# Advanced Microeconomics

## Lecture 5: Incomplete Information, Bayesian Games, and Signaling

Yihu Hou

Renmin University of China

Spring 2026

# Outline

1. Why incomplete information is a different strategic friction
2. Bayesian games and Bayesian Nash equilibrium
3. Signaling, PBE, and the role of off-path beliefs
4. Screening as the mirror image and bridge to mechanism design

# From Chapter 3 to Chapter 4

## Chapter 3: complete information

Everyone knows the game being played:

- action sets are known,
- payoff functions are known,
- strategic difficulty comes from **interdependence of actions**.

## Chapter 4: incomplete information

Players may also know something **private**:

- a seller knows quality,
- a worker knows ability,
- a bidder knows valuation,
- an insured driver knows risk or effort.

## Typical question

"What is my best action given what you may do?"

## New question

"What do I do given my information, and what do I infer from yours?"

## The new object

The game now depends not only on **actions**, but also on **types**, **beliefs**, and how information is revealed.

# Two benchmark frictions: hidden type vs. hidden action

## Adverse selection

Private information is about a player's **type**:

- quality,
- productivity,
- risk,
- cost.

The hidden information is already present **before** contracting or trade.

## Typical consequence

The wrong people trade, the wrong contracts are signed, or good types leave the market.

## Moral hazard

The private information is about a player's **action**:

- effort,
- care,
- maintenance,
- risk-taking.

The action becomes hidden **after** the contract is in place.

## Typical consequence

The contract is signed, but behavior becomes distorted because incentives are misaligned.

# Timing matters: where does the information problem enter?

## Adverse selection timeline

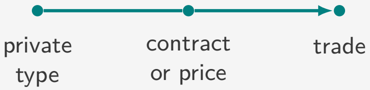


- used-car quality before sale,
- insurance risk before buying coverage,
- borrower type before lending.

## Moral hazard timeline

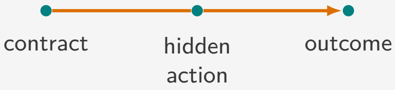


- driver effort after insurance,
- worker effort after employment,
- borrower behavior after receiving funds.

## Why the distinction matters

Hidden types primarily raise problems of **beliefs, signaling, and sorting**. Hidden actions primarily raise problems of **incentive provision, monitoring, and contingent pay**.

# Akerlof's lemons logic: information can destroy trade

## A simple used-car example

There are two qualities:

$$\text{good car: value } 10, \text{ seller reservation } 8$$

$$\text{lemon: value } 4, \text{ seller reservation } 3$$

Half the cars are good and half are lemons.

## If quality is observable

- good cars trade at a price between 8 and 10,
- lemons trade at a price between 3 and 4.

## If buyers only know the average

Expected value to a buyer:

$$\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 4 = 7.$$

So the most buyers are willing to pay is 7.

## Unraveling

- good-car owners reject $7 < 8$ and leave,
- now only lemons remain,
- buyers update and pay at most 4,
- the market shrinks to "lemons only".

**Core mechanism:** Adverse selection changes market **composition**: price affects who trades, not just how surplus is split.

# Three broad responses to information asymmetry

### Signaling

The informed side moves first and sends a costly or otherwise informative signal.
Examples:

- education,
- warranties,
- certification.

### Screening

The uninformed side designs a menu so that different types sort themselves.
Examples:

- insurance menus,
- product versioning,
- nonlinear pricing.

### Performance incentives

When the hidden object is an action, the contract must make the desired behavior privately optimal, not just sort types.
Examples:

- bonuses,
- deductibles,
- performance pay.

### Roadmap for today

We focus on hidden types: first with **Bayesian games**, then with **signaling**, and finally a short bridge to **screening**.

# What a Bayesian game adds to normal form

## Normal-form game

- players $i = 1, \ldots, n$,
- actions $A_i$,
- payoffs $u_i(a_1, \ldots, a_n)$.

## Strategic issue

What do others **do**?

## Bayesian game

- players $i = 1, \ldots, n$,
- types $t_i \in T_i$,
- common prior $p(t_1, \ldots, t_n)$,
- actions $A_i$,
- payoffs $u_i(a_1, \ldots, a_n; t_1, \ldots, t_n)$.

