Chapter 3. Games of Incomplete Information

Xiaoxiao Hu

Games of Incomplete Information

- In this chapter, we will study games of incomplete information, also called Bayesian games.
- In such games, at least one player is uncertain about another player's payoff function.

3.A. Static Games of Incomplete Information

- In this section, we will study simultaneous-move game of incomplete information, also called static Bayesian game.
- In Section 3.A.1, we will study incomplete information Cournot duopoly model.
- In Section 3.A.2, we will develop normal-form representation of general static Bayesian game and corresponding solution concept Bayesian Nash Equilibrium.
- In Sections 3.A.3 to 3.A.6, we will study applications.

3.A.1. Cournot Competition under Asymmetric Information

- Quantities (of a homogeneous product) produced by firms 1 and 2: q_1 and q_2
- Market-clearing price when aggregate quantity is $Q = q_1 + q_2$: $P(Q) = a - Q$.
- Firms choose quantities simultaneously.

Asymmetric Cournot

- Firm 1's cost function is $C_1(q_1) = cq_1$.
- Firm 2's cost function is
  - $C_2(q_2) = c_H q_2$ with probability θ, and
  - $C_2(q_2) = c_L q_2$ with probability $1 - \theta$,

  where $c_L < c_H$.

Asymmetric Cournot

Information is asymmetric:

- Firm 2 knows
  - its own cost function (realization of c_H, c_L) and
  - Firm 1's cost function
- Firm 1 knows
  - its own cost function and
  - only that Firm 2's marginal cost is c_H with probability θ and c_L with probability $1 - \theta$.

Asymmetric Cournot: Analysis

- Naturally, Firm 2 may choose different quantities depending on whether its marginal cost is high or low.
- Moreover, Firm 1 should anticipate this.
- Let
  - $q_2^*(c_H)$ and $q_2^*(c_L)$ denote Firm 2's equilibrium quantity choice;
  - q_1^* denote Firm 1's equilibrium quantity choice.

Asymmetric Cournot: Analysis

Then,

- $q_2^*(c_H)$ solves $\max_{q_2} \Big[(a - q_1^* - q_2) - c_H\Big] q_2$.
- $q_2^*(c_L)$ solves $\max_{q_2} \Big[(a - q_1^* - q_2) - c_L\Big] q_2$.
- q_1^* solves

$$\max_{q_1} \theta \Big[(a - q_1 - q_2^*(c_H)) - c\Big] q_1 + (1-\theta)\Big[(a - q_1 - q_2^*(c_L)) - c\Big] q_1$$

$$\Longrightarrow \max_{q_1} \Big[a - q_1 - \Big(\theta q_2^*(c_H) + (1-\theta) q_2^*(c_L)\Big) - c\Big] q_1$$

Asymmetric Cournot: Analysis

The solution is (assuming that solutions are all positive)

$$q_2^*(c_H) = \frac{a - 2c_H + c}{3} + \frac{1-\theta}{6}(c_H - c_L);$$
$$q_2^*(c_L) = \frac{a - 2c_L + c}{3} - \frac{\theta}{6}(c_H - c_L);$$
$$q_1^* = \frac{a - 2c + \theta c_H + (1-\theta)c_L}{3}.$$

Remark. $q_2^*(c_H) > \frac{a-2c_H+c}{3}$ and $q_2^*(c_L) < \frac{a-2c_L+c}{3}$: Firm 2 not only tailors its quantity to its cost but also responds to the fact that Firm 1 cannot do so.

3.A.2. Static Bayesian Games and Bayesian Nash Equilibrium

To characterize static Bayesian games, we need to capture the idea that

1. each player knows his/her own payoff function;
2. each player may be uncertain about the other players' payoff functions.

Harsanyi (1967) introduced type spaces to model players' information on payoff-relevant parameters.

Type and Belief.

Player i's payoff functions is represented by

$$u_i(a_1, ..., a_n; t_i),$$

where t_i is called Player i's type.

- $t_i \in T_i$
- T_i is the set of possible types, or type space.

Type and Belief.

For the Cournot competition model in Section 3.A.1,

- Firm 2 has two types and its type space is $T_2 = \{c_L, c_H\}$;
- Firm 1 has only one type and its type space is $T_1 = \{c\}$.

Type and Belief.

Given this definition of types,

1. "Player i knows his/her own payoff function" is equivalent to "Player i knows his/her own type".
2. "Player i may be uncertain about the other players' payoff functions" is equivalent to "Player i maybe uncertain about the types of other players, $t_{-i} = \{t_1, ..., t_{i-1}, t_{i+1}, ..., t_n\}$".

Type and Belief.

- We use T_{-i} to denote the set of possible types of other players.
- We use $p_i(t_{-i}|t_i)$ to denote Player i's belief about t_{-i} when his/her own type is t_i.
- Belief is computed by Bayes' rule from prior probability distribution $p(t)$:

$$p_i(t_{-i}|t_i) = \frac{p(t_{-i}, t_i)}{p(t_i)} = \frac{p(t_{-i}, t_i)}{\sum_{t_{-i} \in T_{-i}} p(t_{-i}, t_i)}.$$

Bayes' Rule

Example 3.A.1. Consider a two player game with the following prior distribution of types:

		Player 2	
		Type C	Type D
Player 1	Type A	30%	40%
	Type B	10%	20%

Question 3.1. What is the posterior probability $p_1(C|A)$?

Bayes' Rule

Example 3.A.2. A certain disease affects about 1 out of 10,000 people. There is a screening test to check whether a person has the disease. The test is quite accurate.

- When a person has the disease, it gives a positive result 99% of the time.
- When a person does not have the disease, it gives a negative result 98% of the time.

Question 3.2. A random person gets tested for the disease and the result comes back positive. What is the probability that the person has the disease?

Normal-form Representation

The normal-form representation of n-player static Bayesian game specifies

- players' action spaces $A_1, ..., A_n$;
- their type spaces $T_1, ..., T_n$;
- their beliefs $p_1(t_{-1}|t_1), ..., p_n(t_{-n}|t_n)$;
- their payoff functions $u_i(a_1, ..., a_n; t_i)$ for all i.[1]

[1] More generally, a player's payoff function could also depend on the other players' types. In this case, we write $u_i(a_1, ..., a_n; t_1, ..., t_n)$.

