Game Theory: Static Games of Incomplete Information

Guillaume Pommey

Tor Vergata University of Rome
guillaume.pommey@uniroma2.eu

Table of Contents

Introducing example

Recall the simple **pure coordination game**

X \Y	Jazz Club	Rock Club
Jazz Club	1,1	0,0
Rock Club	0,0	1,1

- Static game of *complete* information
  - Simultaneous play
  - **Perfect knowledge** of the other player's payoffs

Introducing example

- What if you have a doubt about whether your friend Y really likes Jazz music?

1. If Y likes Jazz

X\Y	Jazz	Rock
Jazz	1,1	0,0
Rock	0,0	1,1

2. If Y hates Jazz

X\Y	Jazz	Rock
Jazz	1,-1	0,0
Rock	0,0	1,1

- You are *uncertain* whether you are playing game **1** or game **2**.
- But you have some idea about the **likelihood of each scenario**
  - You think that there is a chance of 2/3 that Y likes Jazz

What's new?

Incomplete information games

- **Incomplete information:** there is uncertainty about the payoffs of other players
  - At least one player must be **uncertain**
  - Here X does not know for sure if Y likes or hates Jazz music.
- That is, payoffs are not **common knowledge**
- Other examples: Bidding games, Cournot with unknown marginal costs, Poker

Reminder: Static games of complete information

Complete information games

- Set of players $N = \{1, \dots, n\}$
- Set of actions $A_1, \dots, A_n$
  - Let $A := A_1 \times \cdots \times A_n$
- Payoffs $u_i : A \to \mathbb{R}$
  - For every vector of actions $(a_1, \dots, a_n) \in A$ player i's payoff is defined by $u_i(a_1, \dots, a_n)$.
- Strategies and actions **coincide** in static game of CI
- (s_i^*, s_{-i}^*) is a Nash equilibrium if

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \text{for all } i \in N, s_i \in S_i$$

Table of Contents

Types and payoffs

We introduce a new notion: **Types**

▷ conveys some information about players' characteristics

- T_i is player i's **Type space**
- $t_i \in T_i$ is player i's type

New definition of payoffs: $u_i : A \times T_i \to \mathbb{R}$

- For each action profile and type we write $u_i(a_1, \ldots, a_n; t_i)$

Jazz/Rock example:

▷ $A_1 = A_2 = \{\text{Jazz}, \text{Rock}\}$

▷ $T_Y = \{t_{Y1}, t_{Y2}\} = \{\text{Likes Jazz}, \text{Hates Jazz}\}$

▷ $u_Y(J, J; t_{Y1}) = 1$ and $u_Y(J, J; t_{Y2}) = -1$

Types and payoffs

Richer payoff definition

We could define $u_i : A \times T \to \mathbb{R}$ where $T := T_1 \times \cdots \times T_n$

- That is $u_i(a_1, \ldots, a_n; t_1, \ldots, t_n)$

In other words, **my payoff can depend on the actions of all players as well as on the type of all players**.

- **Jazz/Rock example:** if player X dislikes the idea that player Y will hate their evening listening to jazz, we may think that player's X utility depends on player Y's type.

We can also use this definition to assume that some **actions are not available to some players** with certain types

Types and non-available actions

Example: Assume you **draw two cards in a regular deck** and then you have to play one.

- Your type represents the **cards you hold in your hands**, e.g., $t_i = \{(Ace, Ace)\}$ or $t_i = \{(Queen, Jack)\}$
- Let $A_i = \{\text{Ace}, \text{King}, \text{Queen}, \ldots, 2, 1\}$
- Obviously you cannot choose $a_i = Ace$ if $t_i = \{(Queen, Jack)\}$
- We can simply say that if you play "Ace" when you do not have one, then your payoff is $-\infty$ so that this action *virtually* disappears

More generally: Any action $a_i \in A_i$ such that $u_i(a_i, a_{-i}; t_i) = -\infty$ is not available to player i with type t_i.

Important: What is known/unknown?

Known: Type spaces $T_1, \ldots, T_n$ and payoff functions $u_1(a_1, \ldots, a_n; t_1, \ldots, t_n), \ldots, u_n(\cdot\,;\cdot)$ are all **common knowledge**.

Unknown: player i does not know what is t_{-i}.

- Each player knows all *possible* games
- But does not know which one is actually played

Next question: How to handle a game that contains several potential games?

Table of Contents

Beliefs

Players are assumed to:

- **Know their own type only** (i knows t_i)
- Have **beliefs** about the types of the other players (t_{-i})

A belief is a **probability distribution** over the types of the other players **given my own type**.

Formally, the belief of player i of type t_i about all other players' types writes

$$p_i(t_{-i} \mid t_i) \in [0, 1]$$

where

$$\sum_{t_{-i} \in T_{-i}} p_i(t_{-i} \mid t_i) = 1$$

Beliefs: Dependent/Independent

The notation $p_i(t_{-i} \mid t_i)$ allows for **interdependent** types.

If types are **independent** we usually use the more convenient notation: $p_i(t_{-i})$.

Examples:

- **Independent types:** You play a game with a friend to decide who will get the chocolate ice cream and who will get the vanilla one. Each knows which her/his favorite, but you are both uncertain about your friend's taste. Your belief about your friend's taste is unlikely to depend on your own taste.

- **Interdependent types:** You and your friend draw one card in a 3-card deck containing only one Ace, one King and one Queen. You look at your card and you see that it is the King. You deduce that your friend can only have the Ace or the Queen. Knowing your type helps you *better predict* the type of the other player.

Beliefs: Priors

How to compute $p_i(t_{-i} \mid t_i)$?

Usually, we first assume that players have a **common prior** on the joint distribution of types.

Let $X_1, \ldots, X_n$ denote the random variables associated with players $1, \ldots, n$ then define:

- $p(t_1, \ldots, t_n) := \mathbb{P}(X_1 = t_1, \ldots, X_n = t_n)$
- for convenience, we sometimes write $p(t_1, \ldots, t_n) =: p(t_i, t_{-i})$

Naturally, $\sum_{t \in T} p(t) = 1$.

