

**UC Berkeley
Haas School of Business
Game Theory
(EMBA 296 & EWMBA 211)
Spring 2025**

**Review of strategic and extensive games
and advanced topics**

**Block IV
Mar 12 and 16, 2025**

Game plan

1. Review of the main ideas
2. Signaling
3. Social learning
4. Auctions
5. Bank runs
6. More, if time permits...

A review of the main ideas

We study two (out of four) groups of game theoretic models:

- [1] Strategic games – all players simultaneously choose their plan of action once and for all.
- [2] Extensive games (with perfect information) – players choose sequentially (and fully informed about all previous actions).

A solution (equilibrium) is a systematic description of the outcomes that may emerge in a family of games. We study two solution concepts:


[1] Nash equilibrium – a steady state of the play of a strategic game (no player has a profitable deviation given the actions of the other players).


[1] Subgame equilibrium – a steady state of the play of an extensive game (a Nash equilibrium in every subgame of the extensive game).

\implies Every subgame perfect equilibrium is also a Nash equilibrium.

Extensive games with perfect information


- The model of a strategic suppresses the sequential structure of decision making.
 - All players simultaneously choose their plan of action once and for all.
- The model of an extensive game, by contrast, describes the sequential structure of decision-making explicitly.
 - In an extensive game of perfect information all players are fully informed about all previous actions.






Subgame perfect equilibrium

- The notion of Nash equilibrium ignores the sequential structure of the game.
- Consequently, the steady state to which a Nash Equilibrium corresponds may not be robust.
- A *subgame perfect equilibrium* is an action profile that induces a Nash equilibrium in every *subgame* (so every subgame perfect equilibrium is also a Nash equilibrium).



"LORETTA'S DRIVING BECAUSE I'M DRINKING,
AND I'M DRINKING BECAUSE SHE'S DRIVING."



Games with imperfect (and asymmetric) information

Markets with asymmetric information

- The traditional theory of markets assumes that market participants have complete information about the underlying economic variables:
 - Buyers and sellers are both perfectly informed about the quality of the goods being sold in the market.
 - If it is not costly to verify quality, then the prices of the goods will simply adjust to reflect the quality difference.

⇒ This is clearly a drastic simplification!!!

- There are certainly many markets in the real world in which it may be very costly (or even impossible) to gain accurate information:
 - labor markets, financial markets, markets for consumer products, and more.
- If information about quality is costly to obtain, then it is no longer possible that buyers and sellers have the same information.
- The costs of information provide an important source of market friction and can lead to a market breakdown.

Nobel Prize 2001
“for their analyses of markets with asymmetric information”



The Market for Lemons

Example I

- Consider a market with 100 people who want to sell their used car and 100 people who want to buy a used car.
- Everyone knows that 50 of the cars are “plums” and 50 are “lemons.”
- Suppose further that

	seller	buyer
lemon	\$1000	\$1200
plum	\$2000	\$2400

- If it is easy to verify the quality of the cars there will be no problem in this market.
- Lemons will sell at some price \$1000 – 1200 and plums will sell at \$2000 – 2400.
- But happens to the market if buyers cannot observe the quality of the car?

- If buyers are risk neutral, then a typical buyer will be willing to pay his expected value of the car

$$\frac{1}{2}1200 + \frac{1}{2}2400 = \$1800.$$

- But for this price only owners of lemons would offer their car for sale, and buyers would therefore (correctly) expect to get a lemon.
- Market failure – no transactions will take place, although there are possible gains from trade!

Example II

- Suppose we can index the quality of a used car by some number q , which is distributed uniformly over $[0, 1]$.
- There is a large number of demanders for used cars who are willing to pay $\frac{3}{2}q$ for a car of quality q .
- There is a large number of sellers who are willing to sell a car of quality q for a price of q .

- If quality is perfectly observable, each used car of quality q would be soled for some price between q and $\frac{3}{2}q$.
- What will be the equilibrium price(s) in this market when quality of any given car cannot be observed?
- The unique equilibrium price is zero, and at this price the demand is zero and supply is zero.