Normal-form Representation

Following Harsanyi (1967), timing of a static Bayesian game is as follows:

1. nature draws a type vector $t = (t_1, ..., t_n)$ where t_i is drawn from set of possible types T_i;
2. nature reveals t_i to Player i but not to any other player;
3. players simultaneously choose actions, Player i choosing $a_i \in A_i$; and then
4. payoffs $u_i(a_1, ..., a_n; t_i)$ are received.

Remark 3.1. Note that by introducing the fictional moves by nature, incomplete information game is transformed to imperfect information game.

- Here, Player i does not know complete history of game when actions are chosen in Step 3.
- In particular, Player i does not know what nature has revealed to other players.

Definition 3.A.1 (Strategy)**.** In static Bayesian game, a strategy for Player i is a function $s_i(t_i)$ that specifies action $a_i \in A_i$ when type $t_i \in T_i$ is drawn by nature.

Bayesian Nash Equilibrium

For Cournot competition model in Section 3.A.1,

- Firm 2's strategy is $(q_2^*(c_H), q_2^*(c_L))$;
- Firm 1's strategy is q_1^*.

Bayesian Nash Equilibrium

- Next, we define the solution concept in a static Bayesian game, called Bayesian Nash Equilibrium.
- The central idea is the same: each player's strategy must be a best response to the other players' strategies.

Bayesian Nash Equilibrium

Definition 3.A.2 (Bayesian Nash Equilibrium)**.** In the static Bayesian game, the strategies $s^* = (s_1^*, ..., s_n^*)$ are a (pure strategy) Bayesian Nash Equilibrium (BNE) if for each player i and for each of i's type $t_i \in T_i$, $s_i^*(t_i)$ solves

$$\max_{a_i \in A_i} \sum_{t_{-i} \in T_{-i}} u_i(s_1^*(t_1), ...s_{i-1}^*(t_{i-1}), a_i, s_{i+1}^*(t_{i+1}), ..., s_n^*(t_n); t) p_i(t_{-i}|t_i).$$

3.A.3. Mixed Strategies Revisited

A mixed-strategy Nash equilibrium in a game of complete information can (almost always) be interpreted as a pure-strategy Bayesian Nash equilibrium in a closely related game with a little bit of incomplete information.

Mixed Strategies Revisited

Example 3.A.3. Consider the following battle of the sexes game:

		Bob: Opera	Bob: Movie
Alice	Opera	$(2 + t_a, 1)$	$(0, 0)$
	Movie	$(0, 0)$	$(1, 2 + t_b)$

Figure 3.1: The Battle of the Sexes

where t_a is privately known by Alice; t_b is privately known by Bob; and t_a and t_b are independently drawn from a uniform distribution on $[0, x]$.

Incomplete Information Battle of Sexes Game

- Action spaces: $A_a = A_b = \{\text{Opera, Movie}\}$;
- Type spaces: $T_a = T_b = [0, x]$;
- Beliefs: $p_a(t_b|t_a) = 1/x$ for all $t_a \in [0, x]$ and $t_b \in [0, x]$, $p_b(t_a|t_b) = 1/x$ for all $t_a \in [0, x]$ and $t_b \in [0, x]$;
- Payoff functions: see the payoff matrix Figure 3.1.

Incomplete Information Battle of Sexes Game

We construct a pure-strategy Bayesian Nash equilibrium in which

- Alice plays "Opera" if t_a exceeds a critical value a, and plays "Movie" otherwise.
- Bob plays "Movie" if t_b exceeds a critical value b, and plays "Opera" otherwise.

Incomplete Information Battle of Sexes Game

Given Bob's strategy, Alice's expected payoffs from playing "Opera" and "Movie" are respectively

$$\mathbb{E}u_a(\text{Opera}, (\frac{b}{x}, \frac{x-b}{x})) = \frac{b}{x}(2+t_a);$$
$$\text{and } \mathbb{E}u_a(\text{Movie}, (\frac{b}{x}, \frac{x-b}{x})) = \frac{x-b}{x}.$$

Thus, playing "Opera" is optimal if and only if

$$\frac{b}{x}(2+t_a) \geq \frac{x-b}{x} \implies t_a \geq \frac{x}{b} - 3 = a. \qquad (3.A.1)$$

Incomplete Information Battle of Sexes Game

Similarly, given Alice's strategy, Bob's expected payoffs from playing "Opera" and "Movie" are respectively

$$\mathbb{E}u_b(\text{Opera}, (\frac{x-a}{x}, \frac{a}{x})) = \frac{x-a}{x};$$
$$\text{and } \mathbb{E}u_b(\text{Movie}, (\frac{x-a}{x}, \frac{a}{x})) = \frac{a}{x}(2+t_b).$$

Thus, playing "Movie" is optimal if and only if

$$\frac{a}{x}(2+t_b) \geq \frac{x-a}{x} \implies t_b \geq \frac{x}{a} - 3 = b. \tag{3.A.2}$$

Incomplete Information Battle of Sexes Game

Equations (3.A.1) and (3.A.2) imply

$$a = b = \frac{-3 + \sqrt{9 + 4x}}{2}.$$

The probability that Alice plays "Opera" (i.e., $\frac{x-a}{x}$) and Bob plays "Movie" (i.e., $\frac{x-b}{x}$) both equal

$$1 - \frac{-3 + \sqrt{9 + 4x}}{2x} = 1 - \frac{2}{3 + \sqrt{9 + 4x}},$$

which approaches 2/3 as x approaches 0.

Incomplete Information Battle of Sexes Game

- Recall that mixed strategy NE of complete information battle of the sexes game is $\left((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3})\right)$.
- As incomplete information disappears, players' behavior in this pure-strategy BNE of incomplete-information game approaches their behavior in mixed-strategy NE in the original game of complete information.

3.A.4. First-Price Sealed-Bid Auction

- We have learned second-price auction in Chapter 1.
- Recall that second-price auction is dominant strategy solvable: bidding one's own valuation is a weakly dominant strategy.
- Now, we will study first-price auction.
- Note that first-price and second-price auctions only differ in winner's payments.
- We will consider a simple version of first-price auction with only two bidders.