Beliefs: Priors and Bayes Rule

Common knowledge: The joint probability distribution $p(t_1, \ldots, t_n)$.

Private information: Only player i knows their type t_i (i.e. realization of X_i).

To form their belief, each player uses their own private information to "guess better" the types of other players.

Bayes Rule: For two players for instance

$$p_i(t_j \mid t_i) = \frac{p(t_i, t_j)}{p(t_i)},$$

where $p(t_i) := \mathbb{P}(X_i = t_i)$.

Reminder: Bayes Rule

Reminder: Define a discrete probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

Take any two events A, $B \in \mathcal{F}$ (s.t. $\mathbb{P}(B) \neq 0$) then Bayes' theorem writes:

$$\mathbb{P}(A \mid B) = \frac{\mathbb{P}(A, B)}{\mathbb{P}(B)}.$$

Useful relationships:

- $\mathbb{P}(A, B) = \mathbb{P}(B \mid A)\mathbb{P}(A)$ then $\mathbb{P}(A \mid B) = \frac{\mathbb{P}(B|A)\mathbb{P}(A)}{\mathbb{P}(B)}$
- $\mathbb{P}(A) = \sum_{i=1}^{m} \mathbb{P}(A, B_i)$ where $B_1, \ldots, B_m$ is a partition of $\mathcal{F}$
- A and B are independent if and only if $\mathbb{P}(A, B) = \mathbb{P}(A)\mathbb{P}(B)$

Beliefs: An Example

Two-player card game

- Deck: contains only 1 Ace, 1 King, and 1 Queen
- Each player draws one card and can privately look at it

Type spaces are $T_1 = T_2 = \{A, K, Q\}$.

Well shuffled deck: probability of drawing a card is uniform then

$t_1 \backslash t_2$	A	K	Q
A	0	1/6	1/6
K	1/6	0	1/6
Q	1/6	1/6	0

$$\Leftrightarrow \quad p(t_1, t_2) = \begin{cases} 1/6 & \text{if } \ t_1 \neq t_2 \\ 0 & \text{if } \ t_1 = t_2 \end{cases}$$

Beliefs: An Example

Now assume that player 1 draws a King: $t_1 = K$

- a. Player 1 learns their type (King)
- b. Player 1 knows that Player 2 cannot have a King: $t_2 \neq K$

Using this information, Player 1 forms their belief:

$$p_1(t_2 \mid t_1 = K) = \frac{p(t_1 = K, t_2)}{p(t_1 = K)} \text{ for each } t_2 = A, Q$$

Beliefs: An Example

For instance, let us find $p_1(t_2 = A \mid t_1 = K)$. We first compute

(a) $p(t_1 = K, t_2 = A) = 1/6$

(b) $p(t_1 = K) = \sum_{t_2 = A, K, Q} p(t_1 = K, t_2)$

$$= \underbrace{p(K, A)}_{=1/6} + \underbrace{p(K, K)}_{=0} + \underbrace{p(K, Q)}_{=1/6} + = \frac{1}{3}$$

Using Bayes' formula we have

$$p_1(t_2 = A \mid t_1 = K) = \frac{p(t_1 = K, t_2 = A)}{p(t_1 = K)} = \frac{1/6}{1/3} = \frac{1}{2}$$

Similarly $p_1(K \mid K) = 0$ and $p_1(Q \mid K) = 1/2$

Table of Contents

Normal-form representation

Definition: A static Bayesian game is defined by the following **normal-form representation**:

- Set of players $N = \{1, \ldots, n\}$
- Sets of actions $A_1, A_2, \ldots, A_n$
- Sets of types $T_1, T_2, \ldots, T_n$
- Beliefs $p_1, p_2, \ldots, p_n$
- (ex post) Payoffs $u_i(a_1, \ldots, a_n; t_1, \ldots, t_n)$

Normal-form representation: Example

Example: Jazz/Rock club

- Players: $N = \{X, Y\}$
- Actions: $A_i = \{Jazz, Rock\}$
- Types: $T_X = \{t_X\}$,
  $T_Y = \{t_{Y1}, t_{Y2}\} = \{\text{Likes Jazz}, \text{Hates Jazz}\}$
- Beliefs: $p_X(t_{Y1} \mid t_X) = 2/3$, $p_Y(t_X \mid t_{Y1}) = p_Y(t_X \mid t_{Y2}) = 1$
- (ex post) Payoffs:

1. t_X, t_{Y1}

X\Y	Jazz	Rock
Jazz	1,1	0,0
Rock	0,0	1,1

2. t_X, t_{Y2}

X\Y	Jazz	Rock
Jazz	1,-1	0,0
Rock	0,0	1,1

Normal-form representation

What's new: **types**, **beliefs** and **type-dependent payoffs**.

- Fix a type for each player (no uncertainty) and you obtain a static game of *complete* info
- Normal-form of static incomplete info games:
  - **Collection** of static games of *complete* info on which players have some belief on their probability of occurrence.

1. t_X, t_{Y1}

X\Y	Jazz	Rock
Jazz	1,1	0,0
Rock	0,0	1,1

Probability: 2/3

2. t_X, t_{Y2}

X\Y	Jazz	Rock
Jazz	1,-1	0,0
Rock	0,0	1,1

Probability: 1/3

Table of Contents

Equilibrium concept

We also need to redefine the **equilibrium concept** for incomplete info games.

We must redefine:

- What is a strategy;
- How to evaluate payoffs;
- Equilibrium condition.

Strategies

In static games of CI, strategies and actions coincide

- Not here, we have to redefine the notion of **strategy**

For each type $t_i \in T_i$, **a strategy** is

▷ $s_i : T_i \rightarrow A_i \Leftrightarrow$ a strategy is a function $s_i(t_i) \in A_i$

Player i **anticipates** the fact **that different types** t_j of player j **may play differently**

▷ But is still uncertain which type they will face in the game

You can think of types as *virtual* new players in the game

▷ Each type is a player who **forms a strategy on their own**

▷ Yet, you are not sure which player you will play against

Strategies: An example

Card game: Assume players' action set is $A_i = \{Play, Withdraw\}$.

