that:
 - this pooling PBE *violates* the Intuitive Criterion,
 - implying that we can eliminate the pooling PBE (NE^H, NE^L) as a solution to this game.

Separating Equilibria Survive the Intuitive Criterion

- In the separating PBE (E^H, NE^L) of Figure 10.5:
 - there is no off-the-equilibrium message in Step 2 of the Intuitive Criterion.
- In other words,
 - all possible messages (education and no education) are being used by at least one worker type.
 - Alternatively, no message is unused.




Figure 10.5. Separating strategy profile (E^H, NE^L) .

Separating Equilibria Survive the Intuitive Criterion

- In this setting, the separating PBE survives the Intuitive Criterion.
- For a PBE to violate the Intuitive Criterion, we need to go through steps 1 to 6a,
 - and in this case Step 2 does not identify off-the-equilibrium messages,
 - entailing that we cannot continue with steps 3 and beyond.




Figure 10.5. Separating strategy profile (E^H, NE^L) .

Separating Equilibria Survive the Intuitive Criterion

- This argument applies to all separating PBEs:
 - With the same number of types as messages (as in the labor market signaling game).
 - With more sender types than messages.
- It doesn't necessarily apply to separating PBEs:
 - With fewer sender types than messages.
 - At least one of the messages is left unused, becoming an off-the-equilibrium message.
 - This is a “partially separating” PBE, which may survive or violate the Intuitive Criterion.

Separating Equilibria Survive the Intuitive Criterion

<i>Separating PBEs</i>	<i>Survives the Intuitive Criterion?</i>
# sender types = # messages	Yes
# sender types > # messages	Yes
# sender types < # messages	Maybe

D1 Criterion

- The D1 Criterion:
 - considers a similar approach as the Intuitive Criterion, as described in the six steps in Tool 11.1,
 - but differs in Step 3.
- In particular, for a given off-the-equilibrium message, the D1 Criterion seeks to answer the following question:

”If the receiver observes an off-the-equilibrium message,
such as education on the right side of Figure 11.1,
for which sender type are *most of the receiver’s actions beneficial?*”

D1 Criterion

- After that, we restrict the receiver's beliefs as in Step 4,
 - and proceed through Steps 5 and 6 as in Tool 11.1.
- That is, the Intuitive Criterion seeks to identify which sender types *could benefit* from deviating towards the off-the-equilibrium message,
 - which in some games could lead to a large number of sender types
- This approach doesn't restrict the receiver's beliefs in Step 4 enough to rule out PBEs sustained on insensible off-the-equilibrium beliefs.

D1 Criterion

- The D1 Criterion poses a similar question,
 - But, in a sense, it “goes deeper”
 - by finding not only which sender types have incentives to deviate...
 - but which sender type is “the most likely to deviate” as we describe next.

D1 Criterion

1. First, we measure the number of the receiver's responses (e.g., salary offers from the firm) that would weakly improve the sender's payoff, relative to her equilibrium payoff,
 - repeating this process for each sender type.
 2. Second, we compare which sender type has the largest set of responses that would make her better off than in equilibrium.
- The application of these two steps slightly differs depending on whether the receiver has:
 - a discrete or continuous strategy space.
 - We separately consider them next.

D1 Criterion

- **Discrete responses:**

- In a setting with discrete responses, such as when the firm hires the applicant as a CEO, manager, or cashier,

1. In the first step, we would just count!

- How many of the firm responses can improve the high-productivity worker's equilibrium payoff,
- How many responses can improve the low-productivity worker's, ...
- Similarly for each worker's type

2. In the second step:

- we would compare which worker type has more payoff-improving responses,
- identifying her as the one who is “the most likely to deviate.”

D1 Criterion

- **Continuous Responses:**

- In a context with continuous responses, such as when the firm responds with a salary $w > 0$ to the worker's observed education, these two steps are slightly different.
 1. In particular, the first step would measure the set of firm responses that improve the worker's productivity (i.e., a wage interval).
 2. The second step would compare the size of these wage intervals across worker types.
 - In this context, the worker with the longest "utility-improving wage interval" is the one who is the "most likely to deviate"

Applying the D1 Criterion – An Example

- We test whether the (NE^H, NE^L) pooling PBE survives the D1-criterion.

Step 1. Consider a specific PBE, such as the pooling PBE (NE^H, NE^L) we found in section 10.7.2.

Step 2. Identify an off-the-equilibrium message for the sender.

- In (NE^H, NE^L) , the only off-the-equilibrium message is that of Education, on the right-side information set.

Applying the D1 Criterion – An Example

Step 3. Find which sender type is more *likely to benefit* from the deviation of Step 2.

(Recall: this is the main difference relative to the Intuitive Criterion).

To identify this type of sender, we go through the following steps:

- a. When the high-productivity worker acquires education, she only has one response from the firm that can improve her equilibrium payoff, M , where her payoff increases from 4 to 6.
- b. When the low-productivity worker acquires education, however, she has no responses from the firm that could improve her equilibrium payoff of 4.
- c. Comparing the number of payoff-improving responses for the high- and low-productivity workers, we find that the former has more than the latter (1 vs. 0) implying that the high type is more likely to deviate.

Applying the D1 Criterion – An Example

Step 4. (From this point on, all steps coincide with those in the Intuitive Criterion).

- Restrict the off-the-equilibrium beliefs of the responder using the results of Step 3.
- Intuitively, if education is observed, the firm believes it must originate from the high-productivity worker, restricting its off-the-equilibrium beliefs from $\mu \leq \frac{2}{5}$ to $\mu = 1$.

Step 5. Find the optimal response given the restricted beliefs found in Step 4.

- Upon observing education, the firm is convinced of facing a high-productivity worker.
- If it responds hiring her as a manager, M , it earns a profit of 10, which exceeds the profit of hiring her as a cashier, 4.

Applying the D1 Criterion – An Example

Step 6. Given the optimal response found in Step 5, check if there is one or more sender types who can profitably deviate from her equilibrium message.

- a. If there is, we say that the PBE we considered *violates* the D1-Criterion.
- b. Otherwise, the PBE *survives* the D1-Criterion.

- In our ongoing example, after Step 5, we see that:
 - the high-productivity (low-productivity) deviates (does not deviate) as she anticipates the firm to respond to education hiring her as a manager, M .
- Therefore, we can conclude that pooling PBE (NE^H, NE^L) violates the D1-Criterion.

Comparing the Intuitive and D1 Criterion

- Therefore, we found the same refinement result as when applying the Intuitive Criterion;
 - but that is not generally the case.
- Both refinement criteria yield the same results in signaling games with only two sender types,
 - explaining why the Intuitive Criterion is used in this class of games,
 - as it is often easier to apply and provides the same refinement power.
- When the sender has three or more possible types, however,
 - the D1 Criterion gives us more restricted equilibrium predictions.
- Formally, this means that, if strategy profile $s = (s_1, s_2, \dots, s_N)$ can be sustained as a PBE,

s survives the D1 criterion $\Rightarrow$ s survives the Intuitive Criterion

$\Leftarrow$

Comparing the Intuitive and D1 Criterion

- We return to this point in our analysis of signaling games with continuous action spaces in Chapter 12 (section 12.7).
- But intuitively, their difference arises from Step 3.
 - In the Intuitive Criterion, Step 3 only helps us to eliminate those sender types for which deviations are equilibrium dominated
 - i.e., they are better off behaving as prescribed by the PBE.
 - In the D1-Criterion, in contrast, Step 3 identifies which sender type is the “most likely to benefit from the deviation,”
 - implying that we focus our attention on a specific type of sender (or types) found in Step 3 of the Intuitive Criterion.

Other Refinement Criteria

- The literature offers other refinement criteria in signaling games which you may consider
 - when analyzing the PBEs where the Intuitive or D1 Criterion have no bite,
 - or when a large number of PBEs survive these two criteria.
- Examples include the:
 - “divinity” criterion and “universal divinity” criterion, both after Banks and Sobel (1988).

Other Refinement Criteria

- **Divinity Criterion.** It is a weakening of the D1 Criterion:
 - It requires that the receiver's posterior belief after observing an off-the-equilibrium message cannot increase the likelihood ratio of the sender being type θ relative to θ' .
 - Note that the D1 Criterion sets a stronger requirement, as it puts zero probability weight on this type of sender.

Other Refinement Criteria

- **University Divinity Criterion.** It requires that a sender type θ is eliminated, making updated beliefs independent on priors.
- These criteria are, however, beyond the scope of this book, and the advanced reader can refer to Banks and Sobel (1988) for more details.
- For experiments testing if subjects behave as prescribed by the Intuitive and D1 Criteria, among other refinements, see:
 - Brandts and Holt (1992) and
 - Banks et al. (1994).

Sequential Equilibrium

- While refining PBEs, it yields similar results as the Intuitive or D1 Criteria and it is not as straightforward to apply as these two criteria, where we only need to check for sender types who could benefit from deviating to an off-the-equilibrium message.
- Consider a behavioral strategy profile $b = (b_i, b_{-i})$.
- Recall that player i 's behavioral strategy b_i prescribes which action she selects when she is called to move at information set h_i
- Allowing her to choose an action among those available at information set h_i , or instead, to randomize over two or more actions available to her at h_i .

Sequential Equilibrium

- Definition. **Sequential Equilibrium (SE).** Behavioral strategy profile $b = (b_i, b_{-i})$ and beliefs μ over all information sets are sustained as a sequential equilibrium if:
 1. *Sequential rationality.* Behavioral strategy b_i specifies optimal actions for every player i at each information set where she is called to move, given the strategies of the other players, b_{-i} , given the system of beliefs μ ; and
 2. *Belief consistency.* There is a sequence $\{(b^k, \mu^k)\}_{k=1}^{\infty}$ such that, for all k :
 - a. b^k is a totally mixed behavioral strategy, assigning a positive probability weight to all available actions at each information set h_i where player i gets to play
 - b. μ^k is consistent with Bayes' rule given b^k
 - c. Sequence $\{(b^k, \mu^k)\}_{k=1}^{\infty}$ converges to (b, μ)

Sequential Equilibrium

- While sequential rationality (point 1) is analogous to the requirement on the PBE definition, belief consistency (point 2) differs because it used totally mixed behavioral strategies.
- Graphically, point 2a implies that every information on the game tree is reached with positive probability, which did not necessarily happen in PBEs where some information sets may not be reached in equilibrium
- By reaching all information sets, Bayes' rule is no longer undefined, and helps us update beliefs given players' behavior
- Finally, point 2c states that the perturbations in the sequence become smaller as k increases, ultimately converging to (b, μ)

Sequential Equilibrium

- Intuitively, both PBE and SE require players to choose sequentially rational actions at every information set they are called to move.
- The SE solution concept, instead, required that off-the-equilibrium beliefs are consistent with Bayes' rule when players deviate from their equilibrium strategies with a positive probability
- Therefore, SE helps us identify which PBEs are sustained on consistent off-the-equilibrium beliefs.
- But exhibits no refining power on separating PBEs where all information sets are reached in equilibrium
- Every separating PBE is also a SE, but not all pooling PBEs are necessarily SEs as some may be sustained on insensible off-the-equilibrium beliefs.

$$(b, \mu) \text{ is a SE} \Rightarrow (b, \mu) \text{ is a PBE}$$
$$\nLeftarrow$$

Finding Sequential Equilibria

- Consider a variation of the labor market signaling game where the probability of the worker's productivity being high is now $\frac{2}{3}$, as depicted in Figure 11.2.
- In this setting, three strategy profiles can be sustained as PBEs




Figure 11.2. Modified labor market signaling game.

Finding Sequential Equilibria

1. The separating PBE $(E^H, NE^L; M, C')$, where only the high-productivity worker acquires education, sustained with equilibrium beliefs $\mu = 1$ ($\gamma = 0$) after observing education (no education, respectively).
2. The pooling PBE $(NE^H, NE^L; M, M')$, where no types of worker acquires education, the firm holds off-the-equilibrium beliefs $\mu > \frac{2}{5}$ and responds with M upon observing education.
3. The pooling PBE $(NE^H, NE^L; C, M')$, where no type of worker acquires education, the firm holds off-the-equilibrium beliefs $\mu \leq \frac{2}{5}$ and responds with C upon observing education.

Separating PBE that are also SEs

- We consider the separating strategy profile $(E^H, NE^L; M, C')$, as depicted in Figure 11.3, which is supported as a PBE with beliefs $\mu = 1$ and $\gamma = 0$.




