

**UC Berkeley
Haas School of Business
Game Theory
(EMBA 296 & EWMBA 211)
Summer 2015**

Extensive form games with perfect information, imperfect (incomplete and asymmetric) information, oligopoly games, auctions, social learning and more...

**Block 3
Jul 10-11, 2015**

Food for thought

LUPI

Many players simultaneously chose an integer between 1 and 99,999. Whoever chooses the lowest unique positive integer (LUPI) wins.

Question What does an equilibrium model of behavior predict in this game?

The field version of LUPI, called Limbo, was introduced by the government-owned Swedish gambling monopoly Svenska Spel. Despite its complexity, there is a surprising degree of convergence toward equilibrium.

Morra

A two-player game in which each player simultaneously hold either one or two fingers and each guesses the total number of fingers held up.

If exactly one player guesses correctly, then the other player pays her the amount of her guess.

Question Model the situation as a strategic game and describe the equilibrium model of behavior predict in this game.

The game was played in ancient Rome, where it was known as “micatio.”

Maximal game (sealed-bid second-price auction)

Two bidders, each of whom privately observes a signal X_i that is independent and identically distributed (i.i.d.) from a uniform distribution on $[0, 10]$.

Let $X^{\max} = \max\{X_1, X_2\}$ and assume the ex-post common value to the bidders is X^{\max} .

Bidders bid in a sealed-bid second-price auction where the highest bidder wins, earns the common value X^{\max} and pays the second highest bid.

**Strategic games
(review)**

A two-player (finite) strategic game

The game can be described conveniently in a so-called bi-matrix. For example, a generic 2×2 (two players and two possible actions for each player) game

	<i>L</i>	<i>R</i>
<i>T</i>	a_1, a_2	b_1, b_2
<i>B</i>	c_1, c_2	d_1, d_2

where the two rows (resp. columns) correspond to the possible actions of player 1 (resp. 2). The two numbers in a box formed by a specific row and column are the players' payoffs given that these actions were chosen.

In this game above a_1 and a_2 are the payoffs of player 1 and player 2 respectively when player 1 is choosing strategy T and player 2 strategy L .

Best response and dominated actions

Action T is player 1's *best response* to action L player 2 if T is the optimal choice when 1 *conjectures* that 2 will play L .

Player 1's action T is *strictly* dominated if it is never a best response (inferior to B no matter what the other players do).

In the Prisoner's Dilemma, for example, action *Work* is strictly dominated by action *Goof*. As we will see, a strictly dominated action is not used in any Nash equilibrium.

Nash equilibrium

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

Put differently, a NE is a set of actions such that all players are doing their best given the actions of the other players.

Mixed strategy Nash equilibrium in the BoS

Suppose that, each player can randomize among all her strategies so choices are not deterministic:

		q	$1 - q$
		L	R
p	T	pq	$p(1 - q)$
$1 - p$	B	$(1 - p)q$	$(1 - p)(1 - q)$

Let p and q be the probabilities that player 1 and 2 respectively assign to the strategy *Ball*.

Player 2 will be indifferent between using her strategy B and S when player 1 assigns a probability p such that her expected payoffs from playing B and S are the same. That is,

$$\begin{aligned}1p + 0(1 - p) &= 0p + 2(1 - p) \\ p &= 2 - 2p \\ p^* &= 2/3\end{aligned}$$

Hence, when player 1 assigns probability $p^* = 2/3$ to her strategy B and probability $1 - p^* = 1/3$ to her strategy S , player 2 is indifferent between playing B or S any mixture of them.

Similarly, player 1 will be indifferent between using her strategy B and S when player 2 assigns a probability q such that her expected payoffs from playing B and S are the same. That is,

$$\begin{aligned}2q + 0(1 - q) &= 0q + 1(1 - q) \\2q &= 1 - q \\q^* &= 1/3\end{aligned}$$

Hence, when player 2 assigns probability $q^* = 1/3$ to her strategy B and probability $1 - q^* = 2/3$ to her strategy S , player 2 is indifferent between playing B or S any mixture of them.