For each hand, player i decides what to do. For instance

$$s_i(A) = P,\ s_i(K) = P,\ s_i(Q) = W.$$

Player i must also consider the fact that player j will have $s_j(A)$, $s_j(K)$ and $s_j(Q)$ potentially different.

Notation: For convenience, we will often denote the **strategy profile** of a player like this: PPP, PPW, PWP, ...

- ▷ PPP means that types A, K and Q all decide to Play
- ▷ PWP means that types A, K and Q decide to Play, Withdraw and Play, respectively

Expected Payoffs

We now have to model **how players evaluate their utility**.

▷ Depends on what they know.

Information stages

- **Ex ante:** Player i knows nothing about types
- **Interim:** Player i knows their type but not the one of the other players
- **Ex post:** Player i knows all players' types

Recall the notation $u_i(a_1, \ldots, a_n; t_1, \ldots, t_n)$.

▷ This is defined at the **ex post** stage

Expected Payoffs

Complete info: Players are at the **ex post stage**, they know

- ▷ their own type and all other players' types
  - ▷ they know which ex post payoff is relevant
    - ▷ they know which game they play for sure

Incomplete info: Players are at the **interim stage**, they know

- ▷ their own type but not the other players' types
- ▷ the probability of each $t_{-i} \in T_{-i}$
- ▷ the **ex post payoffs** but not which one is relevant
- ▷ that different types may have different strategies

Expected Payoffs

Incomplete info: At the interim stage, player i of type t_i who plays s_i can compute:

- $p_i(t_{-i} \mid t_i)$ **for each** $t_{-i} \in T_{-i}$
- the ex post payoff $u_i(s_i, s_{-i}(t_{-i}); t_i, t_{-i})$ **for each** $t_{-i} \in T_{-i}$ **and each** $s_{-i}(t_{-i})$

Then the **interim expected payoff** for given s_i and $s_{-i}(t_{-i})$ is obtained by *averaging* over all possible $t_{-i} \in T_{-i}$

$$\sum_{t_{-i} \in T_{-i}} p_i(t_{-i} \mid t_i) u_i(s_i, s_{-i}(t_{-i}); t_i, t_{-i})$$

Expected Payoffs: An Example

Previous card game: Assume that if you play and win you get 1, play and loose you get -1/2, withdraw you get 0 and the other player 1 (if they played).

Interim stage: Assume player 1 knows $t_1 = K$ and plays $s_1(K) = P$.

- If player 1 anticipates $s_2(A) = P$ and $s_2(Q) = W$ ($s_2(K)$ is irrelevant), the relevant ex post payoffs are:

$$u_1(P, P; K, A) = -\frac{1}{2},$$
$$u_1(P, W; K, Q) = 1.$$

- Player 1 computes $p_1(A \mid K) = p_1(Q \mid K) = 1/2$ and $p_1(K \mid K) = 0$

The interim expected payoff is

$$p_1(A \mid K)u_1(P, P \mid K, A) + p_1(Q \mid K)u_1(P, W \mid K, Q)$$
$$= \frac{1}{2}(-\frac{1}{2}) + \frac{1}{2}1 = \frac{1}{4}$$

Table of Contents

Bayesian Nash Equilibrium: Definition

First, recall that a strategy s_i is now a function $s_i : T_i \to A_i$

▷ So when we say s_i, we refer to $s_i(t_i)$ for each $t_i \in T_i$

Definition: In a static Bayesian game, the strategy profile $(s_1^*, \ldots, s_n^*)$ is a **pure-strategy Bayesian Nash Equilibrium** if

$$s_i^*(t_i) \in \arg\max_{a_i \in A_i} \sum_{t_{-i} \in T_{-i}} p_i(t_{-i} \mid t_i) u_i(a_i, s_{-i}^*(t_{-i}); t_i, t_{-i})$$

for all $t_i \in T_i$ and $i \in N$.

Bayesian Nash Equilibrium: Intuition

More intuitively, player i

- ▷ considers all the possible games for every possible $t_{-i} \in T_{-i}$
- ▷ computes their interim expected payoff (averaging over games using their belief) as a function of their own action a_i, type t_i and other player's strategies $s_{-i} \in S_{-i}$
- ▷ for each of their possible type $t_i \in T_i$, player i selects the action $a_i \in A_i$ that maximizes this interim EP to best-respond to $s_{-i} \in S_{-i}$

A strategy profile $(s_1^*, \ldots, s_n^*)$ is a Bayesian NE when each $s_i^*(t_i)$ is actually a **best-response** to $s_{-i}^*(t_{-i})$ for all $t_i \in T_i$ and all $i \in N$

- ▷ In other words: no player has a **profitable deviation**

BNE: Equilibrium Strategies

Why do we consider $s_i(t_j)$ for player i of type $t_i \neq t_j$?

- ▷ Why should we define a strategy for a type that does not exist in the game?

Simply because only player i knows that their type is t_i

- ▷ All other players are **uninformed** about player i's type and so they must be able to anticipate what would a *potential* player i of type t_j would do, that is, $s_i(t_j)$

Parallel with **dynamic games of CI**

- ▷ Recall that we had to define a player's strategy at each decision node, even on those never played at equilibrium

BNE: Implicit timing

It is possible to represent Bayesian Games in **extensive form**.

