Microeconomics IV (Game Theory)

Lecture 7 – Dynamic Games with Incomplete Information

Lasha Chochua

TA: Luka Makharadze

2026

Introduction

Dynamic Games with Incomplete Information

- Our final topic is **dynamic games with incomplete information**.
- These games encompass many interesting economic models:
    - **Market signalling**
    - Cheap talk
    - Reputation
- To study these problems:
    - We first introduce "new" **solution concepts**.
    - Then we apply these concepts to the specific game as a demonstration.

Formal Definitions (Reminder)

- As in Lecture 4, a finite extensive-form game Γ^e consists of:
  - A set of players N, with typical player $i \in N$.
  - For each player i, a set of information sets S_i, disjoint across players ($S_i \cap S_j = \emptyset$ for $i \neq j$), with $S^* = \bigcup_{i \in N} S_i$.
  - For each $s \in S^*$: the set of nodes Y_s at which it occurs, and the set of moves D_s available there.
  - A game tree with initial node x^0 and terminal nodes Ω, with payoffs $w_i(y)$ to player i at each $y \in \Omega$.

Information Sets

- An information set $s \in S_i$ collects the nodes Y_s that player i cannot tell apart when called to move.
- The available moves depend only on the information set: every node in Y_s shares the same move set D_s.
- The information sets partition the decision nodes: each lies in exactly one Y_s, and $S^* = \bigcup_{i \in N} S_i$ collects all of them.

The Failure of Subgame Perfection

Problems with Subgame Perfection

- In extensive-form games with incomplete information, subgame perfection is not fully satisfactory.
- Subgame perfection may fail to rule out actions that are sub-optimal given any beliefs about uncertainty.
- Consider the following games:

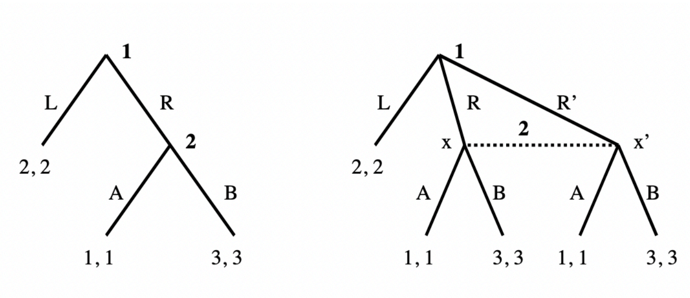

Figure 1: Example 1

SPE and Multiple Equilibria

- In the game on the left, (R, B) is the **only SPE**.
- In the game on the right:
  - $(pR + (1-p)R', B)$ for any p
  - $(L, qA + (1-q)B)$ for any $q \geq 1/2$
- The game on the right has no proper subgames other than the game itself.

Why SPE Fails

- Subgame perfection has **no bite** when subgames do not exist.
- One solution is to require players to choose optimally at *every information set*.
- To do this, we need to introduce the idea of **beliefs**.

- **Example 1 (cont.):**
  - Suppose an information set s with nodes $Y_s = \{x, x'\}$.
  - Let $\pi(x)$ and $\pi(x')$ be beliefs, where $\pi(x) + \pi(x') = 1$.
  - Player 2 must choose the action that maximizes his **expected** payoff.
  - Given any belief, choosing B is optimal.

Example 2

- Consider the second example:

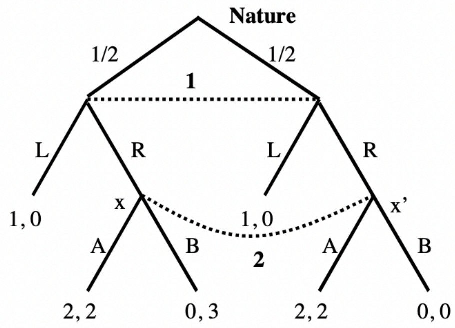

Figure 2: Example 2

Example 2 (Cont.)