Figure 11.3. Separating strategy profile $(E^H, NE^L; M, C')$.

Separating PBE that are also SEs

- First, consider the following totally mixed behavioral strategy

$$b^k = \left(\underbrace{(1 - \varepsilon^k, \varepsilon^k)}_{\text{Worker-H}}, \underbrace{(\varepsilon^k, 1 - \varepsilon^k)}_{\text{Worker-L}}, \underbrace{(1 - 2\varepsilon^k, 2\varepsilon^k)}_{\text{Firm after educ.}}, \underbrace{(2\varepsilon^k, 1 - 2\varepsilon^k)}_{\text{Firm after no educ.}} \right)$$

where $\varepsilon^k > 0$ denotes a small perturbation, the first component of the high (low) types worker's behavioral strategy denotes the probability that she chooses E^H (E^L), and the first component in the third (fourth) parenthesis represents the probability that the firm responds with M (M') upon observing education (no education, respectively).
- This totally mixed behavioral strategy converges to $(1,0; 1,0)$ when $k \rightarrow +\infty$, yielding outcome $(E^H, NE^L; M, C')$.

Separating PBE that are also SEs

- In addition, the belief system (μ^k, γ^k) is consistent with Bayes rule if

$$\mu^k = \frac{\frac{2}{3}(1-\varepsilon^k)}{\frac{2}{3}(1-\varepsilon^k) + \frac{1}{3}\varepsilon^k} = \frac{2(1-\varepsilon^k)}{2-\varepsilon^k} \text{ upon observing education, and}$$

$$\gamma^k = \frac{\frac{2}{3}\varepsilon^k}{\frac{2}{3}\varepsilon^k + \frac{1}{3}(1-\varepsilon^k)} = \frac{\frac{2}{3}\varepsilon^k}{1+\varepsilon^k} \text{ upon observing no education}$$

Separating PBE that are also SEs

- Figure 11.4 plots these beliefs as a function of k , showing that they converge to $\lim_{k \rightarrow \infty} \mu^k = 1$ and $\lim_{k \rightarrow \infty} \gamma^k = 0$.
- Therefore, updated beliefs coincide with those in the separating PBE, $\mu = 1$ and $\gamma = 0$, as required.




Figure 11.4. Updated beliefs μ^k and γ^k as a function of k .

Separating PBE that are also SEs

- Finally, given these beliefs, the firm's responses (M, C') , satisfy sequential rationality, as required by SE, since $\pi(M) = 10 > 4 = \pi(C)$ after observing education, and $\pi(C') = 4 > 0 = \pi(M')$ after no education.
- Therefore, the PBE $(E^H, NE^L; M, C')$ sustained with beliefs $\mu = 1$ and $\gamma = 0$ is also a SE of this game.

Pooling PBE that is not a SE

- We consider the pooling PBE($E^H, NE^L; M, M'$) depicted in Figure 11.5, which is based on insensible off-the-equilibrium beliefs.
- Indeed, $\mu \leq \frac{2}{5}$ means that, upon observing the off-the-equilibrium message of education, the firm assigns a relatively low weight to such as message originating from the high-productivity worker.
- It even allows for the firm to be convinced that only the low-productivity worker acquires education ($\mu = 0$)!




Figure 11.5. Pooling strategy profile ($NE^H, NE^L; M, M'$).

Pooling PBE that is not a SE

- We next demonstrate that this PBE is not a SE.
- To show this result, recall that when players choose a sequence of totally mixed behavioral strategies $\{b^k\}_{k=1}^{\infty}$, the firm's right-hand information set is reached with positive probability (i.e., E^H and E^L are played with positive probabilities).
- Because this information set is reached, we no longer face off-the-equilibrium information sets, and the firm's beliefs can be updated by Bayes' rule.

Pooling PBE that is not a SE

- In particular, if both types of workers choose education with the same small perturbation, ε^k , Bayes' rule entails that

$$\mu^k = \frac{\frac{2}{3}\varepsilon^k}{\frac{2}{3}\varepsilon^k + \frac{1}{3}\varepsilon^k} = \frac{2}{3}$$

implying that posteriors and priors coincide.

- Indeed, when perturbation ε^k is symmetric across players, the two nodes of the right-hand information set are reached with the same (small) probability, as in pooling strategy profiles.

Pooling PBE that is not a SE

- Given this updated belief, the firm responds with M after observing education since $E\pi(M) = \frac{20}{3} > 4 = E\pi(C)$.
- Therefore, responding with C cannot be part of a SE.
- We next provide a formal proof, describing the sequence of totally mixed behavioral strategies.

- Consider the following totally mixed behavioral strategy

$$b^k = \left(\underbrace{(\varepsilon^k, 1 - \varepsilon^k)}_{\text{Worker-H}}, \underbrace{(\varepsilon^k, 1 - \varepsilon^k)}_{\text{Worker-L}}, \underbrace{(2\varepsilon^k, 1 - 2\varepsilon^k)}_{\text{Firm after educ.}}, \underbrace{(1 - 2\varepsilon^k, 2\varepsilon^k)}_{\text{Firm after no educ.}} \right)$$

which converges to $(0,0; 0,1)$ when $k \rightarrow +\infty$, yielding strategy profile $(NE^H, NE^L; C, M')$.

Pooling PBE that is not a SE

- In addition, the belief system (μ^k, γ^k) is consistent with Bayes rule if

$$\mu^k = \frac{\frac{2}{3}\varepsilon^k}{\frac{2}{3}\varepsilon^k + \frac{1}{3}\varepsilon^k} = \frac{2}{3} \text{ upon observing education, and}$$
$$\gamma^k = \frac{\frac{2}{3}(1-\varepsilon^k)}{\frac{2}{3}(1-\varepsilon^k) + \frac{1}{3}(1-\varepsilon^k)} = \frac{2}{3} \text{ upon observing no education}$$

which converge to $\lim_{k \rightarrow \infty} \mu^k = \gamma^k = \frac{2}{3}$.

- Finally, given these beliefs, the firm's responses, (C, M') , *violate* sequential rationality, since $E\pi(M) = \frac{20}{3} > 4 = E\pi(C)$ after observing education and, similarly $E\pi(M') = \frac{20}{3} > 4 = E\pi(C')$ after no education.
- In conclusion, the PBE $(NE^H, NE^L; C, M')$, supported with $\mu \leq \frac{2}{5}$, is not a SE.

Pooling PBE that is also a SE

- Consider now the pooling PBE $(NE^H, NE^L; M, M')$.




Figure 11.6 Pooling strategy profile $(NE^H, NE^L; M, M')$

Pooling PBE that is also a SE

- To see this point, consider the totally mixed behavioral strategy

$$b^k = \left(\underbrace{(\varepsilon^k, 1 - \varepsilon^k)}_{\text{Worker-H}}, \underbrace{(\varepsilon^k, 1 - \varepsilon^k)}_{\text{Worker-L}}, \underbrace{(1 - 2\varepsilon^k, 2\varepsilon^k)}_{\text{Firm after educ.}}, \underbrace{(1 - 2\varepsilon^k, 2\varepsilon^k)}_{\text{Firm after no educ.}} \right)$$

which converges to $(0,0; 1,1)$ when $k \rightarrow +\infty$, yielding strategy profile $(NE^H, NE^L; M, M')$.

- In addition, the belief system (μ^k, γ^k) is consistent with Bayes rule, and produces the same results as b^k in section 11.4.1, that is, $\mu^k = \frac{2}{3}$ upon observing education, and $\gamma^k = \frac{2}{3}$ upon observing no education; converging to $\lim_{k \rightarrow \infty} \mu^k = \gamma^k = \frac{2}{3}$.

Pooling PBE that is also a SE

- Finally, given these beliefs, the firm's responses (M, M') , *satisfy* sequential rationality, since $E\pi(M) = \frac{20}{3} > 4 = E\pi(C)$ after observing education, and similarly, $E\pi(M') = \frac{20}{3} > 4 = E\pi(M')$ after no education.
- Consequently, the PBE $(NE^H, NE^L; M, M')$ supported with $\mu > \frac{2}{5}$, is also a SE.

Chapter 12: Signaling Games with Continuous Messages

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- We extend Spence's Labor Market signaling game to a setting with continuous messages and responses.
- Outline:
 - Identify equilibrium behavior under complete information
 - Find separating and pooling PBEs of the game and apply refinement criteria discussed in previous lectures.
 - Compare equilibrium results under complete and incomplete information
 - Fundamental question:
 - Are workers better off when firms observe their productivity or...
 - when they need to use education to signal their type to the employers?

Utility functions with Continuous Actions

- Consider a worker with utility function

$$u(e, w | \theta_K) = w - c(e, \theta_K)$$

where

- $w \geq 0$ represents the salary that she receives from the firm,
- $c(e, \theta_K)$ denotes her cost of acquiring e units (e.g., years) of education given that her productivity is θ_K ,
 - where $K = \{H, L\}$ is worker's type (high or low productivity).

Utility functions with Continuous Actions

- For simplicity, function $c(e, \theta_K)$ often satisfies the following four assumptions:
 1. *Zero cost of no education.* When the worker acquires no education, she incurs no cost, that is, $c(0, \theta_K) = 0$.
 2. *Cost is increasing and convex.* Cost function $c(e, \theta_K)$ is strictly increasing and convex in education that the worker acquires, i.e., $c_e > 0$ and $c_{ee} > 0$.
 - indicating that additional years of education become progressively more costly.
 3. *Cost is decreasing in productivity.* Cost function $c(e, \theta_K)$ is decreasing in the worker's productivity, θ_K , i.e., $c(e, \theta_H) < c(e, \theta_L)$.
 - implying that a given education e (e.g., a university degree) is easier to complete for the high-productivity than for the low-productivity worker.

Utility functions with Continuous Actions

4. *Marginal cost is decreasing in productivity (single-crossing property).* A similar property applies to the marginal cost of education, $c_e(e, \theta_K)$ which is also lower for the high-productivity worker than the low-productivity worker,

$$c_e(e, \theta_H) < c_e(e, \theta_L).$$

Remark:

- When messages are discrete, the fourth assumption can be expressed as “cost increments”, as follows:

$$c(e', \theta_H) - c(e, \theta_H) < c(e', \theta_L) - c(e, \theta_L)$$

Example 12.1 Cost of Acquiring Education

- Consider a cost function such as

$$c(e, \theta_K) = \frac{Ae^2}{\theta_K}$$

where $A > 0$ is a constant.

As a practice, show that it satisfies the above four assumptions.

Example 12.1 Cost of Acquiring Education

- Graphically, we can depict the worker's indifference curve in the (e, w) – quadrant by, first, solving for w , which yields

$$w = u + c(e, \theta_K) \Rightarrow w = u + \frac{Ae^2}{\theta_K},$$

$$\text{if } c(e, \theta_K) = \frac{Ae^2}{\theta_K}$$

- Since the cost of education is strictly increasing and convex in e ,
 - indifference curves are also monotonically increasing and convex in e , as depicted in Figure 12.1.




Figure 12.1 Indifference curves for a representative worker.

Example 12.1 Cost of Acquiring Education

- Intuitively, a higher education must be accompanied by an increase in her salary for the worker's utility to remain unchanged.
- Indifference curves to the northwest are associated with a higher utility:
 - for a given education e , the worker receives a higher wage (movement from point A to C in Figure 12.1), or
 - for a given wage w , the worker acquires less education (movement from B to C).




Figure 12.1 Indifference curves for a representative worker.

Example 12.1 Cost of Acquiring Education

- If $A = \theta_L = 1$ and $\theta_H = 2$, then the high- and low-type indifference curves cross at

$$u_H + \frac{e^2}{2} = u_L + e^2$$

Or

$$e = 4(u_H - u_L)^{1/2}$$




Figure 12.1 Indifference curves for a representative worker.

Complete Information

- We assume the labor market is perfectly competitive.
- Firm observes the worker's type and pays a salary equal to the worker's productivity $w = \theta_K$
 - If $w < \theta_K$, other firms could offer a slightly higher salary w' such that $w < w' \leq \theta_K$, making a weakly positive profit
 - If $w > \theta_K$, it would attract the worker but make a negative profit
- A similar argument applies when the firm does not observe the worker's type and pays a salary equal to her expected productivity, $w = E[\theta_K]$,
 - where this expectation is based on the firm's beliefs about the worker's type.