First-Price Sealed-Bid Auction

- There is one indivisible good for sale.
- Valuations of two potential buyers are independently drawn from a uniform distribution with support $[0, 1]$.
- Denote Buyer i's valuation by v_i.

First-Price Sealed-Bid Auction

Auction rule is as follows:

- Buyers bid simultaneously and each submits a bid $b_i \in [0, +\infty)$.
- Bidder with the highest bid wins the auction and pays his/her own bid.
- If the two buyers submit the same highest bid, then each of the buyers has 1/2 chance of winning the good. The payment is the highest bid (since there is a tie).

Normal-form Representation

- Buyer i's action: submit a bid $b_i \in A_i \in [0, \infty)$.
- Buyer i's type: her valuation $v_i \in T_i = [0, 1]$.
- Buyer i's belief about Buyer j's type: v_j is uniformly distributed on $[0, 1]$, given any v_i
  (valuations are independent)
- Buyer i's payoff when submitting bid b_i is

$$u_i = \begin{cases} 0 & \text{if } b_i < b_j \\ \frac{v_i - b_i}{2} & \text{if } b_i = b_j \\ v_i - b_i & \text{if } b_i > b_j \end{cases}$$

Bayesian Nash Equilibrium

- A strategy for Buyer i is a function $b_i(v_i)$.
- In a Bayesian Nash Equilibrium, Buyer 1's strategy $b_1(v_1)$ is a best response to Buyer 2's strategy $b_2(v_2)$, and vice versa.
- Thus, $b_i(v_i)$ solves

$$\max_{b_i}(v_i - b_i)\text{Prob}\{b_i > b_j(v_j)\} + \frac{1}{2}(v_i - b_i)\text{Prob}\{b_i = b_j(v_j)\}.$$

Bayesian Nash Equilibrium

We focus on symmetric Bayesian Nash equilibrium where the two players adopt the same strictly increasing, continuous and differentiable bidding strategy $b(\cdot)$.

Bayesian Nash Equilibrium

Suppose Buyer j adopts $b(\cdot)$. Then

- $\text{Prob}\{b_i = b(v_j)\} = 0$ since v_j is uniformly distributed and $b(\cdot)$ is strictly increasing.
- $\text{Prob}\{b_i > b(v_j)\} = \text{Prob}\{b^{-1}(b_i) > v_j\} = b^{-1}(b_i)$ since v_j is uniformly distributed on $[0, 1]$ and $b(\cdot)$ is strictly increasing, continuous and differentiable.

So Buyer i solves

$$\max_{b_i} (v_i - b_i) b^{-1}(b_i).$$

Bayesian Nash Equilibrium

- Use FOC.
- Recognizing solution is $b_i = b(v_i)$ (symmetric Bayesian Nash equilibrium)
- Solve and get $b(v_i)v_i = \frac{1}{2}v_i^2 + c$.
- Use boundary condition $b(0) = 0 \implies c = 0$.
- Thus, equilibrium bidding strategy is

$$b(v_i) = \frac{1}{2}v_i.$$

3.A.5. Common Value Auction

In a common value auction, the value of good for sale is the same for all bidders.

- "Oil well" is an often cited example of common value auctions.

Jar of coins

Let us play the following auction game.

- There is a jar with some coins.
- Every bidder bids for the coins in the jar.
- Rules of the auction are as follows:
  - Do not open the jar.
  - The winner is the bidder with the highest bid.
  - The winner pays his/her own bid and gets the coins in the jar.

This is a common value auction: the amount of money in the jar is certain.

Jar of coins

Question 3.3. What is your bidding strategy?

Should you bid less or more than your estimate?

Winner's curse

In common value auctions, winning bid tends to be higher than true value of the good. Such a phenomenon is called winner's curse.

Winner's curse

Question 3.4. Why winner's curse exists?

Winner's curse

- Let v be the common value, and b_i be Bidder i's bid.
- Then Bidder i's payoff is

$$\begin{cases} v - b_i & \text{if } b_i \text{ is the highest bid;} \\ 0 & \text{otherwise.} \end{cases}$$

Winner's curse

Bidders only have estimates of value of the good.

- Let y_i be Player i's estimate: $y_i = v + \tilde{\varepsilon}_i$, where $\tilde{\varepsilon}_i$ is Bidder i's estimation error.
- y_i is also Bidder i's type.
- Suppose that on average bidders estimate correctly.
- If bidders bid roughly the same as their estimate, winner would be bidder with largest $\tilde{\varepsilon}_i$.
- Then, winning bid would be higher (actually much higher) than true value.

Winner's curse

Question 3.5. After learning winner's curse, how should you bid?

Winner's curse

1. If everyone bids roughly their own estimates, then when you (Player i) win, you know that $y_j < y_i$ for all j.
2. You only care how many coins are in the jar if you win.

So, you should bid based

- not only on your initial estimate y_i;
- but also on the fact that $y_i > y_j$ for all j.

Put differently, you should bid as if you know you win.

3.A.6. Double Auction

- one good for sell
- Buyer's valuation for the good is v_b
- Seller's is v_s
- Valuations are private information, drawn from independent uniform distributions on $[0, 1]$.
- To trade, Seller names an asking price, p_s, and Buyer simultaneously names an offer price, p_b.
  - If $p_b \geq p_s$, then trade occurs at price
  $$p = (p_b + p_s)/2;$$
  - if $p_b < p_s$, then no trade occurs.

Double Auction

Players' payoffs are as follows:

- If Buyer gets the good for price p, then Buyer's utility is $v_b - p$; if there is no trade, then Buyer's utility is 0.
- If Seller sells the good for price p, then Seller's utility is $p - v_s$; if there is no trade, then Seller's utility is 0.