⇒ The asymmetry of information has destroyed the market for used cars. But the story does not end here!!!

Signaling

- In the used-car market, owners of the good used cars have an incentive to try to convey the fact that they have a good car to the potential purchasers.
- Put differently, they would like choose actions that signal that they are offering a plum rather than a lemon.
- In some case, the presence of a “signal” allows the market to function more effectively than it would otherwise.

Example – educational signaling

- Suppose that a fraction $0 < b < 1$ of workers are *competent* and a fraction $1 - b$ are *incompetent*.
- The competent workers have marginal product of a_2 and the incompetent have marginal product of $a_1 < a_2$.
- For simplicity we assume a competitive labor market and a linear production function

$$L_1 a_1 + L_2 a_2$$

where L_1 and L_2 is the number of incompetent and competent workers, respectively.

- If worker quality is observable, then firm would just offer wages

$$w_1 = a_1 \text{ and } w_2 = a_2$$

to competent workers, respectively.

- That is, each worker will be paid his marginal product and we would have an efficient equilibrium.
- But what if the firm cannot observe the marginal products so it cannot distinguish the two types of workers?

- If worker quality is unobservable, then the “best” the firm can do is to offer the average wage

$$w = (1 - b)a_1 + ba_2.$$

- If both types of workers agree to work at this wage, then there is no problem with adverse selection (more below).
- The incompetent (resp. competent) workers are getting paid more (resp. less) than their marginal product.

- The competent workers would like a way to signal that they are more productive than the others.
- Suppose now that there is some signal that the workers can acquire that will distinguish the two types
- One nice example is education – it is cheaper for the competent workers to acquire education than the incompetent workers.

- To be explicit, suppose that the cost (dollar costs, opportunity costs, costs of the effort, etc.) to acquiring e years of education is

$$c_1 e \text{ and } c_2 e$$

for incompetent and competent workers, respectively, where $c_1 > c_2$.

- Suppose that workers conjecture that firms will pay a wage $s(e)$ where s is some increasing function of e .
- Although education has no effect on productivity (MBA?), firms may still find it profitable to base wage on education – attract a higher-quality work force.

Market equilibrium

In the educational signaling example, there appear to be several possibilities for equilibrium:

- [1] The (representative) firm offers a single contract that attracts both types of workers.
- [2] The (representative) firm offers a single contract that attracts only one type of workers.
- [3] The (representative) firm offers two contracts, one for each type of workers.

- A separating equilibrium involves each type of worker making a choice that separate himself from the other type.
- In a pooling equilibrium, in contrast, each type of workers makes the same choice, and all getting paid the wage based on their average ability.

Note that a separating equilibrium is wasteful in a social sense – no social gains from education since it does not change productivity.

Example (cont.)

- Let e_1 and e_2 be the education level actually chosen by the workers.
Then, a separating (signaling) equilibrium has to satisfy:

[1] zero-profit conditions

$$s(e_1) = a_1$$

$$s(e_2) = a_2$$

[2] self-selection conditions

$$s(e_1) - c_1 e_1 \geq s(e_2) - c_1 e_2$$

$$s(e_2) - c_2 e_2 \geq s(e_1) - c_2 e_1$$

- In general, there may be many functions $s(e)$ that satisfy conditions [1] and [2]. One wage profile consistent with separating equilibrium is

$$s(e) = \begin{cases} a_2 & \text{if } e > e^* \\ a_1 & \text{if } e \leq e^* \end{cases}$$

and

$$\frac{a_2 - a_1}{c_2} > e^* > \frac{a_2 - a_1}{c_1}$$

\implies Signaling can make things better or worse – each case has to be examined on its own merits!

