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
$$\Theta_i = \{\theta_i^1, \theta_i^2, \ldots, \theta_i^K\}$$
- 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, $(p_i^1, p_i^2, \ldots, p_i^K)$ satisfy $p_i^k \in [0,1]$ for every $k = \{1,2, \ldots, K\}$ and $\sum_{k=1}^{K} p_i^k = 1$, we can omit the probability of the last type, writing it, instead, as
$$p_i^K = 1 - \sum_{k \neq K}^{K} p_i^k.$$
- 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).

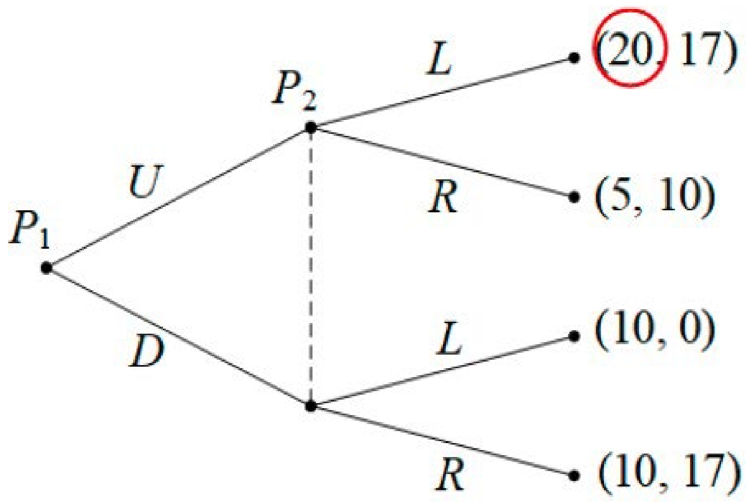

Figure 8.1a. Simultaneous move game where $x = 20$

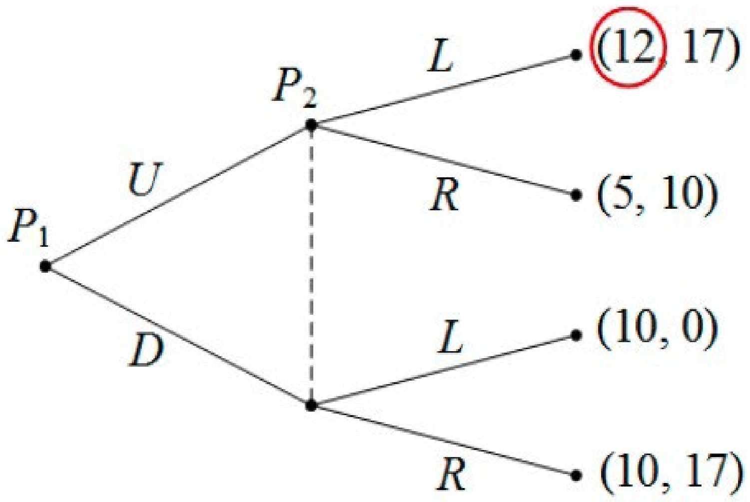

Figure 8.1b. Simultaneous move game where $x = 12$

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).