In terms of best responses:

$$B_1(q) = \begin{cases} p = 1 & \text{if } p > 1/3 \\ p \in [0, 1] & \text{if } p = 1/3 \\ p = 0 & \text{if } p < 1/3 \end{cases}$$

$$B_2(p) = \begin{cases} q = 1 & \text{if } p > 2/3 \\ q \in [0, 1] & \text{if } p = 2/3 \\ q = 0 & \text{if } p < 2/3 \end{cases}$$

The *BoS* has two Nash equilibria in pure strategies $\{(B, B), (S, S)\}$ and one in mixed strategies $\{(2/3, 1/3)\}$. In fact, any game with a finite number of players and a finite number of strategies for each player has Nash equilibrium (Nash, 1950).

Three Matching Pennies games in the laboratory

		.48	.52
		a_2	b_2
.48	a_1	80, 40	40, 80
.52	b_1	40, 80	80, 40

		.16	.84			.80	.20
		a_2	b_2			a_2	b_2
.96	a_1	320, 40	40, 80	.08	a_1	44, 40	40, 80
.04	b_1	40, 80	80, 40	.92	b_1	40, 80	80, 40

Extensive games with perfect information


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).


An example: entry game




Subgame perfect and backward induction



Two entry games in the laboratory





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).

⇒ Every subgame perfect equilibrium is also a Nash equilibrium.

**Incomplete and asymmetric information
(an illustration – the market for lemons)**

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 sold 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

$$c_1e \text{ and } c_2e$$

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

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

⇒ 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 collage.
- High school graduates produce essentially the same amount of output as non-graduates.

**Oligopolistic competition
(in strategic and extensive forms)**

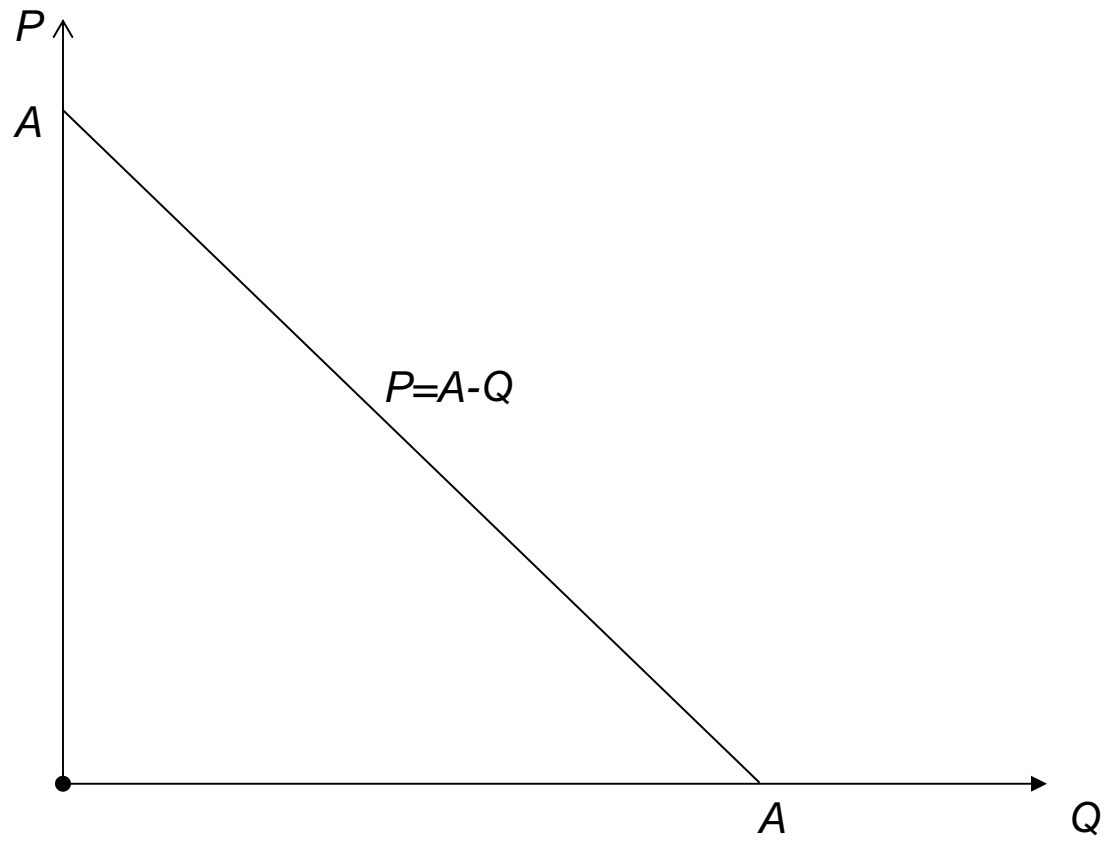
Cournot's oligopoly model (1838)