Complete Information

- As a benchmark, we consider a complete information context where the firm observes the worker's productivity, θ_K .
- In this setting, the only subgame perfect equilibrium has the worker acquiring zero education regardless of her productivity, $e_H = e_L = 0$.
- The firm responds with salary $w(\theta_H) = \theta_H$ when the worker's productivity is high, and $w(\theta_L) = \theta_L$ when her productivity is low.
- Intuitively, the worker cannot use education as a signal about her type since the firm observes her productivity.

Incomplete Information

- Under incomplete information, education can become an informative (although costly) signal.
- Assume that the worker privately observes her type θ_K before choosing her education level e .
- Firm observes education e but does not know the worker's type θ_K .
- However, it assigned a prior probability $p \in [0,1]$ to the worker's productivity being high and $1 - p$ to her productivity being low.
 - This probability distribution is common knowledge among players.

Separating PBE

- We check if a separating strategy profile where each type of worker acquires a different education level can be sustained as a PBE.
- Follow the same four steps as in the discrete version of the game.

First step - strategy profile

- High-productivity worker chooses education level e_H
- Low-productivity worker chooses e_L , where $e_H \neq e_L$.

Second step – updating beliefs

- Upon observing e_H , firm assigns full probability to facing a high-productivity worker, i.e., $\mu(\theta_H|e_H) = 1$.
- In contrast, after observing e_L , firm assigns no probability to facing a high-productivity worker, i.e., $\mu(\theta_H|e_L) = 0$.
- If the firm observes the worker acquiring an off-the-equilibrium education level, i.e., an education e different than e_H and e_L ,
 - it cannot update its beliefs using Bayes' rule, leaving them unrestricted, i.e., $\mu(\theta_H|e) \in [0,1]$ for all $e \neq e_H \neq e_L$.

Separating PBE

Third step – optimal responses

- If firm observes e_H , it pays $w(e_H) = \theta_H$, and if it observes e_L , it pays $w(e_L) = \theta_L$
- Salaries coincide with the worker's productivity, thus being the same as under complete information.
- Yet, education levels in the separating PBE do not coincide with those under complete information.
 - Upon observing $e \neq e_H \neq e_L$, a firm's beliefs are $\mu(\theta_H|e) \in [0,1]$
 - Therefore, the firm pays a salary between
 - $w(e) = \theta_L$ when $\mu(\theta_H|e) = 0$ and
 - $w(e) = \theta_H$ when $\mu(\theta_H|e) = 1$
 - That is, $w(e) \in [\theta_L, \theta_H]$

Separating PBE

- As illustrated in Figure 12.2, this wage schedule means that $w(e_L) = \theta_L$ at education level e_L , and $w(e_H) = \theta_H$ at e_H .
- However, for all $e \neq e_H \neq e_L$, the wage can lie strictly above θ_L or below θ_H (but of course bounded between θ_L and θ_H)




Figure 12.2. Separating PBE - Two wage schedules.

Separating PBE

Fourth step – optimal messages

Let's identify conditions under which worker of type θ_L has the incentives to choose e_L while θ_H type chooses e_H .

Low productivity worker: Let's find when e_L satisfies $e_L = 0$

- Figure 12.3 depicts the indifference curve of the low-productivity worker, IC_L , that passes through point $(w, e) = (\theta_L, 0)$.




Figure 12.3 Separating PBE - Low productivity worker.

Separating PBE

- At $e_L = 0$, the firm is convinced to deal with a low-productivity worker, paying a wage $w(e_L) = \theta_L$.
- Wage schedule $w(e)$ guarantees that this type of worker has no incentives to deviate from education level $e_L = 0$.
- Consider point A :
 - At e_1 , the salary the worker receives is above θ_L ,
 - but the indifference curve passing through A , IC_L^A is associated to a lower utility level than IC_L




Figure 12.3 Separating PBE - Low productivity worker.

Separating PBE

- Additional cost of education she incurs when deviating from $e_L = 0$ to $e_1 > 0$ offsets the extra salary she receives.
- A similar argument applies for any education levels $e > e_L$.
- More generally, for the low-productivity worker to stick to $e_L = 0$, we need that the firm's wage schedule to lie below IC_L




Figure 12.3 Separating PBE - Low productivity worker.

Separating PBE

High productivity worker chooses the education level prescribed in this separating strategy profile, e_H , rather than deviating to the low-type's education, $e_H = 0$ if

$$\theta_H - c(e_H, \theta_H) \geq \theta_L - c(0, \theta_H)$$

since she can anticipate that acquiring e_H identifies her as a high-productivity worker, yielding a salary $w(e_H) = \theta_H$,

while acquiring $e_L = 0$ identifies her as a low-productivity worker, receiving $w(e_L) = \theta_L$

Separating PBE

- Since cost $c(0, \theta_H) = 0$ by assumption, we can rearrange the above inequality as

$$\underbrace{\theta_H - \theta_L}_{\substack{\text{Wage} \\ \text{Differential}}} \geq \underbrace{c(e_H, \theta_H)}_{\substack{\text{Additional cost} \\ \text{of education}}}$$

- Similarly, her incentive to not deviate off the equilibrium education level $e \neq e_H \neq e_L$

$$\theta_H - c(e_H, \theta_H) \geq w(e) - c(e, \theta_H)$$

Separating PBE

- Figure 12.4 depicts a wage schedule $w(e)$ that provides the high-productivity worker with incentives to choose education level e_H , where she reaches IC_H , rather than deviating towards any other education $e \neq e_H$.
- High productivity worker does not have incentives to deviate from e_H to e_1 (lower utility level), and this applies to $e < e_H$ (including $e_L = 0$).
- No incentive either to deviate above e_H such as e_2 where her salary is represented by the height of point B .




Figure 12.4 Separating PBE - Low and high-productivity worker.

Separating PBE

- Other wage schedules also induce the high-productivity worker to choose education e_H , such as that of Figure 12.5a.
- Intuitively, this wage schedule indicates that the firm pays:
 - the lowest salary $w(e) = \theta_L$ upon observing $e < e_H$,
 - but pays the highest salary otherwise.
- The figure also depicts education levels e_1 and e_2 , confirming that the worker does not have incentives to deviate from e_H .




Figure 12.5a. Separating PBE – Case I

Separating PBE

- Different wage schedules help us support different education levels e_H for the high-productivity worker, such as that in Figure 12.5b, where the firm only pays the high salary θ_H when e_H is extremely high, that is, $e_H = e_4$.
- In this case, the high-productivity worker, despite receiving salary $w(e_4) = \theta_H$, is indifferent between acquiring education e_4 or deviating to $e_L = 0$.




Figure 12.5b. Separating PBE – Case II

Separating PBE

- Specifically, e_4 solves
$$\theta_H - \theta_L = c(e_4, \theta_H)$$
- The least-costly separating PBE, e_3 , cannot be supported with the wage schedule depicted in Figure 12.5b.
- It could be sustained with wage schedules such as those in Figures 12.4 or 12.5a.
- Overall, a wage schedule may help us sustain one equilibrium but not another one.




Figure 12.5b. Separating PBE – Case II

Example 12.2. Identifying the highest education that achieves separation

Following the setting in Example 12.1.

- Consider that $\theta_H = 2, \theta_L = 1$
- Cost of education is $c(e, \theta_K) = \frac{e^2}{\theta_K}$
- Solving for e_4 in $\theta_H - \theta_L = c(e, \theta_K)$
$$\Rightarrow \theta_H - \theta_L = \frac{e^2}{\theta_H}$$
$$\Rightarrow e_4 = \sqrt{2} \approx 1.41.$$

Example 12.2. Identifying the highest education that achieves separation

- More generally, in separating PBEs of the labor market signaling game,
 - Low-productivity worker chooses $e_L = 0$, the high-productivity worker selects an education level e_H in the range $[e_3, e_4]$
 - $e_H = e_4$ represents the “most-costly separating PBE”
 - $e_H = e_3$ represents the “least-costly separating PBE”, since the high type conveys her type to the firm acquiring the lowest education level
- The least-costly separating education level $e_H = e_3$ solves

$$\theta_H - c(e_3, \theta_L) = \theta_L - c(0, \theta_L)$$

- Intuitively, the low-productivity worker is indifferent between her equilibrium strategy $e_L = 0$, receiving a wage of θ_L , and deviating to education level e_3 which provides her with wage θ_H .

Example 12.3. Identifying the highest education that achieves separation

- Since $c(0, \theta_L) = 0$ by assumption, the above equation simplifies to

$$\theta_H - \theta_L = c(e_3, \theta_L)$$

- where e_3 solves $\theta_H - \theta_L = \frac{e^2}{\theta_L}$
- Since, $\theta_H = 2$ and $\theta_L = 1$, we obtain $e_3 = 1$.
- Therefore, $e_L = 0$ but e_H lies between $e_3 = 1$ and $e_4 = \sqrt{2} \approx 1.41$.

Separating PBE- Applying the Intuitive Criterion

- The most-costly separating PBE where the high-productivity worker chooses $e_H = e_4$ can only be sustained if:
 - the firm, upon observing off-the-equilibrium education levels in the interval $e_H \in [e_3, e_4)$, strictly below e_4 , believes it doesn't face a high-productivity worker, i.e., $\mu(\theta_H | e_H) < 1$, and thus pays her strictly less than θ_H .
- But are these off-the-equilibrium beliefs sensible? No!
- We show this using the **Cho and Kreps' (1987) Intuitive Criterion** following the six-step approach.

Cho and Kreps' (1987) Intuitive Criterion

Step 1 – Identify a PBE to test.

Consider a specific PBE, such as the most-costly separating PBE, where $(e_L, e_H) = (0, e_4)$.

Step 2 – Off-the-equilibrium message. Identify an off-the-equilibrium message for the worker, such as $e' \in [e_3, e_4)$, as depicted in Figure 12.6.




Figure 12.6. Applying the Intuitive Criterion to the separating PBEs.

Cho and Kreps' (1987) Intuitive Criterion

Step 3 – Types who profitably deviate. Find which types of worker can benefit by deviating to e' .

Low productivity worker has no incentive to deviate, because the cost of effort is too large for this type of worker even when the firm responds paying her the highest salary, $w(e') = \theta_H$. Formally,

$$\theta_L - 0 \geq \theta_H - c(e', \theta_L) \Rightarrow c(e', \theta_L) \geq \theta_H - \theta_L$$

Cho and Kreps' (1987) Intuitive Criterion

Step 3 – Types who profitably deviate. Find which types of worker can benefit by deviating to e' .

High productivity worker can benefit from choosing e' .

- Consider that firm keeps paying her highest salary $w(e') = \theta_H$
- If $e' < e_4$, she will deviate to e' . Formally,

$$\theta_H - c(e', \theta_H) \geq \theta_H - c(e_4, \theta_H) \Rightarrow c(e_4, \theta_H) \geq c(e', \theta_H)$$

which holds given $e_4 > e'$.

- Overall, this means that education level e' can only originate from the high-productivity worker.

Cho and Kreps' (1987) Intuitive Criterion

Step 4 – Restricting off-the-equilibrium beliefs.

- We can now restrict the off-the-equilibrium beliefs of the firm. If e' is observed, it can only originate from the high-productivity worker, i.e., $\mu(\theta_H|e') = 1$.

Step 5 – Updated Responses.

- Let us find the optimal response given the restricted belief $\mu(\theta_H|e') = 1$.
- As the firm is convinced of dealing with a high-productivity worker, it optimally responds paying $w(e') = \theta_H$.

Step 6 – Conclusion.

- Given the optimal response in step 5, high productivity worker has incentive to deviate from e_4 to e' .
- Therefore, the most-costly separating PBE $(e_L, e_H) = (0, e_4)$ violates the Intuitive Criterion.

Summary: Cho and Kreps' (1987) Intuitive Criterion

- A similar argument applies to any other separating PBE where $e_H \in (e_3, e_4]$ since only the high-productivity worker has incentives to deviate to $e' < e_4$.
- At the least-costly separating PBE, her choice is $e_H = e_3$, and it survives the Intuitive Criterion. How?
 - Both types of worker can benefit from deviating to education levels strictly below e_3 , implying that the firm cannot restrict its beliefs upon observing off-the-equilibrium education e' satisfying $e' < e_3$, keeping its beliefs unaltered relative to those in the separating PBE.
 - Only one separating PBE survives the Intuitive Criterion, namely, the least-costly separating PBE where $(e_L, e_H) = (0, e_3)$.