Double Auction: Analysis

- Buyer's strategy is a function $p_b(v_b)$
- Seller's strategy is a function $p_s(v_s)$
- For each $v_b \in [0,1]$, $p_b(v_b)$ solves

$$\max_{p_b} \left[v_b - \frac{p_b + \mathbb{E}[p_s(v_s) | p_b \geq p_s(v_s)]}{2} \right] \text{Prob}\{p_b \geq p_s(v_s)\};$$

- For each $v_s \in [0,1]$, $p_s(v_s)$ solves

$$\max_{p_s} \left[\frac{p_s + \mathbb{E}[p_b(v_b) | p_b(v_b) \geq p_s]}{2} - v_s \right] \text{Prob}\{p_b(v_b) \geq p_s\}.$$

- There are many Bayesian Nash equilibria of this game.

One-price equilibrium

Trade occurs at a single price $x \in [0, 1]$ if it occurs at all.

- Buyer's strategy: offer x if $v_b \geq x$ and offer 0 otherwise;
- Seller's strategy: demand x if $v_s \leq x$ and demand 1 otherwise.

Question 3.6. Can you check that the above strategy profile constitutes a BNE?

One-price equilibrium

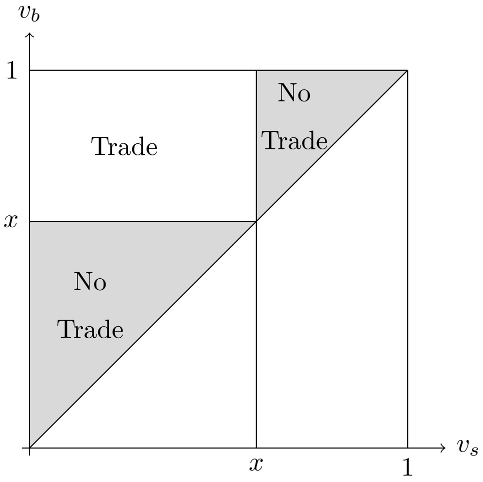

Linear equilibrium

Suppose Seller's and Buyer's strategies are

$$p_s(v_s) = a_s + c_s v_s;$$

$$\text{and } p_b(v_b) = a_b + c_b v_b.$$

- p_s is uniformly distributed on $[a_s, a_s + c_s]$;
- p_b is uniformly distributed on $[a_b, a_b + c_b]$.

Linear equilibrium

$p_b(v_b)$ and $p_s(v_s)$ solves:

$$\max_{p_b}\left[v_b - \frac{p_b + \frac{a_s+p_b}{2}}{2}\right]\left(\frac{p_b - a_s}{c_s}\right);$$
$$\text{and }\max_{p_s}\left[\frac{p_s + \frac{a_b+c_b+p_s}{2}}{2} - v_s\right]\left(\frac{a_b + c_b - p_s}{c_b}\right).$$

First-order conditions imply

$$p_b = \frac{2}{3}v_b + \frac{1}{3}a_s;$$
$$\text{and } p_s = \frac{2}{3}v_s + \frac{1}{3}(a_b + c_b).$$

Linear equilibrium

We have

$$p_s(v_s) = \frac{2}{3}v_s + \frac{1}{4};$$
$$\text{and } p_b(v_b) = \frac{2}{3}v_b + \frac{1}{12}.$$

Trade occurs when

$$p_b(v_b) \geq p_s(v_s) \implies v_b \geq v_s + \frac{1}{4}.$$

Linear equilibrium

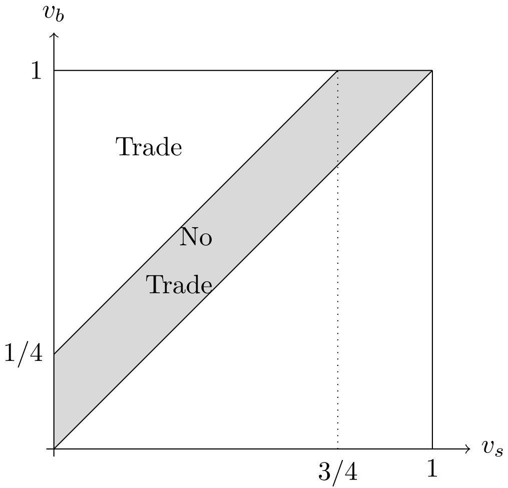

3.B. Dynamic Games of Incomplete Info.

- We will study three specific models:
  - asymmetric information Cournot model with verifiable information (Section 3.B.1),
  - job market signaling model (Section 3.B.2)
  - a screening model (Section 3.B.3)
- In Section 3.B.4, we will study the theory and formally define solution concept Perfect Bayesian Equilibrium.
- In Section 3.B.5, we will discuss refinements of Perfect Bayesian Equilibrium.

3.B.1. Asymmetric Cournot with Verifiable Information

- Quantities (of a homogeneous product) produced by firms 1 and 2: q_1 and q_2
- Market-clearing price when aggregate quantity is $Q = q_1 + q_2$: $P(Q) = a - Q$.
- Firms choose quantities simultaneously. (Cournot model)

Asymmetric Cournot with Verifiable Information

Firm 1's cost function is

$$C_1(q_1) = c_M q_1.$$

Firm 2's cost function is

$$C_2(q_2) = \begin{cases} c_H q_2 = (c_M + x) q_2 & \text{with probability } 1/3 \\ c_M q_2 & \text{with probability } 1/3 \\ c_L q_2 = (c_M - x) q_2 & \text{with probability } 1/3 \end{cases}$$

Asymmetric Cournot with Verifiable Information

Information is asymmetric:

- Firm 2 knows
  - its own cost function (realization of c_H, c_M, c_L) and
  - Firm 1's cost function
- Firm 1 knows
  - its own cost function and
  - only that Firm 2's marginal cost is c_H, c_M or c_L, each with 1/3 probability.

Asymmetric Cournot with Verifiable Information

Before the firms choose quantities, Firm 2 can costlessly and verifiably reveal its cost information to Firm 1.

Asymmetric Cournot with Verifiable Information

Question 3.7. Should Firm 2 reveal its cost information?

Asymmetric Cournot with Verifiable Information

Perhaps it is easier to first consider the following question:

Question 3.8. Would Firm 2 want Firm 1 to know if it has high, middle, or low cost?