The **implicit timing** is as follows:

0. Nature draws types $(t_1, \ldots, t_n)$ according to $p(t_1, \ldots, t_n)$
1. Each player i privately learns t_i, computes beliefs and interim expected payoffs as a function of a_i and s_{-i}
2. Players simultaneously choose the action that maximizes their interim EP
3. Information is revealed and ex post payoffs $u_i(s_i^*, s_{-i}^*; t_i, t_{-i})$ are received

BNE: An Example

Example: Jazz/Rock Club

- One type for X, two types for Y (likes or hates Jazz)
- Beliefs: $\mathbb{P}(\text{"Player Y likes Jazz"}) = \alpha \in [0, 1]$

1. t_X, t_{Y1}

X\Y	Jazz	Rock
Jazz	1,1	0,0
Rock	0,0	1,1

Prob $= \alpha$

2. t_X, t_{Y2}

X\Y	Jazz	Rock
Jazz	1,-1	0,0
Rock	0,0	1,1

Prob $= 1 - \alpha$

BNE: An Example

First, consider the two possible games separately.

- ▷ Two pure-strat Nash in **1.**
- ▷ Unique pure-strat Nash in **2.**

1. t_X, t_{Y1}

X\Y	Jazz	Rock
Jazz	1,1	0,0
Rock	0,0	1,1

Prob $= \alpha$

2. t_X, t_{Y2}

X\Y	Jazz	Rock
Jazz	1,-1	0,0
Rock	0,0	1,1

Prob $= 1 - \alpha$

BNE: An Example

Information: Nature draws types and

- ▷ Only Y learns if they like or hate jazz.
- ▷ Player X has only a belief $\alpha \in [0, 1]$ that Y likes jazz.

Player Y can **distinguish** which game is actually played
Player X must **form expectations** over the outcomes

Strategies

- ▷ Player X strategy is simply choosing J or R
- ▷ Player Y strategy is choosing a couple of strategies JJ, JR, RJ or RR

Notation: here JR means $s_2(t_{Y1}) = J$ *and* $s_2(t_{Y2}) = R$ *for instance. The first (resp. second) letter is the strategy for the first (resp.second) type*

BNE: An Example

Compute X's interim expected payoff for each action $a_X \in \{J, R\}$ and each strategy of Y (JJ, JR, RJ and RR)

When Player X plays J we have:

$$\mu_X(J, JJ) = \alpha \cdot 1 + (1-\alpha) \cdot 1 = 1$$
$$\mu_X(J, JR) = \alpha \cdot 1 + (1-\alpha) \cdot 0 = \alpha$$
$$\mu_X(J, RJ) = \alpha \cdot 0 + (1-\alpha) \cdot 1 = 1-\alpha$$
$$\mu_X(J, RR) = \alpha \cdot 0 + (1-\alpha) \cdot 0 = 0$$

1. t_X, t_{Y1}

X\Y	Jazz	Rock
Jazz	1,1	0,0
Rock	0,0	1,1

Prob $= \alpha$

2. t_X, t_{Y2}

X\Y	Jazz	Rock
Jazz	1,-1	0,0
Rock	0,0	1,1

Prob $= 1-\alpha$

BNE: An Example

When Player X plays R we have:

$$\mu_X(R, JJ) = \alpha \cdot 0 + (1 - \alpha) \cdot 0 = 0$$
$$\mu_X(R, JR) = \alpha \cdot 0 + (1 - \alpha) \cdot 1 = 1 - \alpha$$
$$\mu_X(R, RJ) = \alpha \cdot 1 + (1 - \alpha) \cdot 0 = \alpha$$
$$\mu_X(R, RR) = \alpha \cdot 1 + (1 - \alpha) \cdot 1 = 1$$

1. t_X, t_{Y1}

X\Y	Jazz	Rock
Jazz	1,1	0,0
Rock	0,0	1,1

Prob $= \alpha$

2. t_X, t_{Y2}

X\Y	Jazz	Rock
Jazz	1,-1	0,0
Rock	0,0	1,1

Prob $= 1 - \alpha$

BNE: An Example

We can create a payoff matrix as follows.

X \Y	JJ	JR	RJ	RR
J	$1\,;(1,-1)$	$\alpha\,;(1,0)$	$1-\alpha\,;(0,-1)$	$0\,;(0,0)$
R	$0\,;(0,0)$	$1-\alpha\,;(0,1)$	$\alpha\,;(1,0)$	$1\,;(1,1)$

Best-responses

- ▷ Player X cannot distinguish types of Y: Best-responds to JJ, JR, RJ and RR by choosing a single action J or R
- ▷ Player Y best-responds to J and R by choosing two actions: One when $t_Y = t_{Y1}$ and one when $t_Y = t_{Y2}$

BNE: An Example

Assume $\alpha > 1 - \alpha$

X \Y	JJ	JR	RJ	RR
J	$1\,;(1,-1)$	$\alpha\,;(1,0)$	$1-\alpha\,;(0,-1)$	$0\,;(0,0)$
R	$0\,;(0,0)$	$1-\alpha\,;(0,1)$	$\alpha\,;(1,0)$	$1\,;(1,1)$

Player X's best-responses

- J to JJ
- J to JR
- R to RJ
- R to RR

Player Y's best-responses

Type t_{Y1}

- J to J
- R to R

Type t_{Y2}

- R to J
- R to R

BNE: An Example

Assume $\alpha > 1 - \alpha$

X \Y	JJ	JR	RJ	RR
J	$\mathbf{1}$; $(\mathbf{1}, -1)$	α ; $(\mathbf{1}, \mathbf{0})$	$1-\alpha$; $(0, -1)$	0 ; $(0, \mathbf{0})$
R	0 ; $(0, 0)$	$1-\alpha$; $(0, \mathbf{1})$	α ; $(\mathbf{1}, 0)$	$\mathbf{1}$; $(\mathbf{1}, \mathbf{1})$

Player X's best-responses

- **J** to *JJ*
- **J** to *JR*
- **R** to *RJ*
- **R** to *RR*

Player Y's best-responses

Type t_{Y1}

- **J** to *J*
- **R** to *R*

Type t_{Y2}

- **R** to *J*
- **R** to *R*

BNE: An Example

Two **pure-strat. Bayesian Nash Equilibria** for $\alpha > 1 - \alpha$.

- ▷ (J, JR)
- ▷ (R, RR)

Notice that $\alpha > 1 - \alpha \Leftrightarrow \alpha > 1/2$.