Separating PBE – Applying the D1 Criterion

- From section 11.3, we know that, in signaling games with only two sender types, the Intuitive and D1 Criterion have the same refinement power.
 - Same set of PBEs survives both refinement criteria.
- As a practice, show that the separating PBE in Figure 12.7, where $e_L = 0$ and $e_H = e_2$, violates the Intuitive Criterion.
- We next show that it also violates the D1 criterion.




Figure 12.7. Applying the D1 Criterion to the separating PBEs.

Separating PBE – Applying the D1 Criterion

Step 1 —Identify a PBE to test.

- Consider a specific PBE, such as the separating PBE where $(e_L, e_H) = (0, e_2)$.

Step 2 —Off-the-equilibrium message.

- Identify an off-the-equilibrium message for the worker, such as $e' \in [0, e_1)$, as depicted in Figure 12.7.




Figure 12.7. Applying the D1 Criterion to the separating PBEs.

Separating PBE – Applying the D1 Criterion

Step 3 — Type most likely to deviate.

- Find which worker type is most likely to benefit from deviating to e'
- $D(\theta_L, \hat{\Theta}, e')$ and $D(\theta_H, \hat{\Theta}, e')$ respectively denotes set of wages that make low productivity and high productivity worker better off after deviating to e' .
- Comparing $D(\theta_L, \hat{\Theta}, e')$ and $D(\theta_H, \hat{\Theta}, e')$, we see that $D(\theta_L, \hat{\Theta}, e')$ is a subset of $D(\theta_H, \hat{\Theta}, e')$, indicating that:
 - after deviating to e' , there are more wage offers that improve the equilibrium payoff of the high-productivity worker than those improving the low-productivity worker's.




Figure 12.7. Applying the D1 Criterion to the separating PBEs.

Separating PBE – Applying the D1 Criterion

Step 3 — Type most likely to deviate.

- We can, then, conclude that the high-productivity worker is more likely to send message e' ,
- which helps the firm update its off-the-equilibrium beliefs:
 - upon observing e' , $\mu(\theta_H|e') = 1$.




Figure 12.7. Applying the D1 Criterion to the separating PBEs.

Separating PBE – Applying the D1 Criterion

Step 3 — Type most likely to deviate.

- For completeness, the figure also includes deviations to off-the-equilibrium message e'' , where $e'' \in (e_1, e_2)$.
- The low-productivity worker would never benefit from deviating to e'' , entailing that the set of wage offers improving her equilibrium utility $D(\theta_L, \hat{\Theta}, e'')$, is in this case nil.




Figure 12.7. Applying the D1 Criterion to the separating PBEs.

Separating PBE – Applying the D1 Criterion

Step 3 — Type most likely to deviate.

- In contrast, sending e'' might be profitable for the high-productivity worker, as depicted in the figure.
- Therefore, upon observing e'' , the firm's updated beliefs are also $\mu(\theta_H|e'') = 1$.




Figure 12.7. Applying the D1 Criterion to the separating PBEs.

Separating PBE – Applying the D1 Criterion

Step 4 — Restricting off-the-equilibrium beliefs.

- We can now restrict the off-the-equilibrium beliefs of the firm.
- From Step 3, we found that, any deviation from the separating PBE $(e_L, e_H) = (0, e_2)$, where $e' < e_2$, is more likely to originate from the high- than the low-productivity worker,
 - so the firm's beliefs become $\mu(\theta_H|e') = 1$.

Step 5 — Updated Responses.

- Because the firm is convinced of dealing with a high-productivity worker, it optimally responds offering a salary $w(e') = \theta_H$.

Separating PBE – Applying the D1 Criterion

Step 6 — Conclusion.

- Given the optimal response found in Step 5, the high-productivity worker deviate from her equilibrium strategy of e_2 to e' .
- Therefore, separating PBE $(e_L, e_H) = (0, e_2)$ violates the D1 Criterion.
- Like in the Intuitive Criterion, one can show that all separating PBEs in this game can be eliminated using the D1 Criterion, except for the efficient (Riley) outcome, where:
 - the low-productivity worker acquires no education, $e_L = 0$, and
 - the high-productivity worker acquires the lowest education level that helps convey her type to the firm, $e_H = e_1$.

Pooling PBE

First Step – Strategy Profile.

- In a pooling strategy profile, both the high- and low-productivity worker choose education level e_P , where the subscript P denotes pooling equilibrium.

Second Step – Updating Beliefs.

- Upon observing education level e_P :
 - The firm's posterior beliefs coincide with its prior, i.e., $\mu(\theta_H|e_P) = p$.
 - Intuitively, the observation of education level e^P provides the firm with no additional information about the worker's productivity since all worker types acquire the same education.
- Upon observing the off-the-equilibrium education level $e \neq e_P$:
 - The firm cannot update its beliefs using Bayes' rule,
 - leaving its off-the-equilibrium beliefs unrestricted, that is $\mu(\theta_H|e) \in [0,1]$.

Pooling PBE

Third step – Optimal responses.

- Upon observing the pooling education level e^P , the firm optimally responds with salary

$$w(e^P) = \underbrace{p\theta_H + (1-p)\theta_L}$$

*Worker's expected
productivity $\equiv E[\theta]$*

- After observing any off-the-equilibrium education $e \neq e_P$, the firm responds with $w(e) \in [\theta_L, \theta_H]$

Technically,

$$w(e) = \mu(\theta_H|e)\theta_H + (1 - \mu(\theta_H|e))\theta_L$$

since its off-the-equilibrium beliefs are $\mu(\theta_H|e) \in [0,1]$.

Pooling PBE

Third step – Optimal responses.

- Figure 12.8 depicts, as an example, a wage schedule that satisfies the above two properties:
 - at the pooling education level e^P , the salary is $w(e^P) = E[\theta]$, and
 - for all other education levels $e \neq e_P$, the firm pays salaries that are bounded between θ_L and θ_H .




Figure 12.8. Pooling PBE - Example of wage schedule.

Pooling PBE

Fourth step – Optimal messages.

- Identify under which conditions both types of workers choose the same education level e^P

Low-productivity worker

- IC_L originates at $(\theta_L, 0)$, where the worker acquires no education and receives the lowest salary θ_L , and passes through point A , which represents $(w, e) = (E[\theta], e^P)$.




Figure 12.9. Pooling PBE – Low-productivity worker

Pooling PBE

Fourth step – Optimal messages.

Low-productivity worker

- This type of worker is, then, indifferent between:
 - identifying herself as a low-productivity worker (i.e., acquiring no education and receiving salary θ_L) and
 - acquiring the pooling level e^P .




Figure 12.9. Pooling PBE – Low-productivity worker

Pooling PBE

Fourth step – Optimal messages.

Low-productivity worker

- The IC_L curve in Figure 12.9 must satisfy

$$\theta_L - c(0, \theta_L) = E[\theta] - c(e^P, \theta_L)$$

and, given that $c(0, \theta_L) = 0$ by definition,

$$c(e^P, \theta_L) = E[\theta] - \theta_L$$

where $E[\theta] > \theta_L$ since $p > 0$.




Figure 12.9. Pooling PBE – Low-productivity worker

Pooling PBE

Fourth step – Optimal messages.

Low-productivity worker

- For this to be the case, we need the wage schedule $w(e)$ to lie weakly below IC_L :
 - since that implies that, a deviation towards any education level $e \neq e_P$ produces an overall utility lower than that at e^P for the low-productivity worker.




Figure 12.9. Pooling PBE – Low-productivity worker

Example 12.4. Finding a pooling education level

- Consider the parametric example of Example 12.1, where $\theta_H = 2$, $\theta_L = 1$, and the cost of education is $c(e, \theta_K) = \frac{e^2}{\theta_K}$.

- If we assume $p = \frac{1}{3}$, we obtain an expected productivity of

$$E[\theta] = \frac{1}{3}2 + \frac{2}{3}1 = \frac{4}{3}.$$

- The above equation becomes:

$$c(e^P, \theta_L) = E[\theta] - \theta_L \Rightarrow (e^P)^2 = \frac{4}{3} - 1$$

which yields a pooling education level of $e^P = 0.58$.

Pooling PBE

High-productivity worker

- Must have incentives to choose e^P rather than deviating to any other education level $e \neq e_P$.
- Figure 12.10a illustrates a wage schedule $w(e)$ that leaves no incentives to deviate from e^P to this type of worker.
- Deviation to e_2 yields a lower salary than $w(e^P) = E[\theta]$ and a higher education cost!
- Deviation to e_1 yields a lower salary and lower education cost, but it's not worthy (point C lies to the southeast of A).




Figure 12.10a. Pooling PBE – High-productivity worker

Pooling PBE

High-productivity worker

- Figure 12.10b superimposes:
 - Figure 12.9 from our analysis of the low-productivity worker, and
 - Figure 12.10a, from our discussion of the high-productivity worker.
- Wage schedule $w(e)$ must lie weakly below both IC_L and IC_H
 - for both types of workers to have incentives to choose e^P instead of deviating,
- The pooling education level lies at the point where IC_L and IC_H cross each other.
 - Alternatively, at the point where IC_L that originates at $(\theta_L, 0)$ crosses $E[\theta]$.




Figure 12.10b. Pooling PBE – Both worker types

Other Pooling PBE

- Are there any other pooling PBEs? Yes!
- If IC_L originates strictly above θ_L :
 - the crossing point of IC_L and IC_H happens closer to the origin, as depicted in Figure 12.11.
- A similar argument applies if we keep increasing the origin of IC_L until we reach a crossing point at $e^P = 0$.
- In this case,
 - IC_L originates at a height of $E[\theta]$.
 - Both types of workers acquire zero education, and they receive a salary equal to their expected productivity $E[\theta]$.
 - Same as in incomplete information game where workers cannot acquire education to signal their type.




Figure 12.11. Pooling PBE - Other equilibria.

Other Pooling PBE

- Therefore, we can summarize all pooling PBEs as (w, e) -pairs where both workers choose $e^P \in [0, e_A]$, where education e_A is the “most costly pooling PBE” and solves

$$\theta_L - c(0, \theta_L) = E[\theta] - c(e^P, \theta_L)$$

or, given that $c(0, \theta_L) = 0$ by definition,

$$c(e^P, \theta_L) = E[\theta] - \theta_L$$

- This equation just identifies the education level where the IC_L starting at a height of θ_L crosses the horizontal line $E[\theta]$ in Figure 12.10b.




Figure 12.11. Pooling PBE - Other equilibria.

Other Pooling PBE

- Intuitively, at this education level, the low-productivity worker is indifferent between:
 - the pooling education e^P , earning a salary $w(e^P) = E[\theta]$, and
 - a zero-education level receiving a salary of θ_L .
- In all these pooling PBEs, the firm responds with:
 - a salary $w(e^P) = E[\theta]$ upon observing e^P , and
 - a wage schedule $w(e)$ that lies below both IC_L and IC_H for all $e \neq e_P$.




Figure 12.11. Pooling PBE - Other equilibria.

Other Pooling PBE

- Therefore, we have a range of pooling PBEs:
 - from the least-costly pooling PBE where both types of workers acquire zero education, $e^P = 0$, to...
 - the most-costly pooling PBE where both acquire the highest education $e^P = e_A$.




Figure 12.11. Pooling PBE - Other equilibria.

Pooling PBE – Applying the Intuitive Criterion

- We identified under which conditions we can sustain pooling PBEs, but...
 - Do they survive the Intuitive Criterion?
No.
- Consider the off-the-equilibrium beliefs for any education level $e \neq e_P$.
- The condition that $w(e)$ lies below IC_L and IC_H for all $e \neq e_P$ means that:
 - upon observing a deviation from e^P , the firm infers that such a deviation is not likely originating from the high-productivity worker
 - thus, paying a relatively low wage as depicted by the height of $w(e)$ in the right-hand side of Figure 12.12.




Figure 12.12. Applying the Intuitive Criterion to Pooling PBEs.

Pooling PBE – Applying the Intuitive Criterion

- This off-the-equilibrium beliefs are, of course, not sensible:
 - one could argue that deviations towards high education levels (higher than e_P in figure 12.12)
 - are more likely to stem from the high-type worker.
- We next show this point, following the six-steps approach to the Intuitive Criterion.




Figure 12.12. Applying the Intuitive Criterion to Pooling PBEs.