Asymmetric Cournot with Verifiable Information

- In Cournot model, one firm's profit would be higher if the other firm produces less.
- Result is: compared to not knowing Firm 2's cost, Firm 1 produces less (more) if it knows that Firm 2 has low (high) cost.
- Firm 2 would want Firm 1 to know if it has low cost.
- Firm 2 with low cost would reveal its cost information.

Asymmetric Cournot with Verifiable Information

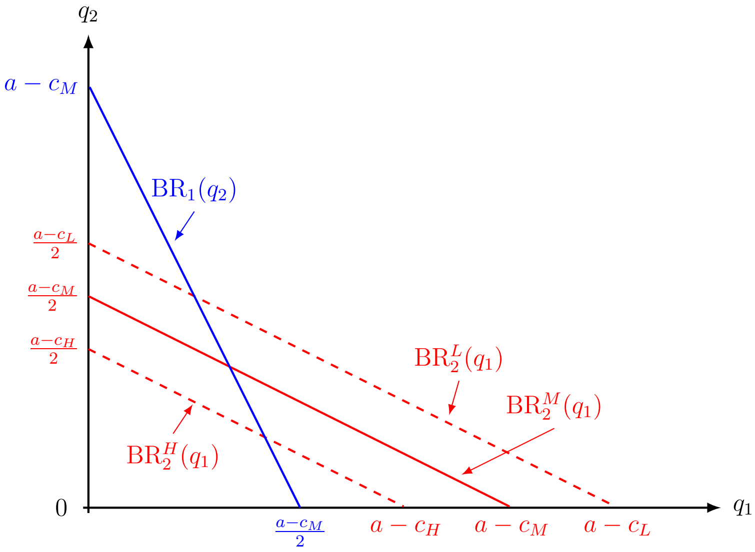

Asymmetric Cournot with Verifiable Information

- The argument is not over.
- Let us now consider whether Firm 2 should reveal its cost information when it has middle cost.
- If Firm 2 doesn't reveal that it has middle cost, then Firm 1 knows that cost is not low.
  - Firm 2 would reveal its cost information if it has low cost, as is argued previously.
- Put it differently, Firm 1 knows that the cost is either middle or high.

Asymmetric Cournot with Verifiable Information

- As a result, Firm 2 with middle cost would want Firm 1 to know it so that Firm 1 would produce less.
- Firm 2 with middle cost would also reveal its cost information.

Asymmetric Cournot with Verifiable Information

For Firm 2 with high cost, it really doesn't matter whether it reveals or not.

- Even if it does not reveal, since Firm 2 with middle or low costs would reveal, the fact of no revealing reveals that Firm 2 has high cost.

Asymmetric Cournot with Verifiable Information

Remark 3.2. The same argument goes through if Firm 2 has more types.

This idea is called information unraveling.

3.B.2. Job-Market Signaling

- Suppose that there are two types of workers, high-ability and low-ability.
- They differ in productivity: high-ability worker has productivity of 100 whereas low-ability worker has productivity of 60.
- In the population, 20% of workers are high-ability and 80% are low-ability.

	Productivity	Proportion
High-ability Worker	100	20%
Low-ability Worker	60	80%

Job-Market Signaling

- Suppose that firms are competitive.
- Firms would offer 100 to a high-ability worker and 60 to a low-ability worker if they could identify worker's types.
- If firms cannot identify worker's types, they would offer $100 * 20\% + 60 * 80\% = 68$.

Question 3.9. Suppose that you are a high-ability worker, how can you make the firms know it?

In particular, would it work if you simply tell the firms "I am a high-ability worker"?

Job-Market Signaling

- Spence (1973) brings up the idea that "education" could be used as a costly signal to differentiate high-ability workers from low-ability ones.
- The crucial assumption in Spence's model is that low-ability workers find education more costly than high-ability workers.
- Assume for a year of education:

	Cost
High-ability Worker	9
Low-ability Worker	21

Job-Market Signaling

When three-year graduate education is available, we argue that there exists an equilibrium where

- High-ability workers take the education but low-ability workers do not.
- Employers identify those workers with graduate degrees as high-ability workers and those without degrees as low-ability workers.
  - Employers offer 100 to a worker with degree and 60 to a worker without degree.

Job-Market Signaling

The solution concept is Perfect Bayesian Equilibrium (PBE).

In essence, PBE requires

1. strategies to be best responses given belief system, and
2. beliefs to be consistent with strategy profile.

For this particular game, we need to check

1. Both types of workers would not deviate in their respective education choices.
2. Employers' beliefs are consistent with the equilibrium behavior.

The second point is obvious.

Job-Market Signaling

For the first point,

- A high-ability worker obtains $100 - 9 * 3 = 73$ if he/she takes education and 60 if not.
- A low-ability worker obtains $100 - 21 * 3 = 37$ if he/she takes education and 60 if not.

Thus, a high-ability worker would not deviate to not taking education and a low-ability worker would not deviate to taking education.

Remark 3.3. This is called a separating equilibrium because in equilibrium the types separate and get identified.

Question 3.10. What if the education program only takes two years? How about one year?

Remark 3.4. For separation to work, there must be enough differences in costs for two types of workers.

Remark 3.5. If standard of obtaining education becomes lower, then probably we will see qualification inflation.

Remark 3.6. Education increases inequality: Compared to the no education outcome, a three-year education program makes high-ability workers better-off ($73 > 68$) and low-ability workers worse-off ($60 < 68$).

Remark 3.7. It is possible that high-ability workers are also worse-off.

- To see this, consider a four-year education program.
- In separating equilibrium, no education is interpreted as evidence of low ability.

3.B.3. Screening

- In the last section, we have seen a signaling model in which informed parties (i.e., workers) move first.
- Signaling models are closely related to screening models, in which uninformed parties take the lead.
- Classic references of screening models concern insurance markets.
- In this course, we still take job market as example.

Job Market Example

Now consider the following timing, which corresponds to a screening setting:

1. Two firms simultaneously announce a menu of contracts specifying required years of education and wage offer (e, w).
2. Given these contracts, workers choose which contract to accept, if any.