- A single good is produced by two firms (the industry is a “duopoly”).
- The cost for firm $i = 1, 2$ for producing q_i units of the good is given by $c_i q_i$ (“unit cost” is constant equal to $c_i > 0$).
- If the firms' total output is $Q = q_1 + q_2$ then the market price is

$$P = A - Q$$

if $A \geq Q$ and zero otherwise (linear inverse demand function). We also assume that $A > c$.

The inverse demand function



To find the Nash equilibria of the Cournot's game, we can use the procedures based on the firms' best response functions.


But first we need the firms payoffs (profits):

$$\begin{aligned}\pi_1 &= Pq_1 - c_1q_1 \\ &= (A - Q)q_1 - c_1q_1 \\ &= (A - q_1 - q_2)q_1 - c_1q_1 \\ &= (A - q_1 - q_2 - c_1)q_1\end{aligned}$$

and similarly,

$$\pi_2 = (A - q_1 - q_2 - c_2)q_2$$

**Firm 1's profit as a function of its output
(given firm 2's output)**



To find firm 1's best response to any given output q_2 of firm 2, we need to study firm 1's profit as a function of its output q_1 for given values of q_2 .

Using calculus, we set the derivative of firm 1's profit with respect to q_1 equal to zero and solve for q_1 :

$$q_1 = \frac{1}{2}(A - q_2 - c_1).$$

We conclude that the best response of firm 1 to the output q_2 of firm 2 depends on the values of q_2 and c_1 .

Because firm 2's cost function is $c_2 \neq c_1$, its best response function is given by

$$q_2 = \frac{1}{2}(A - q_1 - c_2).$$

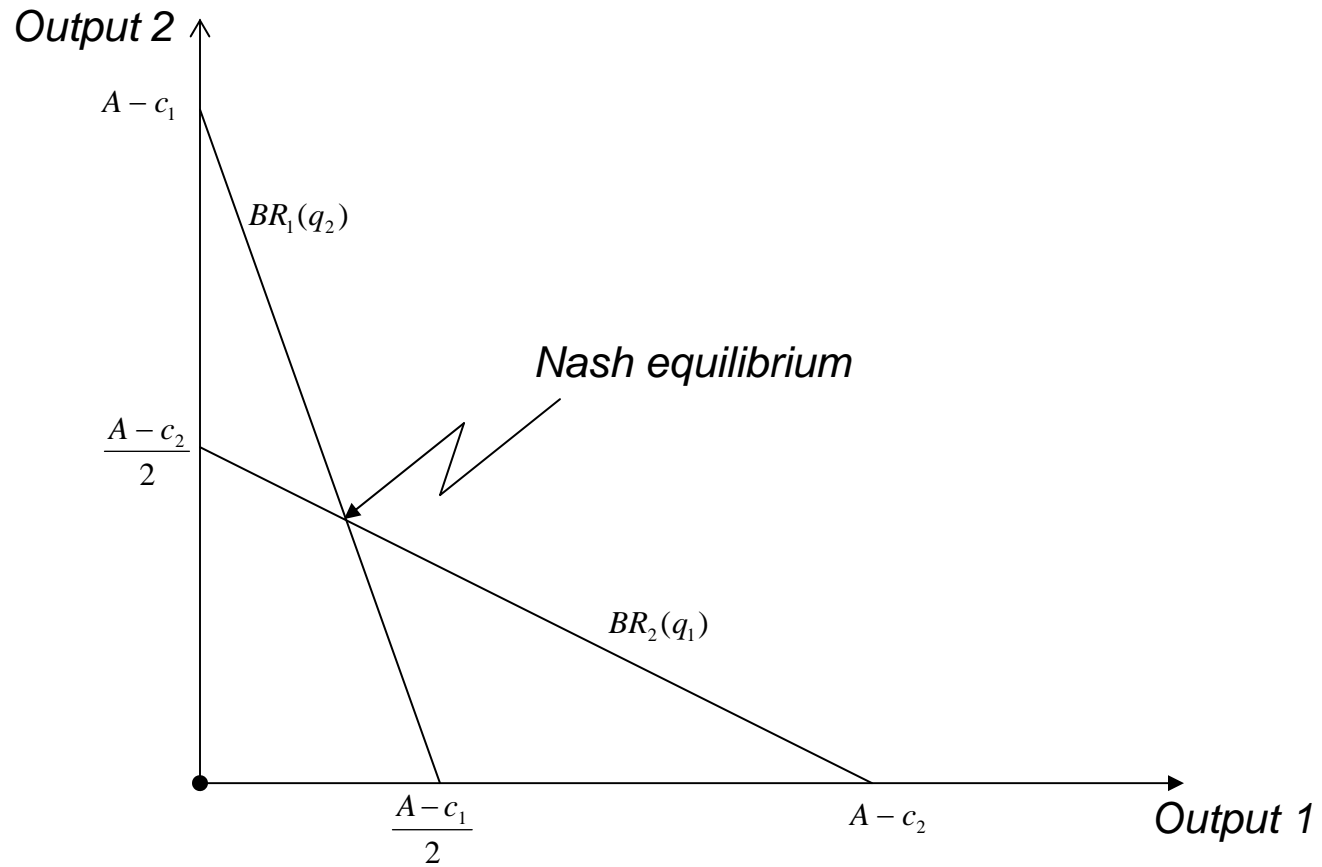
A Nash equilibrium of the Cournot's game is a pair (q_1^*, q_2^*) of outputs such that q_1^* is a best response to q_2^* and q_2^* is a best response to q_1^* .