Each player observes **own type** but not others' types.

## Strategic issue

What do others **know**, and how does that change what they do?

**Interpretation:** A type is the privately observed state that changes payoffs, actions, or both.

# Strategy in a Bayesian game: a contingent plan

## Definition

A pure strategy is a function

$$s_i : T_i \to A_i.$$

You do not choose a single action once and for all. You specify what you would do for **each type you might be**.

- If $t_i$ is valuation, the strategy is a **bid function**.
- If $t_i$ is cost, the strategy says which output or entry decision each cost type takes.
- If $t_i$ is risk, the strategy says which contract or action each risk type chooses.

## Auction example

$$t_i = v_i \in [0, 1]$$

$$s_i(v_i) = b_i$$

Different valuations submit different bids.

## Key shift

Equilibrium is now about **plans indexed by information**, not just about one-shot actions.

# Bayesian Nash equilibrium

## Definition

A strategy profile $s^* = (s_1^*, \ldots, s_n^*)$ is a **Bayesian Nash equilibrium** if, for every player $i$ and every realized type $t_i$,

$$s_i^*(t_i) \in \operatorname{argmax}_{a_i \in A_i} \mathbb{E}\left[u_i\big(a_i, s_{-i}^*(t_{-i}); t_i, t_{-i}\big) \mid t_i\right].$$

## What is being optimized?

Expected payoff given:

- your realized type,
- beliefs about others' types, formed from the prior and your own realized type,
- others' equilibrium strategies.

## What is the logic?

For each type $t_i$, the prescribed action must be a best response to the **behavior induced by other types and strategies**.

## Mutual best response survives

BNE is mutual best response evaluated **type by type**.

# First-price sealed-bid auction as a Bayesian game

## Environment

- two bidders,
- valuation $v_i \sim U[0,1]$ independently,
- bidder $i$ knows only own $v_i$,
- highest bid wins and pays own bid.

## Payoff

If bidder $i$ bids $b_i$ and wins,

$$u_i = v_i - b_i.$$

If bidder $i$ loses,

$$u_i = 0.$$

## Why this is a Bayesian game

- **type**: valuation $v_i$,
- **action**: bid $b_i$,
- **strategy**: bid function $b(v_i)$,
- **belief**: the other bidder's valuation is drawn from the prior.

## Economic tradeoff

Bid higher to win more often, or bid lower to keep more surplus if you win.

# Solving the auction for $n = 2$: expected payoff

Assume a symmetric linear candidate strategy:

$$b(v) = \beta v, \qquad 0 < \beta < 1.$$

Consider bidder 1 with valuation $v$. If bidder 1 deviates and submits some bid $b$, then

$$U(b \mid v) = (v - b)\,\Pr(\text{win with bid } b).$$

## Winning probability

Bidder 1 wins if

$$b > \beta v_2.$$

Since $v_2 \sim U[0, 1]$,

$$\Pr(\text{win}) = \Pr\left(v_2 < \frac{b}{\beta}\right) = \frac{b}{\beta},$$

as long as $b \leq \beta$.

## Expected payoff

$$U(b \mid v) = (v - b)\frac{b}{\beta}.$$

So the type-$v$ bidder solves

$$\max_b \ (v - b)\frac{b}{\beta}.$$

# Solving the auction for $n = 2$: equilibrium bid shading

## First-order condition

$$\frac{\partial U(b \mid v)}{\partial b} = \frac{1}{\beta}(v - 2b) = 0 \quad \Rightarrow \quad b^*(v) = \frac{v}{2}.$$

So the symmetric equilibrium coefficient is

$$\beta^* = \frac{1}{2}.$$

## Equilibrium strategy

$$b^*(v) = \frac{1}{2}v.$$

## How to read it

- a bidder with higher valuation bids more,
- but still bids below value,
- the gap $v - b^*(v) = v/2$ is bid shading,
- more generally, with $n$ bidders and $v_i \sim U[0, 1]$, $b_n^*(v) = \frac{n-1}{n}v$, so bids move closer to value as $n$ rises.