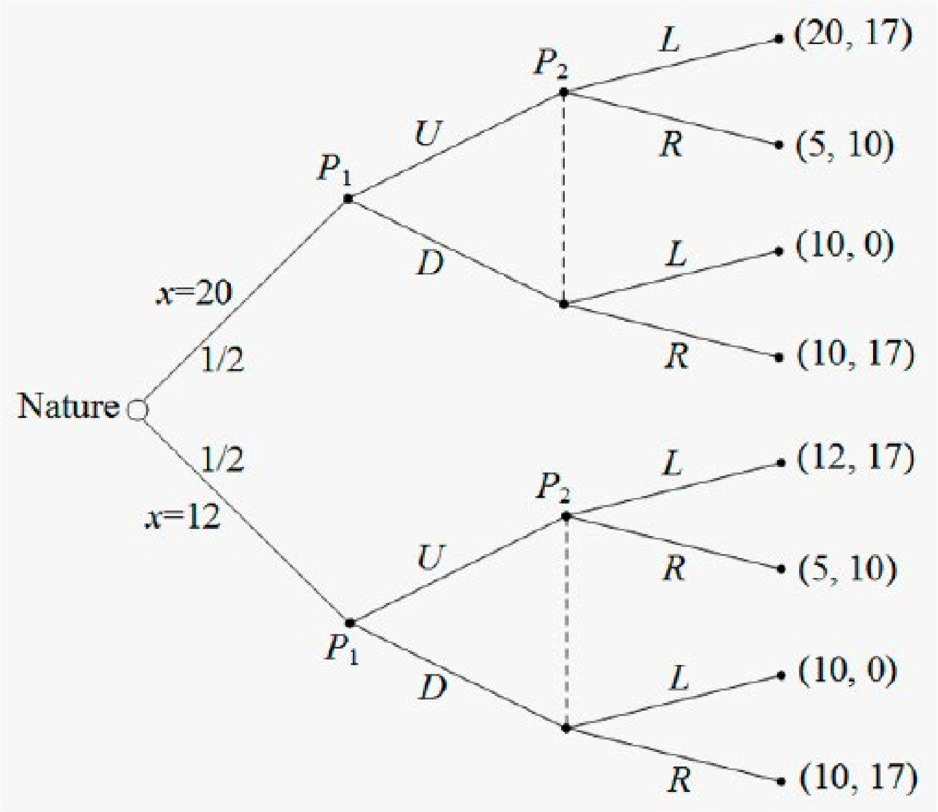

Figure 8.2. Combining both games of figure 8.1 to allow for incomplete information.

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

$$EU_i\big(s_i(\theta_i), s_{-i}(\theta_{-i})\big) \geq EU_i\big(s_i'(\theta_i), s_{-i}(\theta_{-i})\big)$$

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 many 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 $\left(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})\right)$ 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	Player 2: 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

$$EU_1(U^{20}U^{12}, L) = \underbrace{\frac{1}{2}20}_{if\ x=20} + \underbrace{\frac{1}{2}12}_{if\ x=12} = 16$$

$$EU_2(U^{20}U^{12}, L) = \underbrace{\frac{1}{2}17}_{if\ x=20} + \underbrace{\frac{1}{2}17}_{if\ x=12} = 17$$

$$EU_1(U^{20}D^{12}, L) = \frac{1}{2}20 + \frac{1}{2}10 = 15$$

$$EU_2(U^{20}U^{12}, L) = \frac{1}{2}17 + \frac{1}{2}0 = 8.5$$

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

Player 1		L	R
	$U^{20}U^{12}$	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

Player 1		L	R
	$U^{20}U^{12}$	$\underline{16},\underline{17}$	5,10
	$U^{20}D^{12}$	15,8.5	7.5,$\underline{13.5}$
	$D^{20}U^{12}$	11,8.5	7.5,$\underline{13.5}$
	$D^{20}D^{12}$	10,0	$\underline{10},\underline{17}$

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, $s_i^*(\theta_i) = \left(s_i^*(\theta_i^1), \ldots, s_i^*(\theta_i^k)\right)$.
   - Therefore, $\left(s_i^*(\theta_i^1), s_i^*(\theta_i^2)\right)$ 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.*

$$\max_{q_2^H \geq 0} \quad \pi_2^H(q_2^H) = (1 - q_2^H - q_1)q_2^H - \frac{1}{2}q_2^H$$

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.*

$$\max_{q_2^L \geq 0} \quad \pi_2^L(q_2^L) = (1 - q_2^L - q_1)q_2^L$$

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.