From the figure below, we see that there is exactly one such pair of outputs


$$q_1^* = \frac{A+c_2-2c_1}{3} \quad \text{and} \quad q_2^* = \frac{A+c_1-2c_2}{3}$$

which is the solution to the two equations above.

The best response functions in the Cournot's duopoly game



**Nash equilibrium comparative statics
(a decrease in the cost of firm 2)**



A question: what happens when consumers are willing to pay more (A increases)?

In summary, this simple Cournot's duopoly game has a unique Nash equilibrium.

Two economically important properties of the Nash equilibrium are (to economic regulatory agencies):

- [1] The relation between the firms' equilibrium profits and the profit they could make if they act collusively.
- [2] The relation between the equilibrium profits and the number of firms.

- [1] Collusive outcomes: in the Cournot's duopoly game, there is a pair of outputs at which *both* firms' profits exceed their levels in a Nash equilibrium.
- [2] Competition: The price at the Nash equilibrium if the two firms have the *same* unit cost $c_1 = c_2 = c$ is given by

$$\begin{aligned} P^* &= A - q_1^* - q_2^* \\ &= \frac{1}{3}(A + 2c) \end{aligned}$$

which is above the unit cost c . But as the number of firm increases, the equilibrium price decreases, approaching c (zero profits!).

Stackelberg's duopoly model (1934)

How do the conclusions of the Cournot's duopoly game change when the firms move sequentially? Is a firm better off moving before or after the other firm?

Suppose that $c_1 = c_2 = c$ and that firm 1 moves at the start of the game. We may use backward induction to find the subgame perfect equilibrium.

- First, for *any* output q_1 of firm 1, we find the output q_2 of firm 2 that maximizes its profit. Next, we find the output q_1 of firm 1 that maximizes its profit, *given the strategy* of firm 2.

Firm 2

Since firm 2 moves after firm 1, a strategy of firm 2 is a *function* that associate an output q_2 for firm 2 for each possible output q_1 of firm 1.

We found that under the assumptions of the Cournot's duopoly game Firm 2 has a unique best response to each output q_1 of firm 1, given by

$$q_2 = \frac{1}{2}(A - q_1 - c)$$

(Recall that $c_1 = c_2 = c$).