- As in the previous example, (L, B) is an SPE because there are no subgames.
- However, there are beliefs for which B is an optimal choice for player 2:
  - If player 2 places probability at least $2/3$ on being at x given that $Y_s = \{x, x'\}$, then B is optimal.
- These beliefs seem quite unreasonable:
  - If player 1 chooses R, then player 2 should place **equal probability** on being at either x or x'.

Perfect Bayesian Equilibrium

Perfect Bayesian Equilibrium

- Let Γ^e be an extensive-form game.
- Let S_i be the set of information sets at which player i moves.

Behavioral Strategies and Assessments

Definition 1: Behavioral Strategy

A *behavioral strategy* for player i assigns to each information set $s \in S_i$ a probability distribution

$$\sigma_{i,s} \in \Delta(D_s)$$

over the moves available at s. We write $\sigma_i = (\sigma_{i,s})_{s \in S_i}$, so that the support of each $\sigma_{i,s}$ is contained in D_s.

Definition 2: Assessment

An *assessment* (σ_i, π_i) for player i consists of a behavioral strategy σ_i and a belief function π_i that assigns to each $s \in S_i$ a distribution $\pi_{i,s} \in \Delta(Y_s)$ over the nodes in Y_s.

Examples and Beliefs

- **Example 1, cont.**
  - Player two's belief function π_2 must satisfy $\pi_2(x) + \pi_2(x') = 1$.
  - If player 2 uses Bayesian updating, and $\sigma_1(R) + \sigma_1(R') > 0$, then:

$$\pi_2(x) = \frac{\sigma_1(R)}{\sigma_1(R) + \sigma_1(R')}.$$

- **Example 2, cont.**
  - Player two's belief function π_2 again satisfies $\pi_2(x) + \pi_2(x') = 1$.
  - If player 2 uses Bayesian updating and $\sigma_1(R) > 0$, then $\pi_2(x) = \pi_2(x') = 1/2$.

Perfect Bayesian Equilibrium (PBE)

Definition 3: Perfect Bayesian Equilibrium

A profile of assessments (σ, π) is a **perfect Bayesian equilibrium** if:

- For all i and all $s \in S_i$, σ_i is **sequentially rational** – it maximizes player i's expected payoff conditional on reaching s, given π_i and σ_{-i}.
- Beliefs π_i are updated using Bayes' rule whenever it applies (i.e. at any information set on the **equilibrium path**).

Perfect Bayesian Equilibrium – Explained

- A **PBE** can be decomposed into **four key requirements** that make the definition precise.

Requirement 1

Every player will have a well-defined belief over where he is in each of his information sets. That is, the game will have a **system of beliefs**.

Requirement 2

Let $\sigma^* = (\sigma_i^*)_{i \in N}$ be a Bayesian Nash equilibrium profile of strategies. We require that in all information sets, beliefs that are **on the equilibrium path** be consistent with **Bayes' rule**.

Perfect Bayesian Equilibrium – Explained (Cont.)

Requirement 3

At information sets that are **off the equilibrium path**, any belief can be assigned to which Bayes' rule does not apply.

- Off-path beliefs are not pinned down by Bayes' rule.
- Players can have arbitrary beliefs at these information sets.

Requirement 4

Given their beliefs, players' strategies must be **sequentially rational**. That is, in every information set, players will play a **best response** to their beliefs.

Refinements of PBE

- PBE allows for beliefs that seem **unreasonable**.
- Problems occur for beliefs off the **equilibrium path**.
- Game theorists have introduced **refinements** to restrict such beliefs.

Example 2 (cont.)

- Consider the earlier game:
  - If $\sigma_1(R) > 0$, then by Bayes' rule $\pi_2(x) = \pi_2(x') = 1/2$, so player 2 must play A. Hence (R, A) is a PBE.
  - However, if $\sigma_1(R) = 0$, Bayes' rule does not apply. Then π_2 is arbitrary.
  - If $\pi_2(x) > 2/3$, then playing B is optimal, so (L, B) is also a PBE!