Pooling PBE – Applying the Intuitive Criterion

Step 1 – Identify a PBE to test.

Consider a specific PBE, such as the most-costly pooling PBE where $e^P = e_A$.

Step 2 – Off-the-equilibrium message.

Identify an off-the-equilibrium education, such as $e' > e_A$ in figure 12.12.




Figure 12.12. Applying the Intuitive Criterion to Pooling PBEs.

Pooling PBE – Applying the Intuitive Criterion

Step 3 – Types who profitably deviate.

Low-productivity worker: cannot benefit.

Even if the firm responds paying her the highest salary, $w(e') = \theta_H$, the cost of effort is too large for this type of worker.

Formally,

$$E[\theta] - c(e_A, \theta_L) \geq \theta_H - c(e', \theta_L)$$

$$\Rightarrow c(e', \theta_L) - c(e_A, \theta_L) \geq \theta_H - E[\theta]$$

Intuitively, the additional cost that the worker must incur offsets the wage increase she experiences.




Figure 12.12. Applying the Intuitive Criterion to Pooling PBEs.

Pooling PBE – Applying the Intuitive Criterion

Step 3 – Types who profitably deviate.

High-productivity worker: can benefit from choosing e' if

$$E[\theta] - c(e_A, \theta_H) \leq \theta_H - c(e', \theta_L)$$

$$\Rightarrow c(e', \theta_H) - c(e_A, \theta_H) \leq \theta_H - E[\theta]$$

- Intuitively, if the firm responds to the higher education level e' by paying her the highest salary $w(e') = \theta_H$...
 - her wage increase (right-hand side) would offset the additional education cost (left-hand side).




Figure 12.12. Applying the Intuitive Criterion to Pooling PBEs.

Pooling PBE – Applying the Intuitive Criterion

Step 3 – Types who profitably deviate.

- At this point, we can combine the above two inequalities to obtain

$$\begin{aligned} c(e', \theta_H) - c(e_A, \theta_H) &\leq \theta_H - E[\theta] \\ &\leq c(e', \theta_L) - c(e_A, \theta_L) \end{aligned}$$

$$\Rightarrow c(e', \theta_H) - c(e_A, \theta_H) \leq c(e', \theta_L) - c(e_A, \theta_L)$$

- which is true by definition (single-crossing property in discrete settings).
- Overall, this means that education level e' can only originate from the high-productivity worker.




Figure 12.12. Applying the Intuitive Criterion to Pooling PBEs.

Pooling PBE – Applying the Intuitive Criterion

Step 4 – Restricting off-the-equilibrium beliefs.

- We can now restrict the firm's off-the-equilibrium beliefs as follows:
 - If education level e' is observed, it can only originate from the high-productivity worker, i.e., $\mu(\theta_H|e') = 1$.

Step 5 – Updated Responses.

- As the firm is convinced of dealing with a high-productivity worker, it optimally responds paying $w(e') = \theta_H$.

Pooling PBE – Applying the Intuitive Criterion

Step 6 – Conclusion.

- Given the optimal response found in Step 5, we can see that the high-productivity worker has incentives to deviate from her equilibrium strategy of e_A to e' .
- Therefore, the most-costly pooling PBE $e^P = e_A$ violates the Intuitive Criterion.

Pooling PBE – Applying the Intuitive Criterion

Step 6 – Conclusion.

- A similar argument applies to all other pooling PBE where $e^P < e_A$, including the least-costly pooling PBE where $e^P = 0$.
- Hence, no pooling PBE in the labor market signaling game survives the Intuitive Criterion,
 - implying that the only PBE surviving the Intuitive Criterion is the least-costly separating PBE where $(e_L, e_H) = (0, e_3)$.

Can signaling be welfare improving?

- A natural question is whether the worker is better off when:
 - she uses education to signal her type (in the least-costly separating PBE found above) than
 - when such signal is not available.
- Figure 12.13 compares the indifference curve that each type of worker reaches in these two information settings.




Figure 12.13. Utility comparison across information settings.

Can signaling be welfare improving?

Low-productivity worker

- Unambiguously worse off with signaling, where she acquires zero education but receives the lowest salary θ_L , than...
 - without signaling, where she still acquires no education and earns a higher salary $E[\theta]$.
- Graphically, the indifference curve passing through point $(0, E[\theta])$ reaches a higher utility level than IC_L .
- This result holds regardless of the specific probability that the worker is of high-productivity, p ,
 - which graphically means regardless of the height of $E[\theta]$ in the (θ_L, θ_H) interval.




Figure 12.13. Utility comparison across information settings.

Can signaling be welfare improving?

High-productivity worker

- In contrast, the high-type worker is better off with signaling, where she reaches IC_H , than...
 - without signaling, where she acquires no education and earns a salary $E[\theta]$,
 - only if $E[\theta]$ is sufficiently low, which occurs when p is relatively low.
- Intuitively, this type of worker is better off acquiring education, despite its cost, and earning the highest wage θ_H , than
 - not investing in education and receiving $E[\theta]$ when this wage is sufficiently low,
 - which occurs when the firm believes that the high-productivity worker is relatively unlikely.




Figure 12.13. Utility comparison across information settings.

Can signaling be welfare improving?

High-productivity worker

- If, instead, probability p (as then $E[\theta]$) is sufficiently high,
 - the high-type worker is better off in the setting where education cannot be used as a signal to firms.




Figure 12.13. Utility comparison across information settings.

Example 12.5 Firms left in the dark?

- Consider again the setting in examples 12.1-12.3, where $\theta_H = 2$ and $\theta_L = 1$.
- Recall, in the least-costly separating PBE, the high-productivity worker chooses education $e_3 = 1$. Therefore, her utility in equilibrium is:

$$u_H = w - \frac{e^2}{\theta_H} = \theta_H - \frac{1^2}{\theta_H} = 2 - \frac{1}{2} = \frac{3}{2}$$

- However, when signaling is not available, she earns a salary

$$E[\theta] = p2 + (1 - p)1 = 1 + p$$

yielding a utility of $1 + p$ since in this setting she acquires zero education.

- Therefore, this type of worker is better off in the environment where education cannot be used a signal to firms if $1 + p > \frac{3}{2}$, or solving for p , when $p > \frac{1}{2}$.
- Intuitively, the frequency of high types is so high that the low-productivity worker benefits from firms being “in the dark” about her type.

What if the sender has three types?

- Consider an extension of the previous signaling game, where the sender can now have three equally likely types, θ_L , θ_M , and θ_H , where $\theta_L < \theta_M < \theta_H$.
- Outline:
 1. Identify the set of separating PBEs
 2. Show that applying the Intuitive Criterion in this setting would have no bite, meaning that no separating PBE violates the Intuitive Criterion
 3. Show that the D1 Criterion helps us delete some separating PBEs in this signaling game with three sender types. Indeed, only the least-costly separating PBE (Riley outcome) survives the application of the D1 Criterion.

Separating PBEs

- For compactness, we describe here the set of separating PBEs.
- The set of all separating PBEs satisfies $(e_L, e_M, e_H) = (0, e_M^*, e_H^*)$ where $e_M^* \in [e_1^M, e_2^M]$, $e_H^* \in [e_1^H, e_2^H]$.
- Wages satisfy $w(e_K) = \theta_i$ for every $i = \{L, M, H\}$ after observing equilibrium education levels $(0, e_M^*, e_H^*)$ and $w(e)$ after observing the off-the-equilibrium education levels $e \neq e_i$, where $w(e)$ lies below the indifference curve of all worker types, as depicted in Figure 12.14.




Figure 12.14. Separating PBEs with three worker types.

Separating PBEs

- Indifference curves in Figure 12.14 correspond to a specific separating PBE:
 - $(0, e_2^M, e_2^H)$,
 - where both the medium- and high-type worker choose the most-costly education level in their corresponding intervals.
- The high-type worker has no incentives to deviate to the medium-type worker's education level, e_2^M , which would yield a wage of $w = \theta_M$.
- This wage and education pair, (θ_M, e_2^M) , would provide her with the same utility as her equilibrium education, e_2^H , as depicted IC_H in the figure.




Figure 12.14. Separating PBEs with three worker types.

Separating PBE – Applying the Intuitive Criterion

Step 1 – Identify a PBE to test.

Consider a specific PBE, such as the separating equilibrium described above $(e_L, e_M, e_H) = (0, e_M^*, e_H^*)$

Step 2 – Off-the-equilibrium message.

Identify an off-the-equilibrium education, e , where $e \in (\hat{e}, e_H^*)$ as depicted in Figure 12.15.




Figure 12.15. Applying the Intuitive Criterion with three worker types.

Separating PBE – Applying the Intuitive Criterion

Step 3 – Types who profitable deviate.

Find which types of worker can benefit by deviating to e :

Low type

- Utility in equilibrium is

$$u_L^* = \theta_L - c(0, \theta_L) = \theta_L$$

given that $c(0, \theta_L) = 0$.

- Highest utility that she can earn by deviating to e is

$$\theta_H - c(e, \theta_L)$$

- Therefore, $\theta_H - c(e, \theta_L) < \theta_L$ holds if $c(e, \theta_L) > \theta_H - \theta_L$, meaning that the cost of education exceeds highest possible salary gain.
- Then, θ_L has no incentive to deviate to $e \in (\hat{e}, e_H)$.

Separating PBE – Applying the Intuitive Criterion

Step 3 – Types who profitable deviate.

Find which types of worker can benefit by deviating to e :

Medium type

- Medium-productivity workers could send such a message $e \in (\hat{e}, e_H)$ because

$$\theta_M - c(e_M^*, \theta_M) < \theta_H - c(e, \theta_M)$$

or alternatively, $c(e, \theta_M) - c(e_M^*, \theta_M) < \theta_H - \theta_M$, reflecting that the cost of acquiring $e - e_M^*$ additional years of education is offset by her salary increase when the firm offers her a high-productivity wage $w(e) = \theta_H$.

Separating PBE – Applying the Intuitive Criterion

Step 3 – Types who profitable deviate.

Find which types of worker can benefit by deviating to e :

High type

- A similar argument applies to the high-productivity worker, who can benefit by deviating to e since

$$\theta_H - c(e_H^*, \theta_H) < \theta_H - c(e, \theta_H)$$

given that $c(e_H^*, \theta_H) > c(e, \theta_H)$.

- Intuitively, by deviating towards e , she does not modify her salary, but saves on education costs.

Separating PBE – Applying the Intuitive Criterion

Step 4 – Restricting Off-the-equilibrium beliefs.

- If $e \in (\hat{e}, e_H)$ is observed, it can only originate from the high- or medium-productivity worker, but not from the low-types.
- Technically, we say that deviations towards $e \in (\hat{e}, e_H)$ are:
 - not equilibrium dominated for the high-and medium-productivity workers,
 - but equilibrium dominated for the low-productivity worker.
- Therefore, upon observing $e \in (\hat{e}, e_H)$, the firm concentrates its beliefs on these two worker types:

$$\Theta^{**}(e) = \{\theta_M, \theta_H\} \quad \text{for all } e \in (\hat{e}, e_H)$$

Separating PBE – Applying the Intuitive Criterion

Step 5 – Updating Responses.

- Since $\Theta^{**}(e) = \{\theta_M, \theta_H\}$ for all $e \in (\hat{e}, e_H)$, the firm's best response to this education level will be a wage offer somewhere in between $w(e) = \theta_M$ and $w(e) = \theta_H$.
- This is different from what we saw in models with two worker types, where we could put full probability weight on a single worker type.

Separating PBE – Applying the Intuitive Criterion

Step 5 – Updating Responses.

Medium-productivity worker

- Lowest wage she can receive is $w(e) = \theta_M$, implying that her indifference curve will pass through point A,
 - being *below* (to the southeast) of her indifference curve in equilibrium, IC_M .
- Alternatively, even after the firm has restricted its beliefs so that deviations to e must originate from the high- and medium-type...
 - the medium-prod. worker cannot guarantee that her equilibrium utility will be strictly improved when deviating to e .




Figure 12.15 Applying the Intuitive Criterion with three worker types

Separating PBE – Applying the Intuitive Criterion

Step 5 – Updating Responses.

High-productivity worker

- Lowest wage she can earn after deviating to e is $w(e) = \theta_M$, yielding an indifference curve that passes through point A .
 - This indifference curve is also *below* her indifference curve in equilibrium payoff, IC_H , providing her with a lower utility level.
- Then, she cannot guarantee that her equilibrium utility will improve when deviating to e .