Firm 1

Firm 1's strategy is the output q_1 the maximizes

$$\pi_1 = (A - q_1 - q_2 - c)q_1 \quad \text{subject to} \quad q_2 = \frac{1}{2}(A - q_1 - c)$$

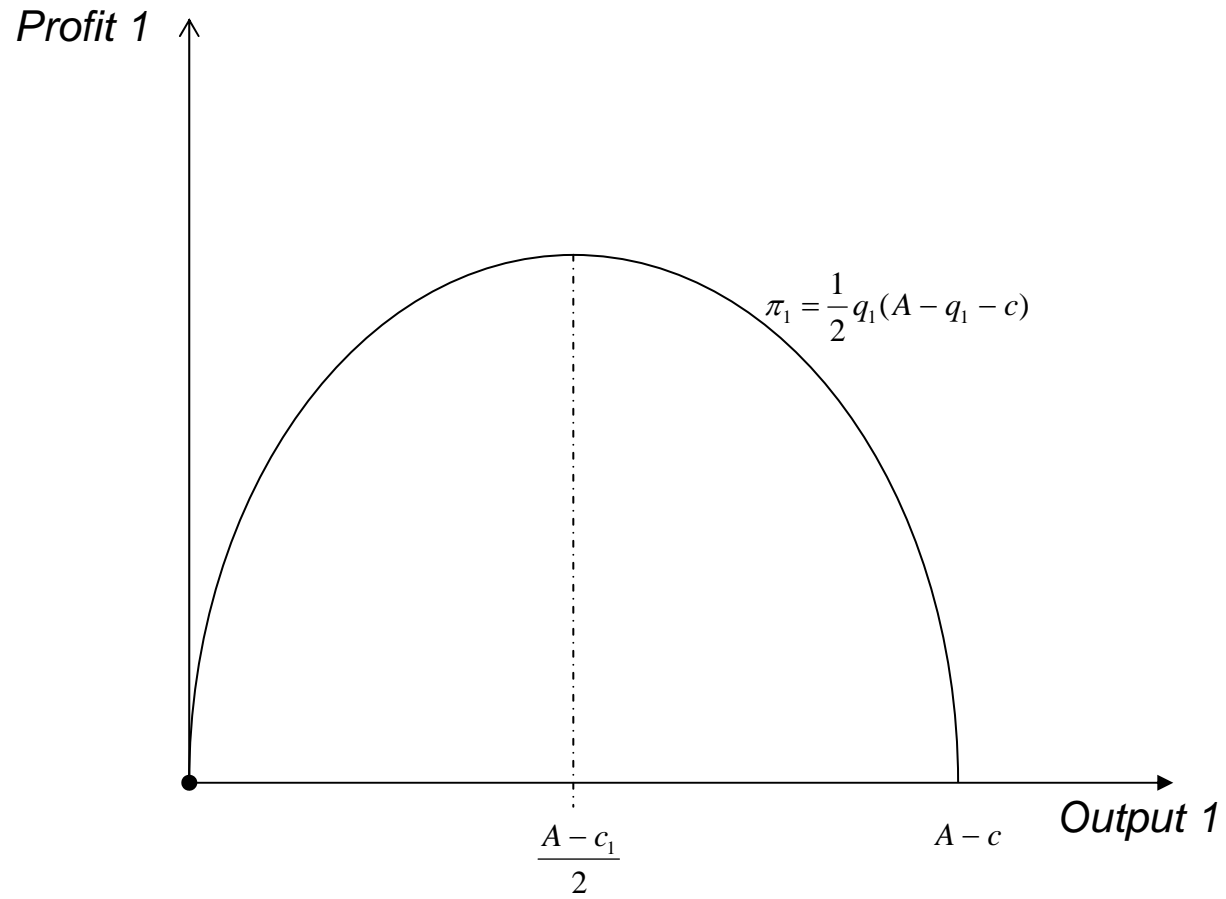
Thus, firm 1 maximizes

$$\pi_1 = (A - q_1 - (\frac{1}{2}(A - q_1 - c)) - c)q_1 = \frac{1}{2}q_1(A - q_1 - c).$$

This function is quadratic in q_1 that is zero when $q_1 = 0$ and when $q_1 = A - c$. Thus its maximizer is

$$q_1^* = \frac{1}{2}(A - c).$$

Firm 1's (first-mover) profit in Stackelberg's duopoly game



We conclude that Stackelberg's duopoly game has a unique subgame perfect equilibrium, in which firm 1's strategy is the output