The Sheepskin (diploma) effect

The increase in wages associated with obtaining a higher credential:

- Graduating high school increases earnings by 5 to 6 times as much as does completing a year in high school that does not result in graduation.
- The same discontinuous jump occurs for people who graduate from college.
- High school graduates produce essentially the same amount of output as non-graduates.

Social learning
herd behavior and informational cascades

“Men nearly always follow the tracks made by others and proceed in their affairs by imitation.” Machiavelli (Renaissance philosopher)

Examples

Business strategy

- TV networks make introductions in the same categories as their rivals.

Finance

- The withdrawal behavior of small number of depositors starts a bank run.

Politics

- The solid New Hampshirites (probably) can not be too far wrong.

Crime

- In NYC, individuals are more likely to commit crimes when those around them do.

Why should individuals behave in this way?


Several “theories” explain the existence of uniform social behavior:

- benefits from conformity
- sanctions imposed on deviants
- network / payoff externalities
- social learning

Broad definition: any situation in which individuals learn by observing the behavior of others.

The canonical model of social learning

- Rational (Bayesian) behavior
- Incomplete and asymmetric information
- Pure information externality
- Once-in-a-lifetime decisions
- Exogenous sequencing
- Perfect information / complete history



Bayes' rule

Let n be the number of a signals and m be the number of b signals. Then Bayes' rule can be used to calculate the posterior probability of urn A :

$$\begin{aligned}\Pr(A | n, m) &= \frac{\Pr(A) \Pr(n, m | A)}{\Pr(A) \Pr(n, m | A) + \Pr(B) \Pr(n, m | B)} \\ &= \frac{(\frac{1}{2})(\frac{2}{3})^n(\frac{1}{3})^m}{(\frac{1}{2})(\frac{2}{3})^n(\frac{1}{3})^m + (\frac{1}{2})(\frac{1}{3})^m(\frac{2}{3})^n} \\ &= \frac{2^n}{2^n + 2^m}.\end{aligned}$$

An example

- There are two decision-relevant events, say A and B , equally likely to occur *ex ante* and two corresponding signals a and b .
- Signals are informative in the sense that there is a probability higher than $1/2$ that a signal matches the label of the realized event.
- The decision to be made is a prediction of which of the events takes place, basing the forecast on a private signal and the history of past decisions.

- Whenever two consecutive decisions coincide, say both predict A , the subsequent player should also choose A even if his signal is different b .
- Despite the asymmetry of private information, eventually every player imitates her predecessor.
- Since actions aggregate information poorly, despite the available information, such herds / cascades often adopt a suboptimal action.

A model of social learning

Signals

- Each player $n \in \{1, \dots, N\}$ receives a signal θ_n that is private information.
- For simplicity, $\{\theta_n\}$ are independent and uniformly distributed on $[-1, 1]$.

Actions

- Sequentially, each player n has to make a binary irreversible decision $x_n \in \{0, 1\}$.

Payoffs

- $x = 1$ is profitable if and only if $\sum_{n \leq N} \theta_n \geq 0$, and $x = 0$ is profitable otherwise.

Information


- Perfect information

$$\mathcal{I}_n = \{\theta_n, (x_1, x_2, \dots, x_{n-1})\}$$


- Imperfect information

$$\mathcal{I}_n = \{\theta_n, x_{n-1}\}$$


Sequential social-learning model:
Well heck, if all you smart cookies agree, who am I to dissent?




Imperfect information:
Which way is the wind blowing?!




A three-agent example




A three-agent example




A three-agent example under perfect information




A three-agent example under imperfect information



A sequence of cutoffs under imperfect and perfect information



A sequence of cutoffs under imperfect and perfect information



- The dynamics of social learning depend crucially on the extensive form of the game.
- The key economic phenomenon that imperfect information captures is a succession of fads starting suddenly, expiring rather easily, each replaced by another fad.
- The kind of episodic instability that is characteristic of socioeconomic behavior in the real world makes more sense in the imperfect-information model.