Sequential Equilibrium

Consistency of Beliefs

Definition 4: Consistency

An assessment (σ, π) is **consistent** if (σ, π) is the limit of some sequence (σ^n, π^n) where σ^n is totally mixed and π^n is derived from σ^n by Bayes' rule.

Consistency of Beliefs (Example 2)

- In this game, the only **consistent** beliefs for player 2 are:

$$\pi_2(x) = \pi_2(x') = \frac{1}{2}.$$

- Why? Consider any sequence of completely mixed strategies (σ_1^n) for player 1:

$$\sigma_1^n(L) > 0, \ \sigma_1^n(R) > 0 \implies \pi_2^n(x) = \frac{\sigma_1^n(R)}{\sigma_1^n(R) + \sigma_1^n(R')} = \frac{1}{2}.$$

- Taking the limit as $n \to \infty$ yields:

$$\lim_{n\to\infty} \pi_2^n(x) = \frac{1}{2}, \text{ so } \pi_2(x) = \pi_2(x') = \frac{1}{2}.$$

Sequential Equilibrium

Definition 5: Sequential Equilibrium

An assessment (σ, π) is a **sequential equilibrium** if (σ, π) is both consistent and a PBE.

- Sequential equilibrium refines PBE by ensuring consistency.
- It is stricter but also harder to apply in practice.
- Even so, some unreasonable outcomes can persist.
- Other refinements (e.g. the **intuitive criterion**) aim to restrict these further.

Application – Signalling Games

Signalling

- We now consider an important class of dynamic games with incomplete information.
- These *signalling models* were introduced by Spence (1974) in his Ph.D. thesis.
- Two players and two periods.

Timing

- **Stage 0**
  - Nature chooses type $\theta \in \Theta$ of player 1 from a distribution p.
- **Stage 1**
  - Player 1 observes θ and chooses $a_1 \in A_1$.
- **Stage 2**
  - Player 2 observes a_1 and chooses $a_2 \in A_2$.
- **Payoffs**
  - $u_1(a_1, a_2, \theta)$ and $u_2(a_1, a_2, \theta)$.

Examples of Signalling Games

- **Example: Job Market Signalling**
  - Player 1: student or worker with privately known ability θ.
  - Action: education level a_1.
  - Player 2: competitive labor market.
  - Market observes a_1 and offers wage based on expected ability.
  - The worker wants to signal high ability.

Examples (cont.)

- **Example: Initial Public Offerings**
  - Player 1: entrepreneur (private firm owner), type $=$ future profitability θ.
  - Action: sell a fraction of company at price a_1.
  - Player 2: investors accept or reject offer.
  - Entrepreneur wants to signal high profitability.
- **Example: Monetary Policy**
  - Player 1: Federal Reserve (preferences for inflation vs unemployment), type $= \theta$.
  - Action: choose inflation level a_1.
  - Player 2: firms set expectations a_2.
  - Fed wants to signal a preference for low inflation.

Examples (cont.)

- **Example: Pretrial Negotiation**
  - Player 1: defendant with private liability θ.
  - Action: settlement offer a_1.
  - Player 2: plaintiff accepts or rejects ($a_2 \in \{A, R\}$).
  - If rejected, trial occurs.
  - Defendant wants to signal a strong case.

Perfect Bayesian Equilibrium in Signalling Models

Definition 6: PBE in Signalling Models

A *perfect Bayesian equilibrium* in the signalling model is a strategy profile $s_1(\theta)$, $s_2(a_1)$ together with beliefs $\pi_2(\theta \mid a_1)$ for player two such that the conditions on the next slide hold.