- ▷ (J, JR) is a BNE as long as X believes it is more likely that Y likes Jazz than hates it.
- ▷ It is easy to show that if we assume instead that $\alpha < 1/2$, then only (R, RR) is a BNE

BNE: Another Example

Consider now the following game.

1\2	L	R
H	2,1	0,0
D	0,0	1,2

Prob $= 1/2$

1\2	L	R
H	2,0	0,2
D	0,1	1,0

Prob $= 1/2$

Assume player 1 has **no information** while player 2 knows which game is played.

- Player 1 has to choose *H* or *D*
- Player 2 has to choose *LL*, *LR*, *RL* or *RR*

BNE: Another Example

Compute 1's interim expected payoff for each action $a_1 \in \{H, Q\}$ and each strategy of Y (LL, LR, RL and RR).

When Player 1 plays H we have:

$$\mu_1(H, LL) = \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 2 = 2$$
$$\mu_1(H, LR) = \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$$
$$\mu_1(H, RL) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 2 = 1$$
$$\mu_1(H, RR) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 0 = 0$$

1\2	L	R
H	2,1	0,0
D	0,0	1,2

Prob = 1/2

1\2	L	R
H	2,0	0,2
D	0,1	1,0

Prob = 1/2

BNE: Another Example

When Player 1 plays D we have:

$$\mu_1(D, LL) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 0 = 0$$
$$\mu_1(D, LR) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 1 = \frac{1}{2}$$
$$\mu_1(D, RL) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 = \frac{1}{2}$$
$$\mu_1(D, RR) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1$$

$1\backslash 2$	L	R
H	2,1	0,0
D	0,0	1,2

Prob $= 1/2$

$1\backslash 2$	L	R
H	2,0	0,2
D	0,1	1,0

Prob $= 1/2$

BNE: Another Example

1 \2	LL	LR	RL	RR
H	2 ; (1, 0)	1 ; (1, 2)	1 ; (0, 0)	0 ; (0, 2)
D	0 ; (0, 1)	$\frac{1}{2}$; (0, 0)	$\frac{1}{2}$; (2, 1)	1 ; (2, 0)

Player 1's best-responses

- ▷ *H* to *LL*
- ▷ *H* to *LR*
- ▷ *H* to *RL*
- ▷ *D* to *RR*

Player 2's best-responses

Type 1

- ▷ *L* to *H*
- ▷ *R* to *D*

Type 2

- ▷ *R* to *H*
- ▷ *L* to *D*

BNE: Another Example

1 \2	LL	LR	RL	RR
H	**2** ; (**1**, 0)	**1** ; (**1**, **2**)	**1** ; (0, 0)	0 ; (0, **2**)
D	0 ; (0, **1**)	$\frac{1}{2}$; (0, 0)	$\frac{1}{2}$; (**2**, **1**)	**1** ; (**2**, 0)

Player 1's best-responses

- **H** to *LL*
- **H** to *LR*
- **H** to *RL*
- **D** to *RR*

Player 2's best-responses

Type 1

- **L** to *H*
- **R** to *D*

Type 2

- **R** to *H*
- **L** to *D*

Unique **pure-strat BNE** is (**H**, **LR**)

Table of Contents

Bayesian Cournot Nash

We can revisit the classical Cournot competition problem in an **incomplete information environment**.

Assume firm 1 is uncertain about firm 2's marginal cost.

- $c_2 \in \{c_L, c_H\}$ with $c_L < c_H$
- $c_1 = c$ is common knowledge

Cournot Game:

- Players: $N = \{\text{Firm } 1, \text{Firm } 2\}$
- Actions: quantities q_1, $q_2 \in [0, +\infty)$
- Types: $c_1 \in T_1 = \{c\}$, $c_2 \in T_2 = \{c_L, c_H\}$
- Beliefs: $p_1(c_H) = \theta \in [0, 1]$
- Ex payoffs: $\pi_i(q_1, q_2; c_i)$

Bayesian Cournot Nash: Setting

Inverse demand: $P(Q) = a - Q$ with $Q = q_1 + q_2$.

Ex post profits write, for $c_1 = c$ and $c_2 = c_L, c_H$,

$$\pi_1(q_1, q_2; c_1) = [a - q_1 - q_2 - c]q_1,$$
$$\pi_2(q_1, q_2; c_2) = [a - q_1 - q_2 - c_2]q_2.$$

Each firm maximizes profits but **be careful!**

- ▷ Firm 1 must consider both types of firm 2, i.e., consider a different strategy for each type of firm 2 and then take the **average profit** between the two possible scenario
- ▷ Firm 2 must consider firm 1's strategy (only one) and **faces two different maximization problems** for each of its type

Bayesian Cournot Nash: Firm 2's problem

Firm 2 faces **two maximization problems** depending on its type

- Two strategies to choose: q_2^L and q_2^H
- Considers only one strategy for firm 1: q_1

$$\text{Type } c_L: \quad \max_{q_2^L} [a - q_1 - q_2^L - c_L] q_2^L,$$
$$\text{Type } c_H: \quad \max_{q_2^H} [a - q_1 - q_2^H - c_H] q_2^H.$$

Bayesian Cournot Nash: Firm 1's problem

Firm 1 faces a **unique maximization problem**.

- ▷ Only sets one strategy: q_1
- ▷ Considers two strategies for firm 2: q_2^L and q_2^H
- ▷ Averages over possible games according to its belief

$$\max_{q_1} \ \theta[a - q_1 - q_2^H - c]q_1 + (1-\theta)[a - q_1 - q_2^L - c]q_1.$$

It should be now clear to you why we **must define strategies for every possible types**, even for the ones that are not in the game after nature have drawn them.

- ▷ Firm 1 needs to consider two different potential strategies for firm 2 to compute its expected payoff even if firm 2 knows that it is of type c_H.

Bayesian Cournot Nash: Firm 2's Best-response

Let us compute the **two possible best-response** for firm 2's.