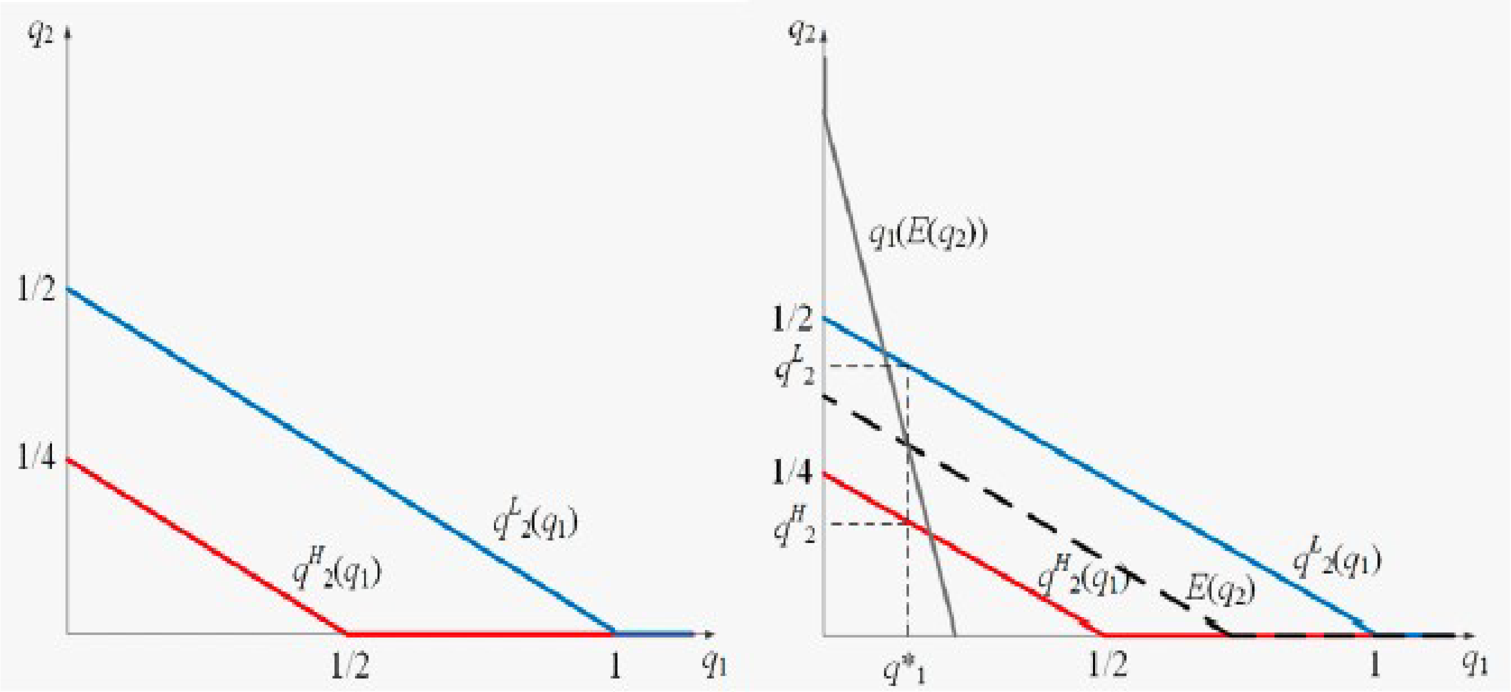

Figure 8.3a. Firm 2 - High or Low costs.

Figure 8.3b. Equilibrium output under incomplete information.

Example 8.2. Cournot competition with asymmetrically informed firms

- *Uninformed firm.*

$$\max_{q_1 \geq 0} \quad \pi_1(q_1) = p \underbrace{[(1 - q_2^H - q_1)q_1]}_{Firm\ 2's\ costs\ are\ high} + (1-p) \underbrace{[(1 - q_2^L - q_1)q_1]}_{Firm\ 2's\ costs\ are\ low} - \frac{1}{2} q_1$$

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!