PBE Conditions

1. Player one's strategy is optimal given player two's strategy:

$$s_1(\theta) \in \arg\max_{a_1 \in A_1} u_1(a_1, s_2(a_1), \theta), \text{ for all } \theta \in \Theta.$$

2. Player two's beliefs are updated by Bayes' rule:

$$\pi_2(\theta \mid a_1) = \frac{\Pr(s_1(\theta) = a_1)\, p(\theta)}{\sum_{\theta' \in \Theta} \Pr(s_1(\theta') = a_1)\, p(\theta')}$$

The prior probability of type θ choosing a_1, divided by the total probability that any type could have chosen a_1.

if $\Pr(s_1(\theta) = a_1) > 0$. Otherwise, $\pi_2(\theta \mid a_1)$ is arbitrary.

3. Player two's strategy is optimal given his beliefs:

$$s_2(a_1) \in \arg\max_{a_2 \in A_2} \sum_{\theta \in \Theta} u_2(a_1, a_2, \theta)\, \pi_2(\theta \mid a_1), \text{ for all } a_1 \in A_1.$$

Mixed Strategies

- It is straightforward to allow for mixed strategies:
  - Denote player i's mixed strategy by σ_i.
  - Beliefs must then be derived from σ_1 using Bayes' rule wherever possible.

Types of Perfect Bayesian Equilibrium

- **Separating**:
  - Different types choose different actions.
  - Player 2 perfectly infers player 1's type.
- **Pooling**:
  - All types choose the same action.
  - Player 2 learns nothing about player 1's type.
- **Semi-Separating**:
  - Some types choose different actions; other types mix.
  - Player 2 partially updates beliefs, imperfect inference.

Job Market Signalling

Job Market Signalling

- **Setting**:
  - One worker with private ability $\theta \in \{\theta_L, \theta_H\}$, where $\theta_H > \theta_L$.
  - Labor market's prior belief that the worker is high-type: λ.
  - The worker chooses education level e.
  - Education is costly: $c(e, \theta)$ depends on ability.

- $c_e > 0$: education is costly at the margin.
- $c_{e\theta} < 0$: marginal cost of education is lower for high-ability workers.

Timing and Payoff

- After choosing e, the firm observes e and offers a competitive wage:

$$w(e) = \pi(e)\theta_H + (1 - \pi(e))\theta_L,$$

- where $\pi(e)$ is the firm's belief (posterior probability) that the worker is high-ability conditional on observing e.
- The worker solves:

$$\max_e \; w(e) - c(e, \theta).$$

  - High-ability workers may choose costly education to **signal** their ability.
  - Even though education has **no direct effect** on productivity, it influences beliefs and wages.

Indifference Curves – Single-Crossing Property (I)

- Worker's utility:
$$u(w, e, \theta) = w - c(e, \theta).$$
- Holding utility fixed at U:
$$w - c(e, \theta) = U.$$
- Implicit differentiation gives:
$$dw - c_e(e, \theta)de = 0,$$
so:
$$\left.\frac{dw}{de}\right|_{u=U} = c_e(e, \theta) > 0.$$

Indifference Curves – Single-Crossing Property (II)

- Indifference curves slope upward: more education requires a higher wage to keep utility constant.
- Since:
$$c_{e\theta}(e,\theta) < 0,$$
we have:
$$\frac{d}{d\theta}\left(\frac{dw}{de}\bigg|_{u=U}\right) = c_{e\theta}(e,\theta) < 0.$$
- Therefore, indifference curves are flatter for high-ability workers:
$$c_e(e,\theta_H) < c_e(e,\theta_L).$$
- This is the **Spence-Mirrlees single-crossing property**: the education signal is less costly at the margin for high-ability workers than for low-ability workers.

Indifference Curves in (e, w) Space

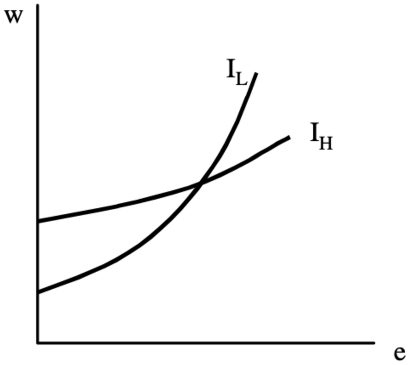

Figure 3: Indifference Curves

Equilibrium

- Consider the key question: what is $w(e)$ (equivalently $\pi(e)$)?
- On the equilibrium path, $w(e)$ is determined by the worker's choice.
- Off the equilibrium path, PBE imposes no restriction on beliefs.