Figure 12.15 Applying the Intuitive Criterion with three worker types

Separating PBE – Applying the Intuitive Criterion

Step 6 – Conclusion.

- Given the optimal responses found in Step 5, we see that no worker type has incentives to deviate from her equilibrium strategy to e .
- In conclusion, the separating PBE specified in Figure 12.15 survives the Intuitive Criterion.
- Take-away lesson:
 - The application of the Intuitive Criterion does not necessarily eliminate separating PBEs in models with more than two sender types (e.g., workers).

Separating PBE – Applying the D1 Criterion

- D1 Criterion to the rescue!
- We now show that the separating PBE examined in the previous section violates D1 Criterion.
- For compactness, we do not reproduce all steps here, as they all coincide for Step 3.
- When deviating to e' the worker with the largest set of wage offers that can improve her equilibrium utility is the medium-productivity worker,
 - i.e., $D(\theta_M, \hat{\theta}, e')$ is longer than that of the other worker types, $D(\theta_L, \hat{\theta}, e')$ and $D(\theta_H, \hat{\theta}, e')$




Figure 12.16. Applying the D1 Criterion with three worker types.

Separating PBE – Applying the D1 Criterion

- **Step 4.** Given our results in Step 3, upon observing e , the firm believes that this education level must originate from the medium-productivity worker alone.
- **Step 5.** Given its restricted beliefs, the firm offers a wage $w(e) = \theta_M$ after observing e .
- **Step 6.** Finally, we obtain that the medium-productivity worker's deviation payoff is

$$w(e) - c(e, \theta_M) = \theta_M - c(e, \theta_M)$$

while her equilibrium payoff, $u_M^* = \theta_M - c(e_M, \theta_M)$.

- And given that $e < e_M$ and $c_e(e, 0) > 0$, we have that $c(e, \theta_M) < c(e_M, \theta_M)$; which implies

$$\theta_M - c(e, \theta_M) > \theta_M - c(e_M, \theta_M)$$

Separating PBE – Applying the D1 Criterion

- Hence, we have found that the medium-productivity worker has incentives to deviate towards e ,
 - entailing that the separating PBE $(e_L, e_M, e_H) = (0, e_M^*, e_H^*)$ violates the D1 Criterion.
- Repeating this process for all off-the-equilibrium messages...
 - we can delete all separating PBEs, except for the efficient (Riley) equilibrium outcome depicted in Figure 12.17.




Figure 12.17. Unique separating PBE surviving the D1 Criterion with three types.

Appendix

Appendix - Equilibrium Refinements

- We present a more formal definition of the Intuitive and D1 Criterion, following Fudenberg and Tirole (2002, pp. 452-453).

Intuitive Criterion

- The first step focuses on those types of senders who can improve their equilibrium utility level by deviating.
- Formally, for any off-the-equilibrium message, m , we construct a subset of types $\Theta^{**}(m) \subset \Theta$ for which sending m is not equilibrium dominated. That is,

$$\Theta^{**}(m) = \left\{ \theta \in \Theta \mid u_i^*(\theta) \leq \max_{a \in A^*(\Theta, m)} u_i(m, a, \theta) \right\}$$

- Intuitively, this expression states that, from all types in Θ , we restrict our attention to those types who *could* improve their equilibrium utility level, $u_i^*(\theta)$.
 - Note that emphasis on “could” since the term $\max_{a \in A^*(\Theta, m)} u_i(m, a, \theta)$ represents the highest payoff that θ -type can achieve by deviating to the off-the-equilibrium message m .

Intuitive Criterion

- In short, we can interpret $\Theta^{**}(m)$ as the subset of senders who *could improve* their equilibrium utility by deviating from m^* to m .
- Given $\Theta^{**}(m)$, we now check if the following inequality holds for some type $\theta \in \Theta^{**}(m)$,

$$\min_{a \in A^*(\Theta^{**}(m), m)} u_i(m, a, \theta) > u_i^*(\theta)$$

- Intuitively, this inequality states that:
 - once beliefs are restricted to $\Theta^{**}(m)$,
 - there is at least one sender type who prefers to deviate to a message m ,
 - since this deviation provides her with a strictly higher utility level than her equilibrium message, m^* , *regardless* of the response of the receiver,
 - i.e., even if the receiver responds in the least beneficial way for the sender, as captured by $\min_{a \in A^*(\Theta^{**}(m), m)} u_i(m, a, \theta)$.

Intuitive Criterion

- In settings where the first step of the Intuitive Criterion helps us restrict off-the-equilibrium beliefs to a single type:
 - the application of the second step is unnecessary.
 - specific response from the sender.
 - as the sender type has incentives to deviate from her equilibrium message.
- However, in contexts where the first step leaves us with two or more sender types in $\Theta^{**}(m)$:
 - off-the-equilibrium beliefs for these types cannot be restricted to a single type,
 - not a specific response from the sender,
 - and the application of the second step is still necessary.
- Formally, an equilibrium strategy profile (m^*, a^*) *violates* the Intuitive Criterion if:
 - there is a type of agent θ and an action m such that condition (2) is satisfied.
 - Otherwise, we say that the equilibrium strategy profile *survives* the Intuitive Criterion.

D1 Criterion

- The D1 Criterion considers that, among all potential deviators, the receiver restricts her beliefs to only those types of senders who *most likely* send the off-the-equilibrium message.
- In particular, we now identify the sender for whom most of the responder's actions (in a setting with discrete actions) or the sender with the largest set of responses (in a context with continuous actions) provide a payoff strictly higher than her equilibrium payoff.
- Formally, for any off-the-equilibrium message m , let us define

$$D(\theta, \hat{\Theta}, m) = \bigcup_{\mu: \mu(\hat{\Theta}|m)=1} \{a \in MBR(\mu, m) | \mu_i^*(\theta) < u_i(m, a, \theta)\}$$

as the set of mixed best responses (MBR) of the receiver for which the θ — *type* of sender is *strictly better-off* deviating towards message m than sending her equilibrium message m^* .

D1 Criterion

- Note that $\mu(\hat{\Theta}|m) = 1$ in the previous expression represents that the receiver believes that message m only comes from types in the subset $\hat{\Theta} \in \Theta$.

- Let us also define

$$D^{\circ}(\theta, \hat{\Theta}, m) = \bigcup_{\mu: \mu(\hat{\Theta}|m)=1} \{a \in MBR(\mu, m) | \mu_i^*(\theta) < u_i(m, a, \theta)\}$$

as the set of MBR of the receiver that make the θ – *type indifferent* between deviating towards m and sending her equilibrium message m^* .

D1 Criterion

- First, we say that θ – *type* can be *deleted* if there is another θ' – *type* such that, when the off-the-equilibrium message m is observed
$$\left[D(\theta, \hat{\Theta}, m) \cup D^\circ(\theta, \hat{\Theta}, m) \right] \subset D(\theta', \hat{\Theta}, m)$$
- That is, for a given message m , the set of receiver's actions which make the θ' – *type* of sender strictly better off, $D(\theta', \hat{\Theta}, m)$, is larger than the set of actions making the θ – *type* of sender strictly better off, $D(\theta, \hat{\Theta}, m)$, or indifferent, $D^\circ(\theta, \hat{\Theta}, m)$.
- Intuitively, after receiving message m , there are more best responses of the receiver that improve the θ' – *type's* equilibrium payoff than there are for the θ – *type*.

D1 Criterion

- As a consequence, the $\theta' - type$ is the sender who is the most likely to deviate from her equilibrium message m^* to the off-the-equilibrium message m .
- We continue this comparison for all types of senders, deleting those for which there is another type of sender who is more likely to deviate towards m .
- Finally, the set of types that cannot be deleted after using this procedure is denoted by $\Theta^{**}(m)$.

D1 Criterion

- Once the set of types has been restricted to $\Theta^{**}(m)$, the Intuitive and the D1 Criterion proceed similarly, seeking to find at least a sender type $\theta \in \Theta^{**}(m)$ with incentives to deviate even if the receiver responds in the least beneficial way for the sender, that is,

$$\min_{a \in A^*(\Theta^{**}(m), m)} u_i(m, a, \theta) > u_i^*(\theta)$$

which coincides with the second expression in our application of the Intuitive Criterion.

Chapter 13: Cheap talk games

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- So far we signals that were costly for the sender
 - Years of education.
- We will now consider settings where signals are costless
 - Such signals are called “cheap talk”, since “talk is cheap” or costless.
- Some examples are
 - A lobbyist (sender) informing a member of Congress (receiver) about the situation of an industry she works for,
 - An investment banker (sender) recommending a client (receiver) which stocks to purchase.
- In chapters 10-12, we assumed utility to be a function of the message. Mathematically, $u(m, a, \theta)$, where m is a costly signal or message.


Introduction

- In this chapter,

$$u(m, a, \theta) = u(m', a, \theta) \text{ for every } m \neq m'$$


- This property allows us to present utility more compactly as $u(a, \theta)$.
- Main goal: To identify if separating PBE is still possible.
- In previous chapters, to sustain separating PBEs, we need the cost of sending messages to differ across sender types.
- In this chapter, we show that separating PBEs can still emerge, but sender and receiver preferences have to be sufficiently similar.

Cheap talk with discrete messages and responses




- Nature determines the sender's type (high or low).
- Sender observes this and sends message m_1 or m_2 . Receiver only observes the signal.
- Receiver response with a , b or c after observing m_1 (or a' , b' or c' after observing m_2).
- Payoffs on the right and left coincide, because m_1 and m_2 do not affect utility (though they may affect the response of the receiver).
- Represents a cheap-talk setting because the sender's payoff is unaffected by her message.

Separating PBE




1. *Specifying a strategy profile.* Consider (m_1, m'_2) .
2. *Bayes' rule.* Upon observing m_1 (m'_2), the receiver believes that this message must originate from the high (low) type, entailing that $\mu(H|m_1) = 1$ at the top left side ($\mu(H|m_2) = 0$ at the top right side).
3. *Optimal response.* Receiver responds with:
 - b upon observing m_1 , since it yields a payoff of $4 > 3$ (payoff from a) > 0 (payoff from c)
 - c' upon observing m'_2 , since it yields a payoff of $5 > 4$ (payoff from a') > 0 (payoff from b')

Separating PBE




4. *Optimal messages.* Based on the receiver's optimal responses:
 - *High type.* If she sends a message of m_1 , as prescribed in this strategy profile, she earns a payoff of 2 (since m_1 is subsequently responded with b), which exceeds her payoff from deviating towards m_2 , 1, where the receiver responds with c' .
 - *Low type.* If she sends a message of m'_2 , as prescribed in this strategy profile, she obtains a payoff of 3 (as m'_2 is responded with c'), which exceeds her payoff from deviating towards m_1 , 1, where the receiver responds with b .
5. *Summary.* Here, (m_1, m'_2) can be supported as a separating PBE.

Pooling PBEs



1. *Specifying a strategy profile.* Consider (m_1, m'_1) .
2. *Bayes' rule.* Upon observing m_1 , the receiver's posterior beliefs coincide with her priors, entailing that $\mu(H|m_1) = 0.6$. Upon observing m_2 , Bayes' rule does not help the receiver update her belief, which remains generic at $(\mu(H|m_2) \in [0,1])$.

Pooling PBEs



3. *Optimal response.* Receiver responds with:


- Upon observing m_1 , responds with a (the expected utility for this is the highest of the three at 3.4).

$$EU_R(a) = (0.6 \times 3) + (0.4 \times 4) = 3.4$$

$$EU_R(b) = (0.6 \times 4) + (0.4 \times 0) = 2.4$$

$$EU_R(c) = (0.6 \times 0) + (0.4 \times 5) = 2.0$$

Pooling PBEs



3. *Optimal response.* Receiver responds with:

- Upon observing m_2 , we need to find expected payoffs from each response:


$$EU_R(a') = (\mu \times 3) + ((1 - \mu) \times 4) = 4 - \mu$$

$$EU_R(b') = (\mu \times 4) + ((1 - \mu) \times 0) = 4\mu$$

$$EU_R(b) = (\mu \times 0) + ((1 - \mu) \times 5) = 5(1 - \mu)$$

- Plot the three lines of $4 - \mu$, 4μ , and $5(1 - \mu)$; all as a function of μ .
 - This helps us see their crossing points.