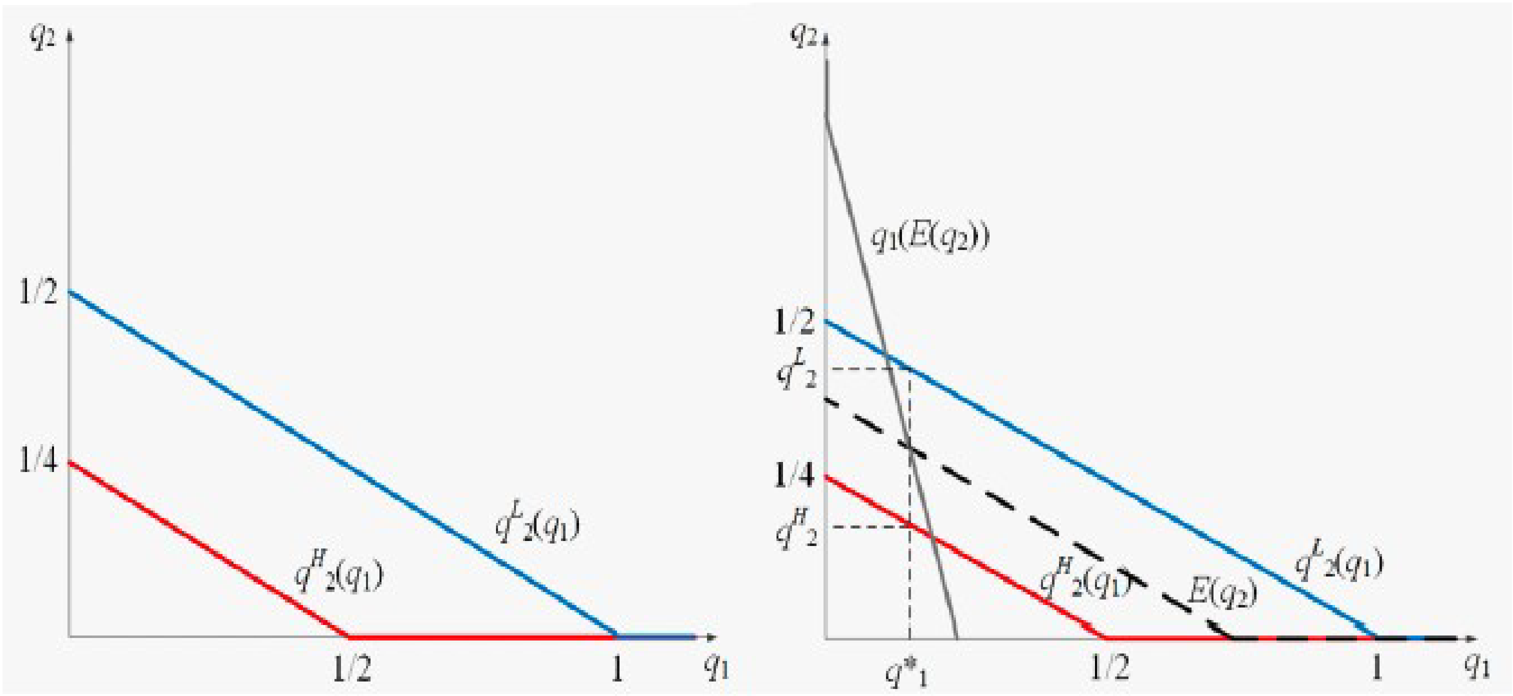

Figure 8.3a. Firm 2 - High or Low costs.

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

$$q_2^H(q_1^*) = \frac{3-p}{12}; \;\; q_2^L(q_1^*) = \frac{6-p}{12}$$

Example 8.2. Cournot competition with asymmetrically informed firms

- Therefore, the BNE of this game is given by the triplet of output levels

$$\left(q_1^*, q_2^{H^*}, q_2^{L^*}\right) = \left(\frac{p}{6}, \frac{3-p}{12}, \frac{6-p}{12}\right)$$

- 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.

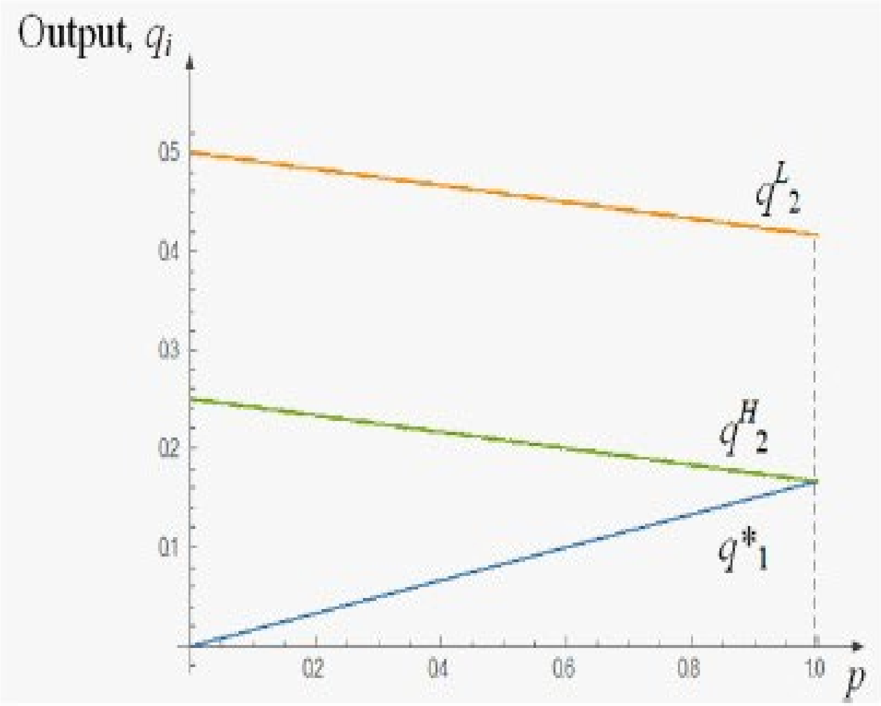

Figure 8.4. Cournot competition with asymmetrically informed firms – Equilibrium output

Example 8.2. Cournot competition with asymmetrically informed firms

- The figure also embodies complete information as special cases.
- When $p \to 0$, firm 2 has (almost certainly) low costs, thus being very competitive.
    - The uninformed firm 1 remains inactive, while firm 2 produces ¼ when its cost is high and ½ when its cost is low.
- When $p \to 1$, firm 2 has (almost certainly) high costs, thus being very competitive.
    - The uninformed firm 1 produces 1/6, as much as firm 2's output.