As such, the imperfect-information model gives insight into phenomena such as manias, fashions, crashes and booms, and better answers such questions as:

- Why do markets move from boom to crash without settling down?
- Why is a technology adopted by a wide range of users more rapidly than expected and then, suddenly, replaced by an alternative?
- What makes a restaurant fashionable over night and equally unexpectedly unfashionable, while another becomes the ‘in place’, and so on?

Auctions

Auction design

Two important issues for auction design are:

- Attracting entry
- Preventing collusion

Sealed-bid auction deals better with these issues, but it is more likely to lead to inefficient outcomes.

European 3G mobile telecommunication license auctions

Although the blocks of spectrum sold were very similar across countries, there was an enormous variation in revenues (in USD) per capita:

Austria	100
Belgium	45
Denmark	95
Germany	615
Greece	45
Italy	240
Netherlands	170
Switzerland	20
United Kingdom	650

United Kingdom

- 4 licenses to be auctioned off and 4 incumbents (with advantages in terms of costs and brand).
- To attract entry and deter collusion – an English until 5 bidders remain and then a sealed-bid with reserve price given by lowest bid in the English.
- later a 5th license became available to auction, a straightforward English auction was implemented.

Netherlands

- Followed UK and used (only) an English auction, but they had 5 incumbents and 5 licenses!
- Low participation: strongest potential entrants made deals with incumbents, and weak entrants either stayed out or quit bidding.

Switzerland

- Also followed UK and ran an English auction for 4 licenses. Companies either stayed out or quit bidding.
 1. The government permitted last-minute joint-bidding agreements. Demand shrank from 9 to 4 bidders one week before the auction.
 2. The reserve price had been set too low. The government tried to change the rules but was opposed by remaining bidders and legally obliged to stick to the original rules.
- Collected 1/30 per capita of UK, and 1/50 of what they had hoped for!

Common-value auctions and the winner's curse

Suppose we all participate in a sealed-bid auction for a jar of coins. Once you have estimated the amount of money in the jar, what are your bidding strategies in first- and second-price auctions?

The winning bidder is likely to be the bidder with the largest positive error (the largest overestimate).

In this case, the winner has fallen prey to the so-called the winner's curse. Auctions where the winner's curse is significant are oil fields, spectrum auctions, pay per click, and more.

The winner's curse has also been shown in stock market and real estate investments, mergers and acquisitions, and bidding on baseball players.

When Goggle launched its IPO by auction in 2004, the SEC registration statement said:

“The auction process for our public offering may result in a phenomenon known as the ‘winner’s curse,’ and, as a result, investors may experience significant losses (...) Successful bidders may conclude that they paid too much for our shares and could seek to immediately sell their shares to limit their losses.”

First-price auction (with perfect information)

To define the game precisely, denote by v_i the value that bidder i attaches to the object. If she obtains the object at price p then her payoff is $v_i - p$.

Assume that bidders' valuations are all different and all positive. Number the bidders 1 through n in such a way that

$$v_1 > v_2 > \cdots > v_n > 0.$$

Each bidder i submits a (sealed) bid b_i . If bidder i obtains the object, she receives a payoff $v_i - b_i$. Otherwise, her payoff is zero.

Tie-breaking – if two or more bidders are in a tie for the highest bid, the winner is the bidder with the highest valuation.

In summary, a first-price sealed-bid auction with perfect information is the following strategic game:

- Players: the n bidders.
- Actions: the set of possible bids b_i of each player i (nonnegative numbers).
- Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \bar{b} & \text{if } b_i = \bar{b} \text{ and } v_i > v_j \text{ if } b_j = \bar{b} \\ 0 & \text{if } b_i < \bar{b} \end{cases}$$

where \bar{b} is the highest bid.

The set of Nash equilibria is the set of profiles (b_1, \dots, b_n) of bids with the following properties:

- [1] $v_2 \leq b_1 \leq v_1$
- [2] $b_j \leq b_1$ for all $j \neq 1$
- [3] $b_j = b_1$ for some $j \neq 1$

It is easy to verify that all these profiles are Nash equilibria. It is harder to show that there are no other equilibria. We can easily argue, however, that there is no equilibrium in which player 1 does not obtain the object.