Job Market Example

Question 3.11. Is it an equilibrium that both firms offer the same two contracts

$$(e_H = 3, w_H = 100) \text{ and } (e_L = 0, w_L = 60)?$$

Job Market Example

- For workers, similar arguments as in signaling model apply.
- Both types of workers would self-select the contracts designed for them.
  - A high-ability worker obtains $100 - 9 * 3 = 73$ if he/she takes contract $(e_H = 3, w_H = 100)$ and 60 if takes $(e_L = 0, w_L = 60)$.
  - A low-ability worker obtains $100 - 21 * 3 = 37$ if he/she takes contract $(e_H = 3, w_H = 100)$ and 60 if takes $(e_L = 0, w_L = 60)$.

Job Market Example

- In the proposed equilibrium, each firm obtains 0.
- A firm could be better-off by offering $(e'_H = 2, w'_H = 95)$ and $(e_L = 0, w_L = 60)$.
  - High-ability workers prefer $(e'_H = 2, w'_H = 95)$ to $(e_H = 3, w_H = 100)$: they obtain $95 - 9 * 2 = 77$ (> 73) if taking $(e'_H = 2, w'_H = 95)$.
  - Low-ability workers would not take $(e'_H = 2, w'_H = 95)$: they obtain $95 - 21 * 2 = 53$ (< 60) if taking $(e'_H = 2, w'_H = 95)$.

The firm obtains $(100 - 95) * 20\% = 1 > 0$.

Job Market Example

Question 3.12. How about both firms offer the same two contracts $(e_H = 2, w_H = 100)$ and $(e_L = 0, w_L = 60)$?

Remark 3.8. Separating equilibria do not always exist. For example, if we change the proportion of high-ability workers to 80%, then there will be no separating equilibria.

3.B.4. Perfect Bayesian Equilibrium

Solution concept associated with dynamic games of incomplete information is Perfect Bayesian Equilibrium (PBE).

- PBE was invented in order to refine BNE in a similar way that SPE refined NE.
- A complementary perspective: PBE strengthens requirements of SPE by explicitly analyzing players' beliefs, as in BNE.
- We will introduce features of PBE from this complementary perspective.

Example 3.B.1. Consider the following game:

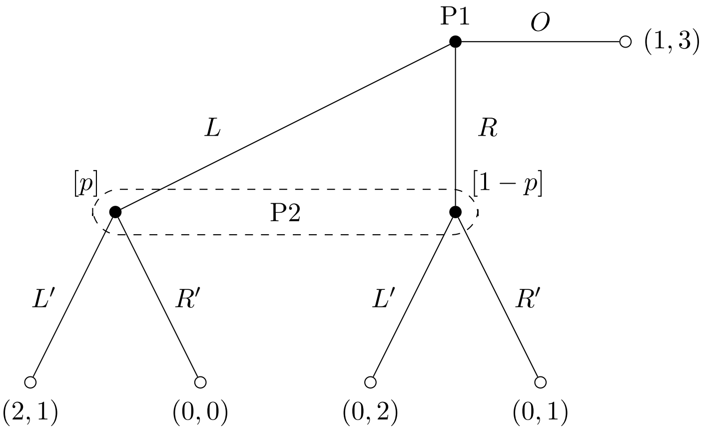

Example 3.B.1

Question 3.13. What are the pure-strategy NEs?

Example 3.B.1

Question 3.14. Are these NEs subgame perfect?

Example 3.B.2. Consider the following modified game:

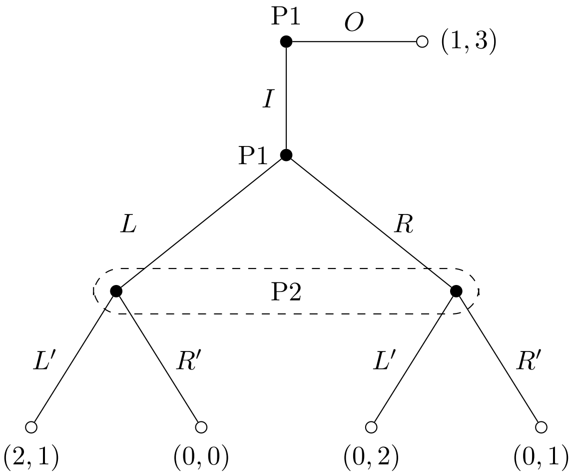

Example 3.B.2

Question 3.15. What are pure strategy NEs of this game?

Are these NEs subgame perfect?

Question 3.16. Are the games in Examples 3.B.1 and 3.B.2 really different?

Question 3.17. Is the equilibrium (O, R') in Example 3.B.1 reasonable?

Sequential Rationality

To rule out the unreasonable prediction (O, R'), we impose sequential rationality requirement:

Requirement 1 (Sequential rationality)**.** At each information set, action taken by player with the move (and player's subsequent strategy) must be optimal given player's belief at information set and other players' subsequent strategies.

Sequential Rationality

Let us apply sequential rationality requirement to the (O, R') equilibrium in Example 3.B.1.

- Let p denote player 2's belief that L has been chosen when the game reaches the information set.
- Given this belief, the expected payoff from choosing
  - R' is $p \cdot 0 + (1-p) \cdot 1 = 1 - p$,
  - L' is $p \cdot 1 + (1-p) \cdot 2 = 2 - p$.
- Since $2 - p > 1 - p$ for any value of p, it is never sequentially rational for player 2 to choose R'.

Perfect Bayesian Equilibrium

- We have only claimed that player 2 should have beliefs and act optimally given the belief.
- However, we have not yet discussed what beliefs are reasonable.
- In order to impose such requirements, we first need to distinguish information sets on the equilibrium path and off the equilibrium path.

Definition 3.B.1. For a given equilibrium in a given extensive-form game, an information set is

- on the equilibrium path if it will be reached with positive probability if game is played according to equilibrium strategies, and
- off the equilibrium path if it is certain not to be reached if game is played according to equilibrium strategies.

Belief Consistency

Requirement 2 (Belief consistency (on the equilibrium path)**).** At information sets on the equilibrium path, beliefs are determined by Bayes' rule and players' equilibrium strategies.

Requirement 3 (Belief consistency (off the equilibrium path)**).** At information sets off the equilibrium path, beliefs are determined by Bayes' rule and players' equilibrium strategies where possible.