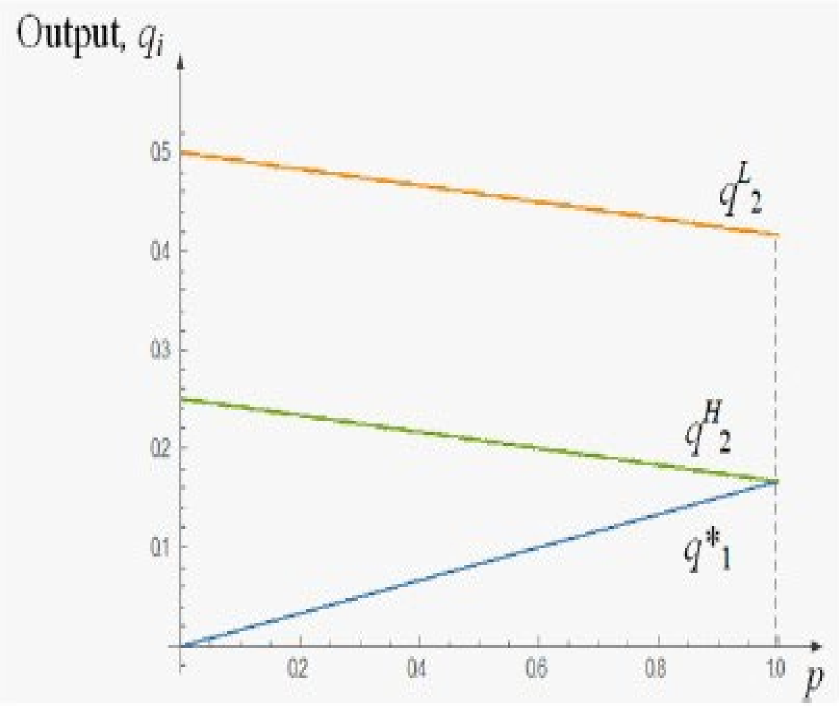

Figure 8.4. Cournot competition with asymmetrically informed firms – Equilibrium output

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:

$$\max_{q_{i \geq 0}} \quad \pi_i^H(q_i) = p \underbrace{\left[\left(1 - q_j^H - q_i\right)q_i\right]}_{Firm\ j's\ costs\ are\ high} + (1-p) \underbrace{\left[\left(1 - q_j^L - q_i\right)q_i\right]}_{Firm\ j's\ costs\ are\ low} - \frac{1}{2} q_i$$

Differentiating with respect to q_i and finding BR function:

$$q_i^H\left(q_j^H, q_j^L\right) = \frac{1}{4} - \frac{pq_j^H + (1-p)q_j^L}{2}$$

What if both players are privately informed?

Low costs. If firm i has, instead, low marginal costs, $c_L = 0$, it solves:

$$\max_{q_{i \geq 0}} \quad \pi_i^L(q_i) = p \underbrace{\left[\left(1 - q_j^H - q_i\right)q_i\right]}_{Firm\ j's\ costs\ are\ high} + (1-p) \underbrace{\left[\left(1 - q_j^L - q_i\right)q_i\right]}_{Firm\ j's\ costs\ are\ low}$$

Differentiating with respect to q_i, yields

$$p\left(1 - q_j^H - 2q_i\right) + (1-p)\left(1 - q_j^L - 2q_i\right) = 0$$

which leads to firm i's best response function:

$$q_i^L\left(q_j^H, q_j^L\right) = \frac{1}{2} - \frac{pq_j^H + (1-p)q_j^L}{2}$$

What if both players are privately informed?

Comparing best response functions:

$$q_i^L\left(q_j^H, q_j^L\right) = \frac{1}{2} - \frac{pq_j^H + (1-p)q_j^L}{2}$$

Relative to $q_i^H\left(q_j^H, q_j^L\right)$, the best response function $q_i^L\left(q_j^H, q_j^L\right)$:

- 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, $\left(q_1^*, q_2^{H^*}, q_2^{L^*}\right)$, 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).