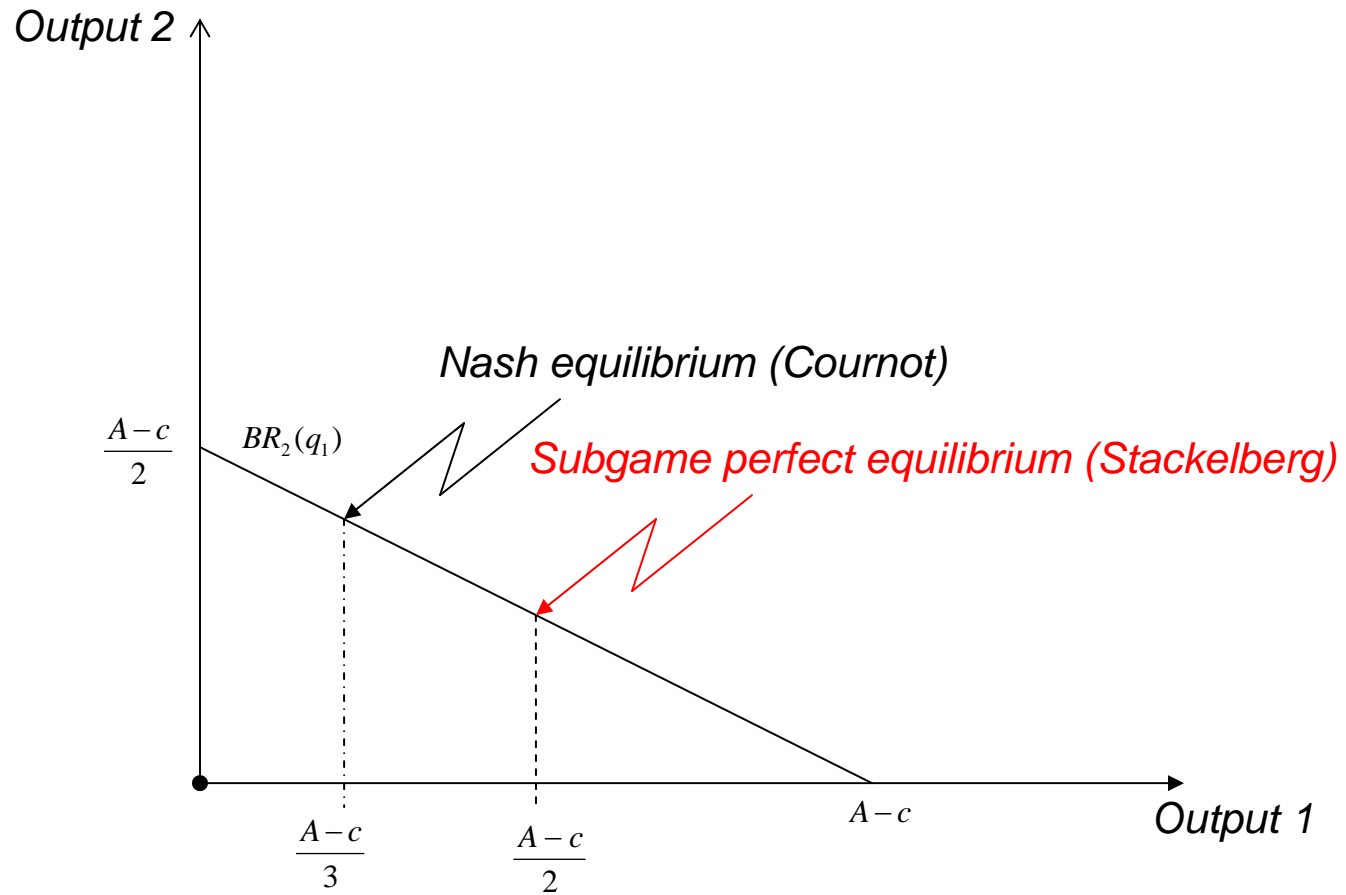
$$q_1^* = \frac{1}{2}(A - c)$$

and firm 2's output is

$$\begin{aligned} q_2^* &= \frac{1}{2}(A - q_1^* - c) \\ &= \frac{1}{2}\left(A - \frac{1}{2}(A - c) - c\right) \\ &= \frac{1}{4}(A - c). \end{aligned}$$

By contrast, in the unique Nash equilibrium of the Cournot's duopoly game under the same assumptions ($c_1 = c_2 = c$), each firm produces $\frac{1}{3}(A - c)$.