Belief Consistency

Let us apply the belief consistency requirements to the (L, L') equilibrium in Example 3.B.1.

- Player 2's belief must be $p = 1$: given player 1's equilibrium strategy (namely, L), player 2 knows which node in the information set has been reached.

Belief Consistency

As an illustration, suppose that in Example 3.B.1 there were a mixed-strategy equilibrium in which player 1 plays

- L with probability q_1,
- R with probability q_2, and
- O with probability $1 - q_1 - q_2$.

Belief consistency requires player 2's belief to be $p = \frac{q_1}{q_1+q_2}$.

Perfect Bayesian Equilibrium

Definition 3.B.2 (Perfect Bayesian Equilibrium)**.** A Perfect Bayesian Equilibrium (PBE) is a strategy profile σ and a belief system $\mu = (\mu_1, \mu_2, ..., \mu_n)$ where μ_i specifies Player i's belief at each of his/her information sets, satisfying Requirements 1 to 3.

Perfect Bayesian Equilibrium

Note that because PBE makes players' beliefs explicit, such an equilibrium often cannot be constructed by working backwards through game tree, as we did to construct a SPE.

Example 3.B.3. Consider the following signaling game:

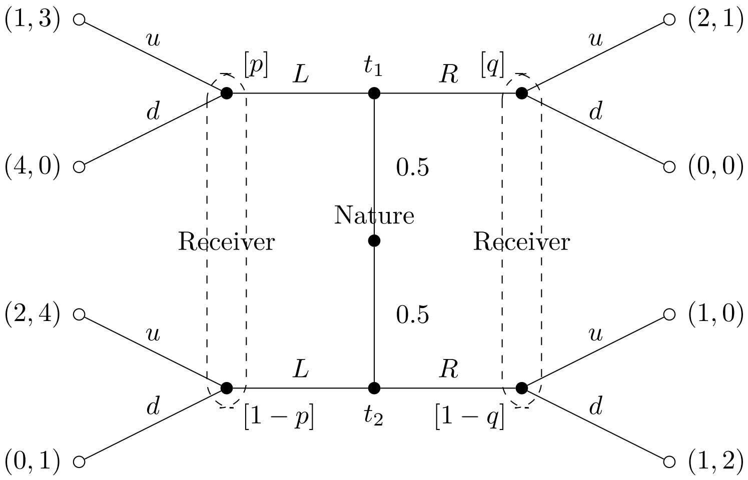

Example 3.B.3

- Sender's type space: $T = \{t_1, t_2\}$;
- Sender's action space: $M = \{L, R\}$ for both t_1 and t_2;
- Receiver's action space: $A = \{u, d\}$, independent of sender's message.

Example 3.B.3

- Recall that (in any game) a player's strategy is a complete plan of action.
- In a signaling game,
  - pure strategy for Sender: a function $m(t_i)$
  - pure strategy for Receiver: a function $a(m_j)$

Example 3.B.3

In this game, there are four possible pure-strategy perfect Bayesian equilibria based on Sender's pure strategy:

1. pooling on L;
2. pooling on R;
3. separation with t_1 playing L and t_2 playing R; and
4. separation with t_1 playing R and t_2 playing L.

Pooling on L

- Receiver's information set corresponding to L is on the equilibrium path and $p = 0.5$.
- Given this belief, Receiver's best response to L is u.
- To ensure both Sender types are willing to choose L, we need to specify how Receiver would react to R.
  - If Receiver's response to R is u, t_1's payoff from playing R is $2\,(> 1)$.
  - If Receiver's response to R is d, t_1 and t_2 earn payoffs of $0\,(< 1)$ and $1\,(< 2)$ from playing R

Pooling on L

- Thus, if there is an equilibrium in which Sender's strategy is (L, L), Receiver's response to R must be d.
- So Receiver's strategy is (u, d).
- It remains to consider Receiver's belief at information set corresponding to R.
  - For playing d to be optimal, we require

    $$q \cdot 0 + (1 - q) \cdot 2 \geq q \cdot 1 + (1 - q) \cdot 0 \implies q \leq 2/3.$$

Thus, $((L, L), (u, d), p = 0.5, q)$ is a pooling perfect Bayesian equilibrium for any $q \leq 2/3$.

Pooling on R

- $q = 0.5$
- Receiver's best response to R is d.
- Again we need to specify Receiver's reaction to L.
  - t_1 can at least obtain $1\,(> 0)$ from playing L
- There is no equilibrium in which Sender plays (R, R).

Separation, with t_1 playing L and t_2 playing R

- $p = 1$ and $q = 0$
- Receiver's best response is (u, d).
- It remains to check whether Sender's strategy is optimal given Receiver's strategy (u, d).
- It is not: if t_2 deviates by playing L, t_2 earns $2\,(> 1)$.

Separation, with t_1 playing R and t_2 playing L

- $p = 0$ and $q = 1$
- Receiver's best response is (u, u).
- It remains to check whether Sender's strategy is optimal given Receiver's strategy (u, u).
  - If t_1 deviates by playing L, t_1 earns $1\,(< 2)$.
  - If t_2 deviates by playing R, t_2 earns $1\,(< 2)$.

Thus, $((R, L), (u, u), p = 0, q = 1)$ is a separating perfect Bayesian equilibrium.

3.B.5. Refinement of Perfect Bayesian Equilibrium

Example 3.B.4. Consider the following signalling game:

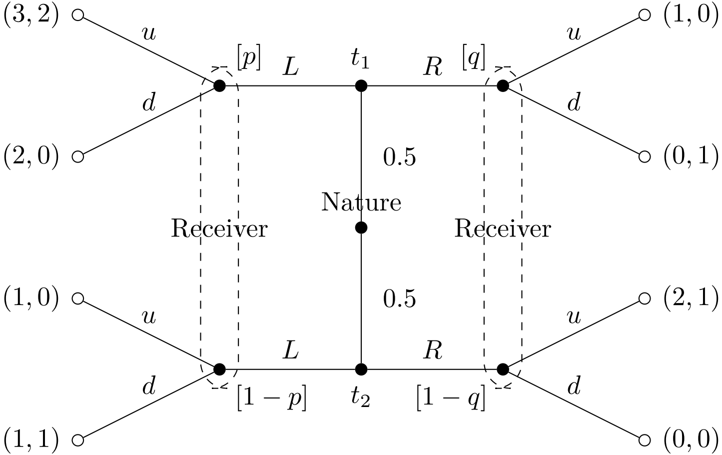

Example 3.B.4

Question 3.18. Can you find pure-strategy PBEs?