## Why not bid your value?

In a first-price auction, raising your bid helps you win, but every extra dollar also raises the amount you must pay if you do win.

# What BNE can do, and where it stops

## BNE is the right tool when

- moves are simultaneous,
- types are private,
- the strategic problem is “type to action”.

## BNE is not enough when

- actions are sequential,
- one player observes a signal,
- later moves depend on updated beliefs.

## Examples

- sealed-bid auctions,
- static quantity choice with unknown cost,
- simultaneous contract choices under private risk.

## Then we need more

We must specify not only strategies, but also what players believe **after every possible history**, including surprising ones.

## Bridge

Static incomplete information leads to **Bayesian Nash equilibrium**. Dynamic incomplete information leads to **Perfect Bayesian equilibrium**.

# Why signaling is different from a static Bayesian game

## Signaling structure

1. Nature draws the sender's type.
2. The sender observes the type and chooses a signal.
3. The receiver observes the signal and chooses an action.

## New issue

The receiver's action depends on beliefs *after seeing the signal*.

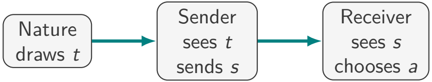


## Canonical examples

- education and wages,
- warranties and product quality,
- certification and hidden quality.

# Why SPE is not enough either

## Recall Lecture 4

In complete-information sequential games, SPE ruled out non-credible threats by requiring optimal play in every subgame.

## What breaks here

In a signaling game, after observing the signal the receiver often sits at an **information set**: the receiver sees the signal but not the sender's true type.

## So what is missing?

There may be no proper subgame that starts there, because the receiver cannot tell which node within the information set has been reached.

## Needed refinement

We need a solution concept that combines:

- strategies,
- beliefs at information sets,
- sequential rationality given those beliefs.

## Name

That solution concept is **Perfect Bayesian equilibrium (PBE)**.

# Spence-style education model: setup

## Environment

- worker type $t \in \{H, L\}$ with prior $\Pr(H) = \Pr(L) = 1/2$,
- productivity is 10 for $H$ and 0 for $L$,
- worker chooses education level $s \geq 0$,
- employer observes $s$ and pays wage equal to expected productivity:

$$w(s) = \mathbb{E}[\text{productivity} \mid s].$$

## Signaling cost

$$c_H(s) = s, \qquad c_L(s) = 2s.$$

The same signal is more expensive for the low type.

## Worker payoff

$$u_t(s) = w(s) - c_t(s).$$

## Source of separation

Education has signaling power here not because it directly raises productivity, but because it has **different costs across types**.

# Perfect Bayesian equilibrium in one slide

## Objects

A PBE consists of:

- a sender strategy,
- a receiver strategy,
- a belief system $\mu(t \mid s)$ at every information set.

## Conditions

1. **Sequential rationality**: each player's continuation action is optimal given beliefs.
2. **Bayes on path**: beliefs are updated by Bayes' rule whenever the signal occurs with positive probability.
3. **Off path**: beliefs must still be specified; refinements may restrict them further.

## Mental model

BNE tells us how to optimize given private types. PBE adds the extra layer needed when later actions depend on **inferences from earlier signals**.

# Separating equilibrium: a whole interval can be supported

Candidate separating PBE:

$$s_L^* = 0, \qquad s_H^* = s^* > 0, \qquad w(0) = 0, \quad w(s^*) = 10.$$

## High type must prefer to separate

Following the signal must beat mimicking $L$:

$$10 - s^* \geq 0 \quad \Rightarrow \quad s^* \leq 10.$$

## Low type must not mimic

Staying at 0 must beat imitating $H$:

$$0 \geq 10 - 2s^* \quad \Rightarrow \quad s^* \geq 5.$$

## Implication

Before refinement, every $s^* \in [5, 10]$ is a separating PBE.

# Pooling equilibria: an interval can be supported

## Candidate pooling PBE

Take any common signal

$$s_H^* = s_L^* = s^*.$$

On path, the employer pays $w(s^*) = 5$, so the payoffs are $u_H(s^*) = 5 - s^*$ and $u_L(s^*) = 5 - 2s^*$.