First-order conditions for each type are:

$$\text{Type } c_L \text{ FOC} : a - q_1 - 2q_2^L - c_L = 0,$$
$$\text{Type } c_H \text{ FOC} : a - q_1 - 2q_2^H - c_H = 0.$$

Then, best-response for each type writes:

$$\text{Type } c_L : \quad q_2^L(q_1) = \frac{a - c_L - q_1}{2},$$
$$\text{Type } c_H : \quad q_2^H(q_1) = \frac{a - c_H - q_1}{2}.$$

Bayesian Cournot Nash: Firm 1's Best-response

Let us now compute the **best-response** of firm 1's.

The **first-order condition** for the interim expected profit writes

$$\text{FOC:} \quad \theta\left[a - 2q_1 - q_2^H - c\right] + (1-\theta)\left[a - 2q_1 - q_2^L - c\right] = 0$$
$$\Leftrightarrow \quad a - 2q_1 - c - \theta q_2^H - (1-\theta) q_2^L = 0$$

Solving for q_1 gives

$$q_1(q_2^L, q_2^H) = \frac{a - c - \theta q_2^H - (1-\theta) q_2^L}{2}$$

Notice that Firm 1's best-response is a function of q_2^L and q_2^H.

- ▷ Each quantity is **weighted by its probability of occurrence** in the BR function

Bayesian Cournot Nash: Equilibrium

To obtain the **BNE**, simply solves the following linear system of **three equations and three unknowns**.

$$\begin{cases} q_1(q_2^L, q_2^H) & = \dfrac{a - c - \theta q_2^H - (1-\theta) q_2^L}{2} \\ q_2^L(q_1) & = \dfrac{a - c_L - q_1}{2} \\ q_2^H(q_1) & = \dfrac{a - c_H - q_1}{2} \end{cases}$$

Put differently we look for a fixed point:

$$\begin{cases} q_1(q_2^L(q_1), q_2^H(q_1)) & = q_1 \\ q_2^L(q_2^L, q_2^{)} & = q_2^L \\ q_2^H(q_2^L, q_2^H) & = q_2^H \end{cases}$$

Bayesian Cournot Nash: Equilibrium

Start by plugging $q_2^L(q_1)$ and $q_2^H(q_1)$ in $q_1(q_2^L(q_1), q_2^H(q_1))$:

$$q_1 = \frac{a-c}{2} - \frac{\theta}{2}\left(\frac{a-c_H-q_1}{2}\right) - \frac{1-\theta}{2}\left(\frac{a-c_L-q_1}{2}\right)$$

$$q_1\left(1-\frac{\theta}{4}-\frac{1-\theta}{4}\right) = \frac{2a-2c-\theta a-(1-\theta)a+\theta c_H+(1-\theta)c_L}{4}$$

Simplifying we obtain

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

Bayesian Cournot Nash: Equilibrium

Now plugging q_1^* in each best-response for firm 2 yields:

$$q_2^L = \frac{a - c_L}{2} - \frac{1}{2}\left(\frac{a - 2c + \theta c_H + (1-\theta)c_L}{3}\right)$$
$$q_2^H = \frac{a - c_H}{2} - \frac{1}{2}\left(\frac{a - 2c + \theta c_H + (1-\theta)c_L}{3}\right)$$

Simplifying and rearranging we get

$$q_2^{L*} = \frac{a - 2c_L + c}{3} - \frac{\theta}{6}\left(c_H - c_L\right)$$
$$q_2^{H*} = \frac{a - 2c_H + c}{3} + \frac{1-\theta}{6}\left(c_H - c_L\right)$$

Bayesian Cournot Nash: Equilibrium

The Bayesian Nash Equilibrium is therefore given by

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

The **uninformed** player, firm 1, plays an **average quantity**.

The **informed** player, firm 2, plays a **different quantity** when of different type

Bayesian Cournot Nash: Equilibrium

If we remove uncertainty, for instance:

▷ $\theta = 1$: Firm 2 is for sure of type c_H.

Then

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

and

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

We exactly obtain the Cournot quantities in a duopoly with certain marginal costs c and c_H for firm 1 and firm 2, respectively

Bayesian Cournot Nash: Equilibrium

Notice also that, under **incomplete information**, Firm 2:

- of type c_H produces **more** than it would produce if information was complete
- of type c_L produces **less** than it would produce if information was complete

This stems from the fact that firm 1 produces an average quantity to adapt each possible cases of firm 2 and firm 2 can use this uncertainty to its own benefit

Table of Contents

Bidding game/Auction

A very interesting application of Bayesian game: **Auctions**

Two potential buyers $i = 1, 2$ are interested in a painting.

- ▷ Each has a **valuation** $v_i \in [0, 1]$ for the painting
- ▷ v_i is **private information** to i

The seller (or auctioneer) proposes the following **selling mechanism**:

- ▷ Each agent $i = 1, 2$ must communicate their **bid** $b_i \in \mathbb{R}_+$ simultaneously
- ▷ The agent with the **highest bid** receives the painting and pays what the amount of their bid to the seller

Bidding game/Auction

We call this bidding game a **sealed-bid first-price** auction.

- **Sealed-bid:** Simultaneous bidding, no player can observe the other player's bid
- **First-price:** The winner pays what they bid, that is, the highest (or first) price

Other types of auctions exist:

- **Sealed-bid second-price auction:** Simultaneous bidding, highest bidder wins but pays the second highest price
- **All-pay auctions:** Several forms but potentially all bidders may pay/receive something even when they loose
- **Descending auctions:** The auctioneer offers a very high price and decreases it until someone raises their hand to buy at this price

Sealed-bid first-price auction: Payoffs

We can write agent i's ex post payoff as follows: v

$$u_i(b_i, b_j; v_i) = \begin{cases} v_i - b_i & \text{if } b_i \geq b_j \\ 0 & \text{if } b_j < b_j \end{cases}$$

(what happens when $b_i = b_j$ does not matter as this event has zero measure).