The subgame perfect equilibrium of Stackelberg's duopoly game



Auctions (the basics)

Auctions

From Babylonia to eBay, auctioning has a very long history.

- Babylon:
 - women at marriageable age.
- Athens, Rome, and medieval Europe:
 - rights to collect taxes,
 - dispose of confiscated property,
 - lease of land and mines,
 - and more...

- Auctions, broadly defined, are used to allocate significant economic resources.

Examples: works of art, government bonds, offshore tracts for oil exploration, radio spectrum, and more.

- Auctions take many forms. A game-theoretic framework enables to understand the consequences of various auction designs.
- Game theory can suggest the design likely to be most effective, and the one likely to raise the most revenues.

Types of auctions

Sequential / simultaneous

Bids may be called out sequentially or may be submitted simultaneously in sealed envelopes:

- English (or oral) – the seller actively solicits progressively higher bids and the item is sold to the highest bidder.
- Dutch – the seller begins by offering units at a “high” price and reduces it until all units are sold.
- Sealed-bid – all bids are made simultaneously, and the item is sold to the highest bidder.

First-price / second-price

The price paid may be the highest bid or some other price:

- First-price – the bidder who submits the highest bid wins and pay a price equal to her bid.
- Second-prices – the bidder who submits the highest bid wins and pay a price equal to the second highest bid.


Variants: all-pay (lobbying), discriminatory, uniform, Vickrey (William Vickrey, Nobel Laureate 1996), and more.

Private-value / common-value


Bidders can be certain or uncertain about each other's valuation:

- In private-value auctions, valuations differ among bidders, and each bidder is certain of her own valuation and can be certain or uncertain of every other bidder's valuation.
- In common-value auctions, all bidders have the same valuation, but bidders do not know this value precisely and their estimates of it vary.

First-price auction class experiment



Second-price auction class experiment



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 $i \neq \{1, j\}$

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.

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 Hampshireites (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.

What have we learned from Social Learning?

- Finding 1

- Individuals 'ignore' their own information and follow a herd.

- Finding 2

- Herds often adopt a wrong action.

- Finding 3

- Mass behavior may be idiosyncratic and fragile.

Informational cascades and herd behavior

Two phenomena that have elicited particular interest are *informational cascades* and *herd behavior*.

- Cascade: agents 'ignore' their private information when choosing an action.
- Herd: agents choose the same action, not necessarily ignoring their private information.

- While the terms informational cascade and herd behavior are used interchangeably there is a significant difference between them.
- In an informational cascade, an agent considers it optimal to follow the behavior of her predecessors without regard to her private signal.
- When acting in a herd, agents choose the same action, not necessarily ignoring their private information.
- Thus, an informational cascade implies a herd but a herd is not necessarily the result of an informational cascade.

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

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


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 decision problem

- The optimal decision rule is given by

$$x_n = 1 \text{ if and only if } \mathbb{E} \left[\sum_{i=1}^N \theta_i \mid \mathcal{I}_n \right] \geq 0.$$

Since \mathcal{I}_n does not provide any information about the content of successors' signals, we obtain

$$x_n = 1 \text{ if and only if } \mathbb{E} \left[\sum_{i=1}^n \theta_i \mid \mathcal{I}_n \right] \geq 0$$

Hence,

$$x_n = 1 \text{ if and only if } \theta_n \geq -\mathbb{E} \left[\sum_{i=1}^{n-1} \theta_i \mid \mathcal{I}_n \right].$$

The cutoff process

- For any n , the optimal strategy is the *cutoff strategy*

$$x_n = \begin{cases} 1 & \text{if } \theta_n \geq \hat{\theta}_n \\ 0 & \text{if } \theta_n < \hat{\theta}_n \end{cases}$$

where

$$\hat{\theta}_n = -\mathbb{E} \left[\sum_{i=1}^{n-1} \theta_i \mid \mathcal{I}_n \right]$$

is the optimal history-contingent cutoff.