Pooling PBEs



3. *Optimal response.* Receiver responds with:

- Upon observing m_2 , we need to find expected payoffs from each response:


$$EU_R(a') = (\mu \times 3) + ((1 - \mu) \times 4) = 4 - \mu$$

$$EU_R(b') = (\mu \times 4) + ((1 - \mu) \times 0) = 4\mu$$


$$EU_R(b) = (\mu \times 0) + ((1 - \mu) \times 5) = 5(1 - \mu)$$

- This implies that the best responses are

- c' if $\mu > \frac{1}{4}$,
- a' if $\frac{1}{4} < \mu \leq \frac{4}{5}$, and
- b' otherwise.




Pooling PBEs



4. *Optimal messages.* Based on the receiver's optimal responses:
 - *High type.* If she sends a message of m_1 , as prescribed in this strategy profile, she earns a payoff of 4 (since m_1 is subsequently responded with a), which exceeds her payoff from deviating towards m_2 , 2, where the receiver responds with b' .
 - *Low type.* If she sends a message of m'_1 , as prescribed in this strategy profile, she obtains a payoff of 2 (as m'_1 is responded with a), which exceeds her payoff from deviating towards m'_2 , 1, where the receiver responds with b' .
5. *Summary.* Here, (m_1, m'_1) can be supported as a separating PBE with responses (a, b') and equilibrium beliefs $\mu(H|m_1) = 0.6$ and off-the-equilibrium beliefs $\mu(H|m_2) > \frac{4}{5}$.

Cheap talk with discrete messages but continuous responses



- We now allow for continuous responses by the receiver.
- After observing her own type privately (θ_H or θ_L), the lobbyist (sender), chooses a message to send to the politician which message to send (still binary).
- Upon observing this message, the politician responds with a policy $p > 0$.

Quadratic loss functions - Sender

- The payoffs of the government (politician) follow a quadratic loss function, as in Crawford and Sobel (1982)

$$U_G(p, \theta) = -(p - \theta)^2$$

- which becomes zero when the politician chooses policy $p = \theta$,
- but is negative otherwise (both when $p < \theta$ and when $p > \theta$, as depicted in the figure two slides from now).

Quadratic loss functions - Receiver

- Similarly, the lobbyist's utility is given by a quadratic loss function

$$U_L(p, \theta) = -(p - (\theta + \delta))^2$$

- which becomes zero when the politician chooses $p = \theta + \delta$, that we can refer to as the lobbyist's ideal policy.
 - Otherwise, the lobbyist's payoff is negative (see figure in next slide).
- *Cases:*
 - When $\delta = 0$, the utility functions of both lobbyist and politician coincide, and we can say that their preferences are aligned.
 - Otherwise, the lobbyist's ideal policy is $p = \theta + \delta$, exceeding the politician's, $p = \theta$.
 - This explains why $\delta > 0$ is known as the lobbyist's "bias" relative to the politician.

Quadratic loss function (graphically)






Figure 13.5 Quadratic loss function for each player

Separating PBE



1. *Specifying a strategy profile.* Consider (θ_H, θ_L) .
2. *Bayes' rule.*
 - Upon observing θ_H , the politician (receiver) believes that this message must originate from the high type lobbyist (sender), entailing that $\mu(\theta_H|\theta_H) = 1$ and $\mu(\theta_L|\theta_H) = 0$.
 - Similarly, upon observing θ_L , the politician's beliefs are $\mu(\theta_H|\theta_L) = 0$ and $\mu(\theta_L|\theta_L) = 1$.

Separating PBE



3. *Optimal response.* Politician responds with responds with:

- $p = \theta_H$ upon observing θ_H , as this policy minimizes her quadratic loss, yielding a payoff of zero.
- $p = \theta_L$ upon observing θ_L , also minimizes her quadratic loss, yielding a payoff of zero.

Separating PBE


4. *Optimal messages.* Based on the receiver's optimal responses:

- *High type.* If she sends a message of θ_H , as prescribed in this strategy profile, she earns a payoff of $-\delta^2$, which exceeds her payoff from deviating towards θ_L , $-(\theta_L - \theta_H - \delta)^2$.
 - Role of bias.
- *Low type.* If she sends a message of θ_L , as prescribed in this strategy profile, she obtains a payoff of $-\delta^2$, which exceeds her payoff from deviating towards θ_H , $-(\theta_H - (\theta_L + \delta))^2$ iff $\delta \leq \frac{\theta_H - \theta_L}{2}$.

5. *Summary.* Hence, (θ_H, θ_L) can be supported as a separating PBE if $\delta \leq \frac{\theta_H - \theta_L}{2}$, or if the lobbyist's bias to report the state of nature is sufficiently small.

- This happens when her preferences and the politician's are sufficiently *aligned*.

Pooling PBEs




1. Specifying a strategy profile. Consider (θ_H, θ_H) .

2. Bayes' rule.

- Upon observing θ_H , the politician's beliefs coincide with her priors, $\mu(\theta_H | \theta_H) = 0.5$.
- Upon observing θ_L , Bayes' rule does not help the receiver update her belief, which remains generic at $(\mu(\theta_L | \theta_L) \in [0, 1])$.

Pooling PBEs



3. *Optimal response.* Receiver responds with:

- $p = \frac{\theta_H + \theta_L}{2}$ upon observing θ_H , from $\max_{p \geq 0} \underbrace{-\frac{1}{2}(p - \theta_H)^2}_{\text{if } \theta = \theta_H} - \underbrace{\frac{1}{2}(p - \theta_L)^2}_{\text{if } \theta = \theta_L}$
- $p = \mu\theta_H + (1 - \mu)\theta_L$ upon observing θ_L , from $\max_{p \geq 0} \underbrace{-\mu(p - \theta_H)^2}_{\text{if } \theta = \theta_H} - \underbrace{(1 - \mu)(p - \theta_L)^2}_{\text{if } \theta = \theta_L}$

which can also be interpreted as the expected state of nature, given the politician's off-the-equilibrium beliefs μ and $(1 - \mu)$ on the high and low state occurring, respectively.

Pooling PBEs

4. *Optimal messages.* From our results in Step 3, we now identify the lobbyist's optimal messages.
- *High type.* The high type sends θ_H and this holds for all values of δ . Recall that this is due to the lobbyist's upward bias.

Pooling PBEs

- *Low type.*

- Chooses between sending message θ_H , which induces that the politician responds with policy $p = \frac{\theta_H + \theta_L}{2}$ as shown in 3(a); or
- Deviating towards the off-the-equilibrium message θ_L , which induces the politician to respond with policy $p = \mu\theta_H + (1 - \mu)\theta_L$ as shown in 3(b). This holds if and only if

$$-\left[\frac{\theta_H + \theta_L}{2} - (\theta_L + \delta)\right]^2 \geq -[(\mu\theta_H + (1 - \mu)\theta_L) - (\theta_L + \delta)]^2$$

After rearranging, it simplifies to $\mu \geq \frac{1}{2}$.

This means that, when the receiver believes the off-the-equilibrium message of θ_L is more likely to originate from the high-type sender, the low-type lobbyist has incentives to choose the pooling message θ_H .

Pooling PBEs

- This includes the case in which message θ_L is responded with policy $p = \theta_H$, i.e., $\mu = 1$.
- Importantly, this entails that the pooling PBE can be sustained for all values of δ .
- This is a typical result in cheap-talk games: while the separating PBE requires sender and receiver to exhibit similar preferences (low δ), the pooling PBE can be sustained regardless of players' preference divergence (for all δ).
- Informally, concealing information from the receiver can easily arise in equilibrium, but conveying it requires more demanding conditions.

Pooling PBEs

5. Summary. From step 4, we found that no sender type has incentives to deviate from the pooling strategy profile, (θ_H, θ_H) , sustaining it as a PBE if $\mu \geq \frac{1}{2}$, for all value of δ .

In this PBE, the politician:

- upon observing message θ_H in equilibrium, holds beliefs $\mu(\theta_H, \theta_H) = \frac{1}{2}$, responding with $p = \frac{\theta_H + \theta_L}{2}$; and
- upon observing θ_L off-the-equilibrium path, her beliefs are unrestricted, $\mu(\theta_L, \theta_L) = \mu$, responding with policy $p = \mu\theta_H + (1 - \mu)\theta_L$.


Cheap talk with continuous messages and responses

- We now extend our cheap talk model to allow for both
 - Continuous messages.
 - Continuous responses.
 - In addition, we allow for the state of nature to be continuous, i.e., $\theta \sim U[0,1]$.

Cheap talk with continuous messages and responses


- Our discussion, based on Crawford and Sobel (1982), helps us confirm that separating strategy profiles can emerge in equilibrium if the sender's and receiver's preferences are relatively aligned.
- This more general setting, however, allows us a new result:
 - That the “quality of information,” understood as the number of different messages that the lobbyist sends, also depends on the players' preference alignment.

Separating PBE



1. *Specifying a strategy profile.* The lobbyist sends message m_k when $\theta \in [\theta_{k-1}, \theta_k)$, where $k \in \{1, 2, \dots, N\}$.
2. *Bayes' rule.*
 - Upon observing m_k , the politician (receiver) believes that:
 - $\mu(\text{Interval } k | m_k) = 1$ and
 - $\mu(\text{Interval } j | m_k) = 0$ for all $j \neq k$.

Separating PBE



3. *Optimal responses.* Politician responds with responds with:

- $p = \frac{\theta_{k-1} + \theta_k}{2}$ upon observing m_k , which minimizes the quadratic loss in that interval.
- Derived from:

- $\max_{p \geq 0} -E[(p - \theta)^2]$ such that $\theta \in [\theta_{k-1}, \theta_k] \Rightarrow \max_{p \geq 0} -\left(p - \frac{\theta_{k-1} + \theta_k}{2}\right)^2$

- Differentiating with respect to policy p , yields:

$$-2 \left(p - \frac{\theta_{k-1} + \theta_k}{2}\right) (-1) = 2 \left(p - \frac{\theta_{k-1} + \theta_k}{2}\right) = 0 \Rightarrow p = \frac{\theta_{k-1} + \theta_k}{2} \text{ (optimal policy)}$$

- In other words, after receiving message m_k , the politician (receiver) responds with a policy that coincides with the expected state of the nature in this interval, $p = \frac{\theta_{k-1} + \theta_k}{2}$.

Separating PBE

4. *Optimal messages.* To check optimal messages, without loss of generality, it suffices to show that the k -th type sender sends neither m_{k-1} or m_{k+1} (in the intervals immediately below and above interval k , respectively).

- *No incentives to overreport* – occurs if

$$\underbrace{-\left(\frac{\theta_{k-1}+\theta_k}{2} - (\theta_k + \delta)\right)^2}_{\text{Utility from sending message } m_k} \geq \underbrace{-\left(\frac{\theta_{k+1}+\theta_k}{2} - (\theta_k + \delta)\right)^2}_{\text{Utility from sending message } m_{k+1}}$$

Term $(\theta_k + \delta)$ lies to right-hand side of $\frac{\theta_{k-1}+\theta_k}{2}$, but to the left side of $\frac{\theta_{k+1}+\theta_k}{2}$, implying that

$$\frac{\theta_{k+1} + \theta_k}{2} > \theta_k + \delta > \frac{\theta_{k-1} + \theta_k}{2}.$$

This implies that the LHS of this inequality is negative while that of the RHS is positive, helping us rewrite it as follows

$$-\left[(-1)\left(\frac{\theta_{k-1} + \theta_k}{2} - (\theta_k + \delta)\right)\right]^2 \geq -\left(\frac{\theta_{k+1} + \theta_k}{2} - (\theta_k + \delta)\right)^2$$

Separating PBE

Multiplying both sides by -1 and taking square roots on both sides, we obtain that

$$\frac{\theta_{k-1} + \theta_k}{2} - (\theta_k + \delta) \geq (\theta_k + \delta) - \frac{\theta_{k+1} + \theta_k}{2}$$

Rearranging the above expression, we find that

$$\theta_{k-1} + \theta_k - 2(\theta_k + \delta) \geq 2(\theta_k + \delta) - \theta_{k+1} - \theta_k$$

And further simplifying, yields

$$\theta_{k+1} \geq 2\theta_k - \theta_{k-1} + 4\delta$$

Separating PBE

- *No incentives to underreport* – checking the $(k + 1)$ – *th* sender, occurs if

$$-\left(\frac{\theta_{k+1} + \theta_k}{2} - (\theta_{k+1} + \delta)\right)^2 \geq -\left(\frac{\theta_k + \theta_{k-1}}{2} - (\theta_{k+1} + \delta)\right)^2$$

- Following the same approach in the case of “no incentives to overreport”, we now have that $\frac{\theta_{k+1} + \theta_k}{2} < \frac{\theta_k + \theta_{k-1}}{2}$, implying that the LHS of the above inequality is negative, whereas the RHS is positive. Rearranging the above expression, we obtain:


$$\theta_{k+1} + \delta - \frac{\theta_{k+1} + \theta_k}{2} \geq \frac{\theta_k + \theta_{k-1}}{2} - (\theta_{k+1} + \delta)$$

$$\Rightarrow \theta_{k+1} \geq \frac{1}{3}(2\theta_k - \theta_{k-1} - 4\delta)$$

- Since, $\theta_{k+1} > \theta_k$ by construction, the k^{th} sender has no incentives to underreport.
 - Intuitively, the lobbyist finds it unprofitable to report a lower type to the politician, for all values of the bias parameter δ .
- Therefore, in general, the condition for the k^{th} sender to send the appropriate message is

$$\theta_{k+1} \geq 2\theta_k - \theta_{k-1} + 4\delta$$

Separating PBE



5. *Summary.* The sender has no incentives to underreport, and does not have incentives to over-report iff $\theta_{k+1} \geq 2\theta_k - \theta_{k-1} + 4\delta$.