- This flexibility in setting $w(e)$ gives rise to **many possible equilibria**.

Separating Equilibrium

- In a separating equilibrium, we have:

$$e(\theta_H) \neq e(\theta_L).$$

Important

Claim 1. In a separating equilibrium, $w(e(\theta)) = \theta$ for $\theta \in \{\theta_L, \theta_H\}$. That is, workers receive their marginal product as wage.

- In a PBE, beliefs are derived from Bayes' rule whenever possible.
- Type θ_L workers always choose $e(\theta_L)$, while type θ_H workers always choose $e(\theta_H)$.
- Thus, if $e(\theta)$ is observed, the market must believe the worker is type θ for sure. Hence, $w(e(\theta)) = \theta$.

Separating Equilibrium

Important

Claim 2. In a separating equilibrium, $e(\theta_L) = 0$.

Proof.

- In a separating equilibrium, $e(\theta_L)$ reveals the worker is low ability, so:

$$w(e(\theta_L)) = \theta_L.$$

- If $e(\theta_L) > 0$, the low type receives:

$$\theta_L - c(e(\theta_L), \theta_L) < \theta_L.$$

- Thus the deviation to $e = 0$ gives at least θ_L, while choosing positive education gives less than θ_L. Therefore:

$$e(\theta_L) = 0.$$

Separating Equilibrium in the Job Market

- We now derive separating equilibria in the job market model.
- Goal: find an education level $e(\theta_H)$ for high-ability workers so that:
    - High-ability workers prefer choosing $e(\theta_H)$ and earning wage θ_H rather than choosing $e(\theta_L) = 0$ and earning $w(0) = \theta_L$. Formally:

$$\theta_H - c(e(\theta_H), \theta_H) \geq \theta_L - c(0, \theta_H). \tag{1}$$

    - Low-ability workers prefer the opposite:

$$\theta_L - c(0, \theta_L) \geq \theta_H - c(e(\theta_H), \theta_L). \tag{2}$$

Separating Equilibrium – Claim 3

Important

Claim 3. For any $e(\theta_H)$ that satisfies (1) and (2), there is a separating equilibrium in which high-ability workers choose $e(\theta_H)$.

Proof.

- If (1) and (2) hold, then no worker will want to mimic a worker of the other type.
- The only concern is that some worker might choose some $e \notin \{e(\theta_L), e(\theta_H)\}$.
- To prevent this, choose $\pi(e)$ so low that for all $e \notin \{e(\theta_L), e(\theta_H)\}$:

$$w(e) - c(e, \theta_H) \leq \theta_H - c(e(\theta_H), \theta_H), \quad w(e) - c(e, \theta_L) \leq \theta_L - c(0, \theta_L).$$

Separating Equilibrium (Proof cont.)

- One simple choice is to set $\pi(e) = 0$ for all $e \notin \{e(\theta_L), e(\theta_H)\}$.
- This implies $w(e) = \theta_L$ for all off-path e.
- Hence, choosing an off-path education level cannot increase any worker's utility.

- **Thus, a separating equilibrium exists under** (1) **and** (2)**.** Q.E.D.