There are two pure-strategy PBEs:

- one pooling: $((L, L), (u, d), p = 0.5, q)$ for any $q \geq 1/2$;
- one separating: $((L, R), (u, u), p = 1, q = 0)$.

Example 3.B.4

Question 3.19. Is the pooling equilibrium reasonable? In particular, is the off-the-equilibrium-path belief $q \geq 1/2$ reasonable?

Example 3.B.4

The pooling equilibrium is not reasonable because

it makes no sense for t_1 to play R:

- if t_1 plays L, the lowest payoff is 2;
- if t_1 plays R, the highest payoff is 1.

However, the belief $q \geq 1/2$ means Receiver believes that a deviation to R is very likely from t_1.

Requirement 4

To eliminate such unreasonable predictions, we impose the following requirement:

Requirement 4 (Signalling)**.** If the information set following m_j is off the equilibrium path and m_j is dominated for type t_i, then (if possible) Receiver's belief $\mu(t_i \mid m_j)$ should place zero probability on type t_i.

Requirement 4

The definition of a message being dominated for a type in the requirement is as follows:

Definition 3.B.3. In a signaling game, the message m_j from M is dominated for type t_i from T if there exists another message $m_{j'}$ from M such that t_i's lowest possible payoff from $m_{j'}$ is greater than t_i's highest possible payoff from m_j:

$$\min_{a_k \in A} U_S(t_i, m_{j'}, a_k) > \max_{a_k \in A} U_S(t_i, m_j, a_k).$$

Requirement 4

Applying Requirement 4 to the pooling equilibrium in Example 3.B.4:

- We require $q = 0$.
- Since the pooling equilibrium is a PBE only if $q \geq 1/2$, this equilibrium cannot satisfy Requirement 4.

On the other hand, the separating equilibrium satisfies Requirement 4 trivially.

Example 3.B.5. Consider the "Beer and Quiche" game:

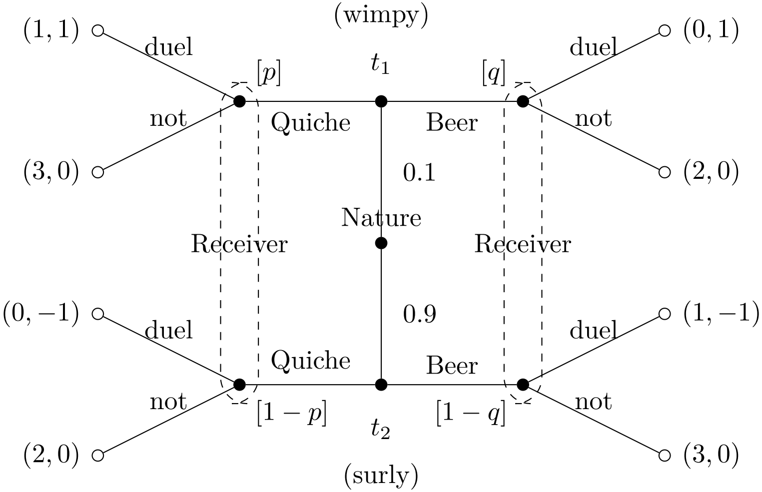

Example 3.B.5

Question 3.20. Can you find pure-strategy PBEs?

There are two pure-strategy PBEs, both are pooling:

- $((\text{Quiche}, \text{Quiche}), (\text{not}, \text{duel}), p = 0.1, q)$

  for any $q \geq 1/2$;
- $((\text{Beer}, \text{Beer}), (\text{duel}, \text{not}), p, q = 0.1)$

  for any $p \geq 1/2$.

Example 3.B.5

Question 3.21. Do these equilibria satisfy Requirement 4?

Yes. Because both Beer and Quiche are not dominated for either Sender type.

Example 3.B.5

Question 3.22. The first pooling equilibrium requires Receiver to believe that Sender is very likely to be of surly type ($q \geq 1/2$) if the off-the-equilibrium-path message Beer is observed. Is it reasonable?

Example 3.B.5

The first pooling equilibrium is not reasonable because:

1. wimpy type cannot possibly improve on the equilibrium payoff of 3
2. surly type could improve on the equilibrium payoff of 2 if Receiver held a belief $q < 1/2$

Intuitive Criterion

To eliminate such unreasonable predictions, we impose the following requirement:

Requirement 5 ("The Intuitive Criterion", Cho and Kreps (1987))**.** If the information set following m_j is off the equilibrium path and m_j is equilibrium-dominated for type t_i, then (if possible) Receiver's belief $\mu(t_i \mid m_j)$ should place zero probability on type t_i.

Intuitive Criterion

The definition of a message being equilibrium-dominated for a type in the requirement is as follows:

Definition 3.B.4. Given a perfect Bayesian equilibrium in a signaling game, the message m_j from M is equilibrium-dominated for type t_i from T if t_i's equilibrium payoff, denoted by $U^*(t_i)$, is greater than t_i's highest possible payoff from m_j:

$$U^*(t_i) > \max_{a_k \in A} U_S(t_i, m_j, a_k).$$

Intuitive Criterion

Applying Requirement 5 to the first pooling equilibrium in Example 3.B.5:

- We require $q = 0$.
- Since the pooling equilibrium is a PBE only if $q \geq 1/2$, this equilibrium cannot satisfy Requirement 5.

On the other hand, the second pooling equilibrium satisfies Requirement 5.

Intuitive Criterion

Remark 3.9. Arguments in the spirit of Requirement 5 are sometimes said to use forward induction: in interpreting a deviation—that is, in forming belief $\mu(t_i \mid m_j)$—Receiver asks whether Sender's past behavior could have been rational.