Each player faces the following **trade-off**:

- ▷ communicate a **low bid** to pay less if awarded the painting
- ▷ communicate a **high bid** to have a higher bid than the other player and be awarded the painting

Sealed-bid first-price auction: Complete information?

Two-sided incomplete information is what makes the auction interesting

Assume each bidder perfectly knows both v_1 and v_2

- ▷ If $v_1 \geq v_2$, bidder 1 will choose $b_1 = v_2$, win the auction for sure and pay the lowest possible price
- ▷ If $v_1 < v_2$, bidder 2 will choose $b_2 = v_1$, win the auction for sure and pay the lowest possible price

It seems more *realistic* to assume that you cannot know for sure the valuation of the other agent

Sealed-bid first-price auction: Normal-form

Normal-form representation

- Players: $N = \{1, 2\}$
- Action space: $A_i = \mathbb{R}_+$
- Type space: $T_i = [0, 1]$, the type is $v_i \in T_i$
- Beliefs: $v_i \sim \mathcal{U}[0, 1]$
- Ex post payoffs:

$$u_i(b_i, b_j; v_i) = \begin{cases} v_i - b_i & \text{if } b_i \geq b_j \\ 0 & \text{if } b_j < b_j \end{cases}$$

SB FP auction: Strategies and Expected Payoffs

Recall that **each type must form a strategy**, i.e., $s_i : T_i \to A_i$

- ▷ A strategy here is $b_i(v_i)$

Player i takes $b_j(v_j)$ as given and takes expectation over all $v_j \in [0, 1]$ to obtain their **interim expected payoff**:

$$\begin{aligned} U_i(b_i; v_i) &= \mathbb{P}(b_i \geq b_j(v_j))[v_i - b_i] + \mathbb{P}(b_i < b_j(v_j)) \cdot 0 \\ &= \mathbb{P}(b_i \geq b_j(v_j))[v_i - b_i]. \end{aligned}$$

The **trade-off** is clear in $U_i(b_i; v_i)$: Increasing b_i,

- ▷ increases $\mathbb{P}(b_i \geq b_j(v_j))$
- ▷ decreases $[v_i - b_i]$

SB FP auction: Equilibrium

There is a **general solution** to this problem.

- ▷ But it requires some technical tricks that are beyond the scope of this class

Instead, let us focus on a **particular solution**:

- ▷ **Linear strategies:** $b_i(v_i) = a_i + c_i v_i$ for $i = 1, 2$

In other words, we **postulate** that $b_j(v_j) = a_j + c_j v_j$ is an equilibrium strategy for player i and we

- ▷ investigate player i's best-response
- ▷ check that player i's best-response is a linear strategy

SB FP auction: Equilibrium

Fix player j's strategy to $b_j(v_j) = a_j + c_j v_j$.

Then we can easily compute

$$\begin{aligned}\mathbb{P}(b_i \geq b_j(v_j)) &= \mathbb{P}(b_i \geq a_j + c_j v_j) \\ &= \mathbb{P}(v_j \leq \frac{b_i - a_j}{c_j}) \\ &= \frac{b_i - a_j}{c_j}.\end{aligned}$$

Again, notice that player i's **winning probability** is increasing in b_i.

Reminder: The CDF of a uniform distribution on $[0, 1]$ writes $F(x) = x$ for all $x \in [0, 1]$. Hence $\mathbb{P}(Z \leq z) = \int_0^1 \mathbb{1}_{\{x \leq z\}} dF(x) = \int_0^z 1 dx = z$.

SB FP auction: Equilibrium

For player i, there is **no uncertainty** on $(v_i - b_i)$.

Therefore, player i's **expected payoff** as a function of a_i and v_i is

$$U_i(b_i; v_i) = \frac{b_i - a_j}{c_j}[v_i - b_i].$$

Player i chooses $b_i(v_i) \in \arg\max_{b_i \in \mathbb{R}_+} U_i(b_i; v_i)$

The solution to the maximization problem $b_i(v_i)$ is a function of v_i, that is, solving the max problem yields player i's best-response for any given v_i

SB FP auction: Equilibrium

The **first-order condition** of i's problem writes

$$\frac{1}{c_j}[v_i - b_i] - \frac{b_i - a_j}{c_j} = 0.$$

Solving for b_i yields

$$b_i(v_i) = \frac{a_j + v_i}{2}$$

By a symmetrical reasoning we get

$$b_j(v_j) = \frac{a_i + v_j}{2}$$

SB FP auction: Equilibrium

We have to find the parameter values a_i, a_j, c_i and c_j.

By simple identification:

As $b(v_i) = a_i + c_i v_i = \frac{a_j}{2} + \frac{1}{2} v_i$, we must have

$$a_i = \frac{a_j}{2} \qquad c_i = \frac{1}{2}.$$

As $b(v_j) = a_j + c_j v_j = \frac{a_i}{2} + \frac{1}{2} v_j$, we must have

$$a_j = \frac{a_i}{2} \qquad c_j = \frac{1}{2}.$$

SB FP auction: Equilibrium

Then $c_i = c_j = \frac{1}{2}$.

Also $a_i = \frac{a_j}{2}$ and $a_j = \frac{a_i}{2}$ are equivalent to

$$2a_i = a_j = \frac{a_i}{2}.$$

It is clear that $a_i = a_j = 0$.

The **linear equilibrium strategies** are then

$$b_i(v_i) = \frac{v_i}{2},$$
$$b_j(v_j) = \frac{v_j}{2}.$$

Table of Contents

Mixed strategies in CI games

Complete information: Recall that in a CI game, a mixed strategy is a probability distribution on pure strategies:

- A mixed strategy is $\sigma_i \in \Delta(A_i)$
- Where $\Delta(A_i) := \left\{(\alpha_1, \ldots, \alpha_{|A_i|}) \in [0,1]^{|A_i|} \mid \sum_{i=1}^{|A_i|} \alpha_i = 1\right\}$

In this context, we assumed that players were

- completely informed on the payoffs of the other players
- randomizing their action at equilibrium

No uncertainty in the **fundamentals** of the game

- Uncertainty arises endogenously as an equilibrium strategy

Mixed strategies in CI games

Sometimes, it may be difficult to interpret what randomizing means.