- $\hat{\theta}_n$ is sufficient to characterize the individual behavior, and $\{\hat{\theta}_n\}$ characterizes the social behavior of the economy.

Overview of results

Perfect information

- A cascade need not arise, but herd behavior must arise.

Imperfect information

- Herd behavior is impossible. There are periods of uniform behavior, punctuated by increasingly rare switches.

- The similarity:
 - Agents can, for a long time, make the same (incorrect) choice.
- The difference:
 - Under perfect information, a herd is an absorbing state. Under imperfect information, continued, occasional and sharp shifts in behavior.

- 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?

The case of perfect information

The optimal history-contingent cutoff rule is

$$\hat{\theta}_n = -\mathbb{E} \left[\sum_{i=1}^{n-1} \theta_i \mid x_1, \dots, x_{n-1} \right],$$

and $\hat{\theta}_n$ is different from $\hat{\theta}_{n-1}$ only by the information reveals by the action of agent $(n - 1)$

$$\hat{\theta}_n = \hat{\theta}_{n-1} - \mathbb{E} \left[\theta_{n-1} \mid \hat{\theta}_{n-1}, x_{n-1} \right],$$

The cutoff dynamics thus follow the cutoff process

$$\hat{\theta}_n = \begin{cases} \frac{-1 + \hat{\theta}_{n-1}}{2} & \text{if } x_{n-1} = 1 \\ \frac{1 + \hat{\theta}_{n-1}}{2} & \text{if } x_{n-1} = 0 \end{cases}$$

where $\hat{\theta}_1 = 0$.

Informational cascades

- $-1 < \hat{\theta}_n < 1$ for any n so any player takes his private signal into account in a non-trivial way.

Herd behavior

- $\{\hat{\theta}_n\}$ has the martingale property by the Martingale Convergence Theorem a limit-cascade implies a herd.

The case of imperfect information

The optimal history-contingent cutoff rule is

$$\hat{\theta}_n = -\mathbb{E} \left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1} \right],$$

which can take two values conditional on $x_{n-1} = 1$ or $x_{n-1} = 0$

$$\begin{aligned} \bar{\theta}_n &= -\mathbb{E} \left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1} = 1 \right], \\ \underline{\theta}_n &= -\mathbb{E} \left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1} = 0 \right]. \end{aligned}$$

where $\bar{\theta}_n = -\underline{\theta}_n$.

The law of motion for $\bar{\theta}_n$ is given by

$$\begin{aligned}\bar{\theta}_n = & P(x_{n-2} = 1 | x_{n-1} = 1) \left\{ \bar{\theta}_{n-1} - \mathbb{E}[\theta_{n-1} | x_{n-2} = 1] \right\} \\ & + P(x_{n-2} = 0 | x_{n-1} = 1) \left\{ \underline{\theta}_{n-1} - \mathbb{E}[\theta_{n-1} | x_{n-2} = 0] \right\},\end{aligned}$$

which simplifies to

$$\begin{aligned}\bar{\theta}_n = & \frac{1 - \bar{\theta}_{n-1}}{2} \left[\bar{\theta}_{n-1} - \frac{1 + \bar{\theta}_{n-1}}{2} \right] \\ & + \frac{1 - \underline{\theta}_{n-1}}{2} \left[\underline{\theta}_{n-1} - \frac{1 + \underline{\theta}_{n-1}}{2} \right].\end{aligned}$$

Given that $\bar{\theta}_n = -\bar{\theta}_n$, the cutoff dynamics under imperfect information follow the cutoff process

$$\hat{\theta}_n = \begin{cases} -\frac{1+\hat{\theta}_{n-1}^2}{2} & \text{if } x_{n-1} = 1 \\ \frac{1+\hat{\theta}_{n-1}^2}{2} & \text{if } x_{n-1} = 0 \end{cases}$$

where $\hat{\theta}_1 = 0$.


Informational cascades

- $-1 < \hat{\theta}_n < 1$ for any n so any player takes his private signal into account in a non-trivial way.

Herd behavior

- $\{\hat{\theta}_n\}$ is not convergent (proof is hard!) and the divergence of cutoffs implies divergence of actions.
- Behavior exhibits periods of uniform behavior, punctuated by increasingly rare switches.

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?!