\implies The first-price sealed-bid auction is socially efficient, but does not necessarily raise the most revenues.

Second-price auction (with perfect information)

A second-price sealed-bid auction with perfect information is the following strategic game:

- Players: the n bidders.
- Actions: the set of possible bids b_i of each player i (nonnegative numbers).
- Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \bar{b} & \text{if } b_i > \bar{b} \text{ or } b_i = \bar{b} \text{ and } v_i > v_j \text{ if } b_j = \bar{b} \\ 0 & \text{if } b_i < \bar{b} \end{cases}$$

where \bar{b} is the highest bid submitted by a player other than i .

First note that for any player i the bid $b_i = v_i$ is a (weakly) dominant action (a “truthful” bid), in contrast to the first-price auction.

The second-price auction has many equilibria, but the equilibrium $b_i = v_i$ for all i is distinguished by the fact that every player’s action dominates all other actions.

Another equilibrium in which player $j \neq 1$ obtains the good is that in which

- [1] $b_1 < v_j$ and $b_j > v_1$
- [2] $b_i = 0$ for all $i \neq \{1, j\}$

Bank runs
The Diamond-Dybvig (1983) model

Nobel Prize 2022



Bernanke, Diamond and Dybvig, have significantly improved our understanding of the role of banks in the economy, particularly during financial crises. An important finding in their research is why avoiding bank collapses is vital...

A simple Diamond-Dybvig (1983) model

“... bank runs are a common feature of the extreme crises that have played a prominent role in monetary history. During a bank run, depositors rush to withdraw their deposits because they expect the bank to fail...

... in fact, the sudden withdrawals can force the bank to liquidate many of its assets at a loss and to fail. In a panic with many bank failures, there is a disruption of the monetary system and a reduction in production...”

The mismatch of liquidity:

- Banks issue demand deposits that allow depositors to withdraw at any time and make loans that cannot be sold quickly even at a high price.
- Because the bank's liabilities are more liquid than its assets, it will face a problem when too many depositors attempt to withdraw at once.

⇒ A situation referred to as a bank run.

Consider the following asset on three dates, $T = \{0, 1, 2\}$:

- If one invests one unit at date 0, it will be worth r_2 at date 2, but only $r_1 < r_2$ at date 1.
- The lower r_1/r_2 is (holding constant market rates), the less liquid is the asset.

\Rightarrow The lower the fraction of the present value of the future cash flow that can be obtained today, the less liquid is the asset.

Investors have an uncertain horizon:

- Each will need to consume either at date $T = 1$ or $T = 2$ but, as of date 0, does not know at which date s/he will need to consume.
- An investor a “type 1” if s/he needs to liquidate at $T = 1$ and a “type 2” otherwise.
- Each investor has a probability t of being of type 1 and $1 - t$ of being of type 2.

⇒ There is no aggregate uncertainty (there will be a fraction t of investors of type 1).

An investor who holds the asset (r_1, r_2) , which gives a choice of r_1 at date 1 or $r_2 > r_1$ at date 2, consumes $c_1 = r_1$ if s/he is type 1 (with prob. t) or $c_2 = r_2$ if type 2 (with prob. $1 - t$).

The investor's expected utility is given by

$$tU(r_1) + (1 - t)U(r_2)$$

and we assume that the investors have the risk-averse utility function

$$U(c) = 1 - \frac{1}{c}.$$

Comparing more and less liquid assets

Consider the following two assets, both of which cost 1 at date 0:

illiquid	$(r_1 = 1, r_2 = R)$
more liquid	$(r_1 > 1, r_2 < R)$

We illustrate the demand for liquidity with the following numerical example for the case where the probability of being of type 1 is $t = \frac{1}{4}$ and

	r_1	r_2
illiquid	1	2
more liquid	1.28	1.813

(we will explain why these particular numerical values are used...).