- ▷ Do players really flip a coin to determine their strategy?

Harsanyi (1973) proposed an **interpretation for mixed strategies** in CI games.

- ▷ Mixed strategies can instead represent **players' payoff uncertainty** as in an **incomplete information** game
- ▷ This holds for a *small* level of uncertainty

Mixed strategies in CI games

More precisely, **Harsanyi (1973)** showed that:

(almost) **every mixed-strategy equilibrium** in a complete information game can be *approached* by a Bayesian game for which uncertainty *is small*.

That is, **equilibrium strategies of the Bayesian game will converge** to the equilibrium mixed strategies of the complete information game.

- ▷ We will be more specific about what we mean by "small" and "converge"

An example

Consider the following game of **complete information**:

1\2	L	R
H	2,1	0,0
D	0,0	1,2

- Two pure-strategy Nash equilibria $\{H, L\}$ and $\{D, R\}$
- **One mixed-strategy Nash equilibrium** in which player 1 plays H with probability $\frac{2}{3}$ and player 2 plays L with probability $\frac{1}{3}$
  - ▷ This is the one that interests us
  - ▷ Do you think anyone would play this instead of pure-strategy NE? Would you?

An example

Consider now the following game of **incomplete information**:

- $N = \{1, 2\}$
- $A_1 = \{H, D\}$, $A_2 = \{L, R\}$
- $T_1 = T_2 = [0, x]$, with $x \in \mathbb{R}_+$
- $t_i \sim \mathcal{U}[0, x]$, $i = 1, 2$

1\2	L	R
H	$2 + t_1$, 1	0,0
D	0,0	1, $2 + t_2$

An example

This game is *almost* like the complete information game except that **two payoffs are now uncertain**

$1\backslash 2$	L	R
H	$2+t_1$, 1	0,0
D	0,0	1, $2+t_2$

Uncertainty is modeled through the $t_i \sim \mathcal{U}[0,x]$.

- ▷ i.e., uniform distribution over the interval $[0,x]$

Notice that uncertainty is *reduced* when x decreases.

- ▷ Extreme case: when x goes to 0 then **uncertainty vanishes**

An example

Our goal is to construct a specific **pure-strategy Bayesian game** and then see what happens **when uncertainty vanishes**.

Assume each player plays the following **threshold strategy**:

- ▷ Player 1 plays H whenever $t_1 \geq a$
- ▷ Player 2 plays R whenever $t_2 \geq b$

for some $a, b \in \mathbb{R}_+$

An example

Player 1 plays H with probability

$$\begin{aligned}\mathbb{P}(t_1 \geq a) &= 1 - \mathbb{P}(t_1 \leq a)\\ &= 1 - \frac{a}{x}\\ &= \frac{x-a}{x}\end{aligned}$$

Similarly, Player 2 plays R with probability

$$\mathbb{P}(t_2 \geq b) = \frac{x-b}{x}$$

An example

Let us now **find the BNE**.

Assume Player 2 plays the threshold strategy "play R if $t_2 \geq b$".

Then player 1's expected payoff when playing H is

$$\begin{aligned}\mu_1(H, \text{play } R \text{ if } t_2 \geq b) &= \mathbb{P}(t_2 < b)\big[2 + t_1\big] + \mathbb{P}(t_2 \geq b) \cdot 0 \\ &= \frac{b}{x}\big[2 + t_1\big].\end{aligned}$$

Player 1's expected payoff when playing D is

$$\begin{aligned}\mu_1(R, \text{play } R \text{ if } t_2 \geq b) &= \mathbb{P}(t_2 < b) \cdot 0 + \mathbb{P}(t_2 \geq b) \cdot 1 \\ &= \frac{x - b}{x}.\end{aligned}$$

An example

For player 1, **playing H is optimal** whenever

$$\begin{aligned} & \mu_1(H, \text{play } R \text{ if } t_2 \geq b) \geq \mu_1(D, \text{play } R \text{ if } t_2 \geq b) \\ \Leftrightarrow \quad & \frac{b}{x}\left[2 + t_1\right] \geq \frac{x-b}{x} \\ \Leftrightarrow \quad & t_1 \geq \frac{x}{b} - 3. \end{aligned}$$

A similar reasoning for player 2 (fixing player 1's strategy to the threshold one) yields that they play R whenever

$$t_2 \geq \frac{x}{a} - 3.$$

An example

Fortunately, the two conditions correspond to the **threshold strategies** that we have **postulated** when we set $a = \frac{x}{b} - 3$ and $b = \frac{x}{a} - 3$.

Solving the system (involves a quadratic equation) yields

$$\mathbb{P}(t_1 \geq a) = \mathbb{P}(t_2 \geq b) = 1 - \frac{-3 + \sqrt{9 + 4x}}{2x}.$$

An example

Now if x converges to 0, **uncertainty vanishes** and we obtain

$$\mathbb{P}(t_1 \geq a) = \mathbb{P}(t_2 \geq b) = \frac{2}{3}.$$

That is, the **pure strategies of the Incomplete information game** are the same **as the mixed strategies of the Complete information game**.

Technical note: $\lim_{x \to 0} \frac{-3+\sqrt{9+4x}}{2x}$ is obtained using l'Hopital's Rule which states that $\lim_{x \to c} \frac{f(x)}{g(x)} = \lim_{x \to c} \frac{f'(x)}{g'(x)}$ whenever $lim_{x \to c} f(x) = lim_{x \to c} g(x) = 0$ and $g'(x) \neq 0$)

An example

Therefore the mixed-strategy Nash equilibrium in the complete information game can also be seen as a pure-strategy Bayesian Nash equilibrium when there is a **very small amount of uncertainty of some payoffs**.

With this interpretation:

- Mixed-strategy in CI game can express the fact that players have a small amount of payoff uncertainty