Incentive Constraints: Low-Ability Type

- To prevent mimicking, we require:

$$\theta_L - c(0, \theta_L) \geq \theta_H - c(e(\theta_H), \theta_L).$$

- Assuming $c(0, \theta_L) = 0$, this simplifies to:

$$c(e(\theta_H), \theta_L) \geq \theta_H - \theta_L.$$

- Hence the **minimum** level of education $\underline{e}$ satisfies:

$$c(\underline{e}, \theta_L) = \theta_H - \theta_L.$$

Incentive Constraints: High-Ability Type

- To ensure the high-ability type is willing to choose $e(\theta_H)$:

$$\theta_H - c(e(\theta_H), \theta_H) \geq \theta_L - c(0, \theta_H).$$

- Assuming $c(0, \theta_H) = 0$, this simplifies to:

$$c(e(\theta_H), \theta_H) \leq \theta_H - \theta_L.$$

- Hence the **maximum** level of education $\bar{e}$ satisfies:

$$c(\bar{e}, \theta_H) = \theta_H - \theta_L.$$

- Combining the two:

$$\underline{e} \leq e(\theta_H) \leq \bar{e}.$$

Range of $e(\theta_H)$

- Hence, there is a separating equilibrium for any $e(\theta_H) \in [\underline{e}, \bar{e}]$.
- No separating equilibrium exists for any $e(\theta_H)$ outside this interval.

> **Note**
>
> **Remark 2.** Note the important role played by the single-crossing property. This is what allows high-ability workers to choose a positive education level that is costly, but less costly for them than it would be for low-ability workers. It is this **differential cost** that allows separation.

Pooling Equilibrium

- In a **pooling equilibrium**, every worker – regardless of ability – chooses the **same education level** e^P with probability one.
- The labor market cannot distinguish between types, so its **beliefs** must satisfy:

$$\pi(e^P) = \lambda,$$

  where λ is the proportion of high-ability workers.
- Hence, the wage offered is the **average productivity**:

$$w(e^P) = \lambda\theta_H + (1-\lambda)\theta_L \equiv \theta^E.$$

Pooling Equilibrium (Definition of $\hat{e}$)

- Let $\hat{e}$ solve:

$$\theta^E - c(\hat{e}, \theta_L) = \theta_L - c(0, \theta_L).$$

$\Rightarrow$ $\hat{e}$ is the level of education where a low-ability worker is indifferent between choosing $\hat{e}$ and earning θ^E, or choosing 0 and earning θ_L.

Important

Proposition 1. For any $e^P \in [0, \hat{e}]$, there is a pooling equilibrium where all workers choose e^P for sure.

Pooling Equilibrium: Proof of Proposition 1

Proof.

- Let $e^P \in [0, \hat{e}]$ be given.
- Suppose that the firm's belief on the equilibrium path is $\pi(\theta_H \mid e^P) = \lambda$, and off the path is $\pi(\theta_H \mid e) = 0$ for all $e \neq e^P$, with wages given accordingly.
- Then $w(e) < w(e^P)$ for any $e \neq e^P$, so clearly no worker wants to deviate to $e > e^P$.
- Moreover, low-ability workers prefer e^P to any $e < e^P$ by definition of $\hat{e}$.
- And since low-ability workers prefer e^P to any lower e, so must high-ability workers (by the single-crossing property).

Q.E.D.

Discussion and Implications

- In a **separating equilibrium**:
  - Education does **not increase productivity**.
  - It **reveals information** about workers' types.
  - Hence it is correlated with wages.
- In a **pooling equilibrium**:
  - Education neither increases productivity nor reveals information.
  - Workers may still incur education costs to avoid the wage penalty for choosing an unexpected action.
  - Pooling with positive education is therefore very **inefficient**.

Efficiency and Model Extensions

- A key feature of the model:
  - Education is **not productive at all**.
- One can generalize the model so that productivity depends on both education and ability, e.g.:

$$\text{productivity} = \theta + e.$$

- Even then, the key idea holds:
  - Signalling causes workers to choose higher education than is **efficient**.

Required Reading

- Yildiz, M. *14.126 Game Theory*, MIT OpenCourseWare – lecture notes on perfect Bayesian and sequential equilibrium.
- Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press – sequential rationality and equilibrium refinements.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press – Chapter 13 (job-market signalling).
- Spence, A. M. (1973). "Job Market Signaling," *Quarterly Journal of Economics*, 87(3), 355–374.