## Necessary condition

Since the low type can always fall back to zero education, we need

$$5 - 2s^* \geq 0 \quad \Rightarrow \quad s^* \leq 2.5.$$

## Implication

Before refinement, the pooling candidates are exactly the signals $s^* \in [0, 2.5]$. Whether they are PBE depends on what beliefs support deviations off path.

# A simple off-path belief supporting pooling at 0

## Focal pooling candidate

Consider the focal candidate $s^* = 0$. On path, the employer pays $w(0) = 5$, so both types get payoff 5.

## Simple off-path belief

Let the employer keep the prior after any unexpected signal:

$$\mu(H \mid s) = \frac{1}{2} \qquad \text{for all } s \neq 0.$$

Then every off-path signal gets wage $w(s) = 5$. A deviation to any $s > 0$ therefore gives the high type payoff $5 - s < 5$ and the low type payoff $5 - 2s < 5$.

## Message

Pooling at 0 is supported with the off-path belief. Other pooling signals in $[0, 2.5]$ can also be supported under suitable off-path beliefs.

# On-path vs. off-path beliefs

## On path

A signal is **on path** if it occurs with positive probability in equilibrium.
Then Bayes' rule tells us how to update:

$$\mu(H \mid s) = \frac{\Pr(s \mid H)\Pr(H)}{\Pr(s \mid H)\Pr(H) + \Pr(s \mid L)\Pr(L)}.$$

## Off path

A signal is **off path** if equilibrium assigns it zero probability.
Then Bayes' rule is not pinned down, because the denominator is zero.

## But the receiver still must act

Even after an off-path signal, the receiver chooses an action, so the equilibrium must still specify beliefs there.

## Why refinements appear

If off-path beliefs are unrestricted, they can artificially support too many equilibria. Refinements restrict which beliefs we should regard as plausible.

# The intuitive criterion: the core logic

## Question to ask

Take an off-path signal $s'$.
Who could plausibly benefit from sending it if the receiver were as favorable as possible to the deviator?

## Best possible belief

Suppose the receiver interpreted $s'$ in the way most favorable to the deviating type.

## If only one type benefits

If type $H$ would gain from $s'$ under that favorable belief, but type $L$ would still not gain, then only $H$ has a reason to use that deviation.

## Criterion

The receiver should then put probability 1 on the types that could plausibly gain from the deviation. Any equilibrium sustained by the opposite belief is considered unreasonable.

# Applying the intuitive criterion: most separating equilibria disappear

Consider the separating PBE with

$$s_L^* = 0, \qquad s_H^* = 6.$$

Then the high type gets payoff $10 - 6 = 4$.

## Deviation to $s' = 5.1$

If the employer interpreted $s' = 5.1$ as coming from $H$, then

$$u_H(5.1) = 10 - 5.1 = 4.9 > 4.$$

So the high type would like to deviate.

## Would the low type also want this?

Even under the same favorable belief,

$$u_L(5.1) = 10 - 2 \cdot 5.1 = -0.2 < 0.$$

So the low type would *not* want the same deviation.

## Conclusion

The deviation $s' = 5.1$ is credible only for the high type. The intuitive criterion therefore rules out the equilibrium with $s_H^* = 6$. By the same logic, every separating equilibrium with $s_H^* > 5$ is eliminated.

# Applying the intuitive criterion: pooling also fails

## Start from any pooling PBE

Take any pooling signal $s^* \in [0, 2.5]$.
On path, the wage is $w(s^*) = 5$, so the payoffs are $u_H(s^*) = 5 - s^*$ and $u_L(s^*) = 5 - 2s^*$.

## A deviation only the high type likes

Consider the off-path signal $s' = s^* + 3$.
If the employer interpreted $s'$ as coming from the high type, then

$$u_H(s') = 10 - (s^* + 3) = 7 - s^* > 5 - s^* = u_H(s^*).$$

So the high type would like to deviate.

## What about the low type?

Under the same favorable belief,

$$u_L(s') = 10 - 2(s^* + 3) = 4 - 2s^* < 5 - 2s^* = u_L(s^*).$$

So only the high type would want the same deviation.