The expected utility from holding the illiquid asset is

$$\frac{1}{4}U(1) + \frac{3}{4}U(2) = 0.375$$

where the expected utility from holding the more liquid asset is

$$\frac{1}{4}U(1.28) + \frac{3}{4}U(1.813) = 0.391 > 0.375$$

⇒ A risk-averse investor prefers holding the more liquid / less risky asset because s/he prefers its more 'smoother' pattern of returns.

But note that if investors were not risk averse $U(c) = c$, they would not prefer this particular liquid asset:

$$\frac{1}{4}U(1) + \frac{3}{4}U(2) = \frac{1}{4}(1) + \frac{3}{4}(2) = 1.75$$

where the expected utility from holding the more liquid asset is

$$\frac{1}{4}U(1.28) + \frac{3}{4}U(1.813) = \frac{1}{4}(1.28) + \frac{3}{4}(1.813) = 1.68 < 1.75.$$

⇒ An investor's demand for liquidity is greater the higher her/his (relative) risk aversion is.

The optimal amount of liquidity

1. Suppose that the bank receives \$1 from each of 100 investors at date $T = 0$ and in return offers to pay $r_1 = 1.28$ to those who withdraw at $T = 1$ or to pay $r_2 = 1.813$ to those who withdraw at $T = 2$.
2. If the bank invests in the illiquid asset, it will need to liquidate $25 \times 1.28 = 32$ assets (32% of the portfolio) at $T = 1$ to pay 1.28 to those who withdraw.
3. Then 68 assets will remain until $T = 2$, when they will be worth $R = 2$ each. The 75 depositors that will remain at $T = 2$ will receive

$$\frac{68 \times 2}{75} = 1.813$$

The optimal levels of r_1 and r_2 maximize the ex-ante expected utility of each investor at date 0. That is,

$$\max_{r_1, r_2} \quad tU(r_1) + (1 - t)U(r_2)$$

$$\text{subject to } r_1 \geq 0, r_2 \geq 0, r_2 \leq \frac{(1 - tr_1)R}{1 - t}$$

- \Rightarrow An investor needs all or none of his liquidity, while the bank knows that a fraction t of its depositors will need liquidity at date 1.
- \Rightarrow When long-term assets are even more illiquid, there is an additional way that banks can help investors and/or make profits...

Bank runs

- How much is left to pay depositors who wait until date 2 to withdraw if a fraction $f \geq t$ of initial depositors withdraw at date 1?
- Each depositor needs a forecast of f denoted by \hat{f} upon which s/he chooses whether to withdraw at date 1.
- Recall that a Nash equilibrium is self-fulfilling, and in the good equilibrium (no bank run) $\hat{f} = f = t = \frac{1}{4}$.
- But there is also a bad equilibrium (bank run) in which all withdraw at date 1 because they all expect each other to do the same...

- The self-fulfilling prophecy of a bank run is $\hat{f} = f = 1$ where all rush to withdraw. That is, if a run is feared, it becomes a self-fulfilling prophecy.
- These two possible equilibrium beliefs (self-fulfilling forecasts of f) are so-called locally stable. The tipping point for a run is a forecast implying that

$$r_1 \geq r_2(\hat{f}) = \frac{(1 - \hat{f} \times r_1)R}{1 - \hat{f}}$$

or

$$\hat{f} > \frac{R - r_1}{r_1(R - 1)}$$

which in the example is 0.5625.

Takeaways

- Moving away from a good equilibrium requires a large change in beliefs \implies a run requires 'bad news' that many depositors see (and believe that others see).
- Not all depositors observe the same news so they do not have a way to tell if others are choosing to panic and run (no common knowledge).
- One way to stop and prevent runs is deposit insurance by government b/c it has taxation authority (ability to take resources without prior contracts).
- But there is also scope for banks to write more refined contracts, e.g. deposits with suspension of convertibility of to cash.