Separating PBE

- We have shown the conditions under which θ is truthfully reported by the sender.
- There are a few things we have not characterized.
 - i. The number of partitions that can be sustained in equilibrium, N
 - ii. How is this number of partitions affected by the preference divergence parameter, δ ; and
 - iii. The length of each of these partitions (intervals), as they are not necessarily equally long.

Equilibrium number of partitions

- $\theta \sim U[0,1]$, first interval starts at 0, last interval finishes at 1.
- Let $d = \theta_1 - \theta_0$.
- We can now arrange the ‘no incentives to over-report’ as
$$\theta_{k+1} - \theta_k \geq (\theta_k - \theta_{k-1}) + 4\delta.$$

- If this condition binds,

$$\theta_2 - \theta_1 = \underbrace{(\theta_1 - \theta_0)}_d + 4\delta$$

- We can keep doing this exercise and get

Equilibrium number of partitions

$$\begin{aligned}\theta_k - \theta_{k-1} &= (\theta_{k-1} - \theta_{k-2}) + 4\delta \\ &= (\theta_{k-2} - \theta_{k-3}) + (2 + 4\delta) \\ &= (\theta_{k-3} - \theta_{k-4}) + (3 + 4\delta) \\ &= \dots \\ &= (\theta_1 - \theta_0) + [(k-1) \times 4\delta] \\ &= d + 4(k-1)\delta\end{aligned}$$

- When $k = N$, $\theta_N - \theta_{N-1} = d + 4(N-1)\delta$
- Since $\theta_N - \theta_0 = 1$

Equilibrium number of partitions

- We can now express the length of the unit interval, $\theta_N - \theta_0 = 1$, as the sum of N partitions as follows:

$$\theta_N - \theta_0 = \underbrace{(\theta_N - \theta_{N-1})}_{d+4(N-1)\delta} + \underbrace{(\theta_{N-1} - \theta_{N-2})}_{d+4(N-2)\delta} + \cdots + \underbrace{(\theta_2 - \theta_1)}_{d+4\delta} + \underbrace{(\theta_1 - \theta_0)}_d$$

$$\underbrace{\theta_N - \theta_0}_{=1} = Nd + 4\delta \underbrace{[(N-1) + (N-2) + \cdots + 1]}_{=\frac{N(N-1)}{2} \text{ since } 1+2+\cdots+(N-1)=\frac{N(N-1)}{2}}$$

Therefore, the above expression simplifies to:

$$\theta_N - \theta_0 = Nd + 4\delta \frac{N(N-1)}{2} = Nd + 2\delta N(N-1)$$

Equilibrium number of partitions

- And since the LHS is $\theta_N - \theta_0 = 1$ (θ lies in the unit interval), we can write the above equation as

$$1 = Nd + 2\delta N(N - 1),$$

- and solve for the bias parameter, δ , to obtain:

$$\underline{\delta}(d) = \frac{1 - Nd}{2N(N - 1)} \Rightarrow \underline{\delta}(0) = \frac{1}{2N(N - 1)}$$

- Cutoff $\underline{\delta}(d)$ decreases in d .
- As d increases, a given number of partitions N is harder to sustain as a PBE.
- Given $\underline{\delta}(0)$, more partitions (higher N) can only be supported as a PBE if the bias parameter δ becomes smaller.
 - That is, more informative PBEs can be sustained when the preferences of lobbyist and politician are more similar (lower δ).

Equilibrium number of partitions

- Solving for N ,

$$N^2 - N - \frac{1}{2\delta} \leq 0$$

- Factorizing yields

$$\left(N - \frac{1 + \sqrt{1 + \frac{2}{\delta}}}{2}\right) \left(N - \frac{1 - \sqrt{1 + \frac{2}{\delta}}}{2}\right) \leq 0$$

- Furthermore, since N is a positive integer, we can rule out the negative root and restrict the result to

$$N \leq \bar{N}(\delta) \equiv \frac{1 - \sqrt{1 + \frac{2}{\delta}}}{2}$$


- Where $\lfloor \cdot \rfloor$ rounds to the next integer from below, e.g., $\lfloor 3.7 \rfloor = 3$.

Example 13.1. Equilibrium number of partitions

- The figure to the right depicts $\bar{N}(\delta)$.
- PBE yields a smaller number of partitions as the bias parameter δ increases. Mathematically,

$$\delta \leq \frac{1}{2N(N-1)}$$

- This also indicates that as we seek a larger number of partitions (higher N) in equilibrium, the preference divergence parameter must be lower (see figure).



Example 13.1. Equilibrium number of partitions.

- *Example 13.1.*

- If we seek to support $N = 2$ partitions, we need

$$\delta \leq \frac{1}{2 \times 2(2 - 1)} = \frac{1}{4}$$

- To sustain $N = 3$ partitions, we need

$$\delta \leq \frac{1}{2 \times 3(3 - 1)} = \frac{1}{12}$$

- The latter imposes a more restrictive condition on players' preference alignment.

Interval lengths in equilibrium

- We now try to find equilibrium length d^* .
- Consider again:

$$\theta_{k+1} - 2\theta_k + \theta_{k-1} = 4\delta.$$

- Since the unit interval starts at $\theta_0 = 0$, for $k = 1$ we can write

$$\theta_2 = 2\theta_1 - \theta_0 + 4\delta = 2\theta_1 + 4\delta$$

- Similarly,

$$\theta_3 = 2\theta_2 - \theta_1 + 4\delta = \underbrace{2(2\theta_1 + 4\delta)}_{\theta_2} - \theta_1 + 4\delta = 3\theta_1 + 12\delta$$

- More generally, for any value value of k , $\theta_k = k\theta_1 + 2k(k-1)\delta$.

Interval lengths in equilibrium

- Evaluating at $k = N$, we obtain $\theta_N = Nd + 2N(N - 1)\delta$.
- We can express $\theta_N - \theta_0 = 1 = Nd + 2N(N - 1)\delta$ since $\theta_0 = 0$ and $\theta_N - \theta_0 = 1$.
- Solving for d yields.

$$d^* = \frac{1}{N} - 2(N - 1)\delta.$$

- This satisfies $d^* = \theta_1 - \theta_0$ since $\theta_0 = 0$.
- This length decreases in the number of partitions in that equilibrium, N , since $\frac{\partial d^*}{\partial N} = -2\delta - \frac{1}{N^2} \leq 0$ which means that the first interval shrinks to “make room” for subsequent partitions to its right side.

Interval lengths in equilibrium

- Example 13.2. First interval decreasing in N .
 - If $\delta = \frac{1}{20}$ and $N = 2$, we obtain $d^* = \frac{1}{2} - 2(2 - 1)\frac{1}{20} = \frac{2}{5}$.
 - When N increases to 3, $d^* = \frac{2}{15}$.
- We now use these results to find the length of the k -th interval, $\theta_k - \theta_{k-1}$.

$$\theta_k = \underbrace{k \left(\frac{1}{N} - 2(N - 1)\delta \right)}_{\theta_1} + 2k(k - 1)\delta = \frac{k}{N} - 2k(N - k)\delta$$

which implies that the length of the k^{th} interval is

$$\begin{aligned} \theta_k - \theta_{k-1} &= \underbrace{\left(\frac{k}{N} - 2k(N - k)\delta \right)}_{\theta_k} - \underbrace{\left(\frac{k-1}{N} - 2(k-1)(N - k + 1)\delta \right)}_{\theta_{k-1}} \\ &= \frac{1}{N} - 2(N + 1 - 2k)\delta \end{aligned}$$

Interval lengths in equilibrium

- We can confirm that the length of the first interval coincides with the expression found above, d^* .

$$\begin{aligned}\theta_1 - \theta_0 &= \frac{1}{N} - 2(N + 1 - 2)\delta \\ &= \frac{1}{N} - 2(N - 1)\delta = d^*\end{aligned}$$

- This coincides with the result of d^* found above.

Example 13.3 Length of each interval in equilibrium

- Following Example 13.2, when $\delta = \frac{1}{20}$ and $N = 2$, we find that $d^* = \frac{2}{5}$.

- The length of the second interval is

$$\theta_2 - \theta_1 = \frac{1}{2} - 2[2 + 1 - (2 \times 2)] \frac{1}{20} = \frac{3}{5}$$

- We get that $\frac{2}{5} + \frac{3}{5} = 1$, as required.

- When $N = 3$:

- The first interval's length is $d^* = \frac{2}{5}$

- That of the second interval is: $\theta_2 - \theta_1 = \frac{1}{3} - 2[3 + 1 - (2 \times 2)] \frac{1}{20} = \frac{1}{3}$;

- That of the third interval is: $\theta_3 - \theta_2 = \frac{1}{3} - 2[3 + 1 - (2 \times 3)] \frac{1}{20} = \frac{8}{15}$

- So when $N = 3$, the interval lengths are $\frac{2}{5}$, $\frac{1}{3}$ and $\frac{8}{15}$ respectively, which add up to 1, spanning the unit interval.

Extensions - I

- Allowing for conversations, Krishna and Morgan (2004).
 - How the above results would be affected if the politician could also send messages to the lobbyist.
 - Authors identify conditions under which this conversation between the players yields more partitions than otherwise.

Extensions - II

- Cheap talk vs. Delegation, Dessein (2002).
 - Delegating the decision of p to the lobbyist entails a trade-off:
 - On one hand, it shifts the policy decision to the informed player (the lobbyist), which should be positive;
 - But, on the other hand, it allows the lobbyist to choose a policy according to her bias, δ , which reduces the politician's utility.
 - When δ is relatively small, the first (positive) effect dominates the second (negative) effect, and the politician is better off delegating.
 - Otherwise, the bias is too large, and the politician is better off not delegating (operating as in the Crawford and Sobel's model we studied above).

Extensions - III

- Open vs. Closed rules.
 - Open rule: Politician can freely amend the bill.
 - Closed rule: Politician faces limited ability to amend the bill.
 - Despite “tying the hands” of the politician, the closed rule can yield more information transmission in equilibrium than the open rule (Gilligan and Krehbiel, 1987).

Extensions - IV

- Several lobbyists.
 - Is information transmission facilitated when the politician receives messages from more than one lobbyist (sender).
 - When lobbyists send their messages simultaneously, they convey the true state to the politician (Krishna and Morgan, 2001).
 - When lobbyists send their messages sequentially, if experts have similar (opposite) biases, the politician is better off ignoring the lobbyist with the smaller bias (considering the messages sent by both lobbyists, respectively).

Other Extensions

- Multiple receivers (politicians) (Farrell and Gibbons, 1989).
- Repeated cheap talk (Ottaviani and Sorensen, 2006).
- Allowing for the expert's private information to be multidimensional (Chakraborty and Harbaugh, 2007).
- The presence of a mediator (Ganguly and Ray, 2005).
- Noise in the lobbyist messages before the politician receives them (Blume et al., 2007).
- Equilibrium refinement criteria in cheap talk environments (Farrell, 1993; Chen et al., 2008; Groot Ruiz et al., 2015).