## Conclusion

For every $s^* \in [0, 2.5]$, the intuitive criterion rules out pooling and selects the **least-cost separating equilibrium** $s_H^* = 5$.

# What signaling teaches us

### 1. Differential cost

A signal separates only if it is cheaper for the high type than for the low type.

### 2. Beliefs matter

Pooling and multiple separating equilibria often exist because off-path beliefs can support them.

### 3. Refinements matter

The intuitive criterion asks whether the off-path belief is economically plausible, not just mathematically admissible.

### Big picture

Signaling turns an adverse-selection problem into a dynamic game in which beliefs, not just optimization, determine the equilibrium.

# Screening is the mirror image of signaling

## Signaling

The **informed side** moves first.

- worker chooses education,
- seller offers warranty,
- firm acquires certification.

Goal: convince the uninformed side about your type.

## Screening

The **uninformed side** designs a menu.

- insurer offers contracts,
- monopolist offers versions,
- employer designs compensation schemes.

Goal: induce different types to sort themselves.

## Common objective

Both signaling and screening try to solve the same adverse-selection problem: **how can private types be revealed or sorted without being directly observed?**

# A simple screening example: insurance menu

## Environment

Suppose there are two risk types:

- low-risk consumers,
- high-risk consumers.

The insurer cannot observe type directly.

## A screening menu

| Contract | Premium | Deductible |
|---|---|---|
| A | low | high |
| B | high | low |

## Desired sorting

- low-risk consumers choose cheap coverage with more risk retained,
- high-risk consumers choose expensive coverage with more protection.

## Design problem

The menu must satisfy:

- **participation**: each type prefers joining to opting out,
- **incentive compatibility**: each type prefers the contract intended for it.

# Why screening leads naturally into mechanism design

## What we learned today

- private information changes the game,
- strategies become type-contingent,
- beliefs matter in dynamic environments,
- sorting can come from signals or menus.

## What comes next

Mechanism design asks:

- what rules or menus should the designer choose,
- how do IC and IR constraints shape the optimal mechanism,
- how are auctions and screening just two versions of the same problem?

## Bridge to the next two lectures

Today we learned the language of **types, beliefs, signals, and self-selection**. Next we put the designer at center stage.

# End of Lecture 5

- Incomplete information means strategic interaction depends on **what players privately know**, not only on what they choose.
- Bayesian games model simultaneous hidden-type problems through type-contingent strategies and Bayesian Nash equilibrium.
- Signaling games require beliefs at information sets, so the natural solution concept is PBE.
- Off-path beliefs can create many candidate equilibria, and refinements such as the intuitive criterion help select the plausible ones.
- Screening is the mirror image of signaling and leads directly into mechanism design.

**Core distinctions to remember**

$$\text{adverse selection} \neq \text{moral hazard}, \qquad \text{BNE} \neq \text{PBE},$$

$$\text{signaling by the informed side} \neq \text{screening by the uninformed side}.$$

# Next lecture

## Where we are going

We now switch from "how private information shapes behavior" to "how a designer chooses rules when information is private."

- auction formats and strategic bidding,
- direct mechanisms and the revelation principle,
- incentive compatibility and individual rationality,
- revenue equivalence and the logic behind VCG.

## Bridge

Lecture 5 treated private information as part of the environment.
Lecture 6 asks how to **design the environment** so that private information is used well rather than wasted.

# References

## Selected references for Lecture 5


- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, esp. Chs. 13–14.
- Jehle, G. A., and Reny, P. J. (2011). *Advanced Microeconomic Theory*. 3rd ed. Financial Times/Prentice Hall, esp. Chs. 9–10.
- Akerlof, G. A. (1970). "The Market for Lemons." *Quarterly Journal of Economics*.
- Spence, M. (1973). "Job Market Signaling." *Quarterly Journal of Economics*.
- Kreps, D. M., and Wilson, R. (1982). "Sequential Equilibria." *Econometrica*.
- Cho, I.-K., and Kreps, D. M. (1987). "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics*.