UC Berkeley
Haas School of Business
Game Theory
(EMBA 296 & EWMBA 211)
Summer 2016

Oligopoly, signaling, auctions, social learning and bargaining

Block 4

Jul 28-30, 2016

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 <u>strategic</u> 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 <u>extensive</u> game (a Nash equilibrium in every <u>subgame</u> of the extensive game).

$\Longrightarrow$ Every subgame perfect equilibrium is also a Nash equilibrium.

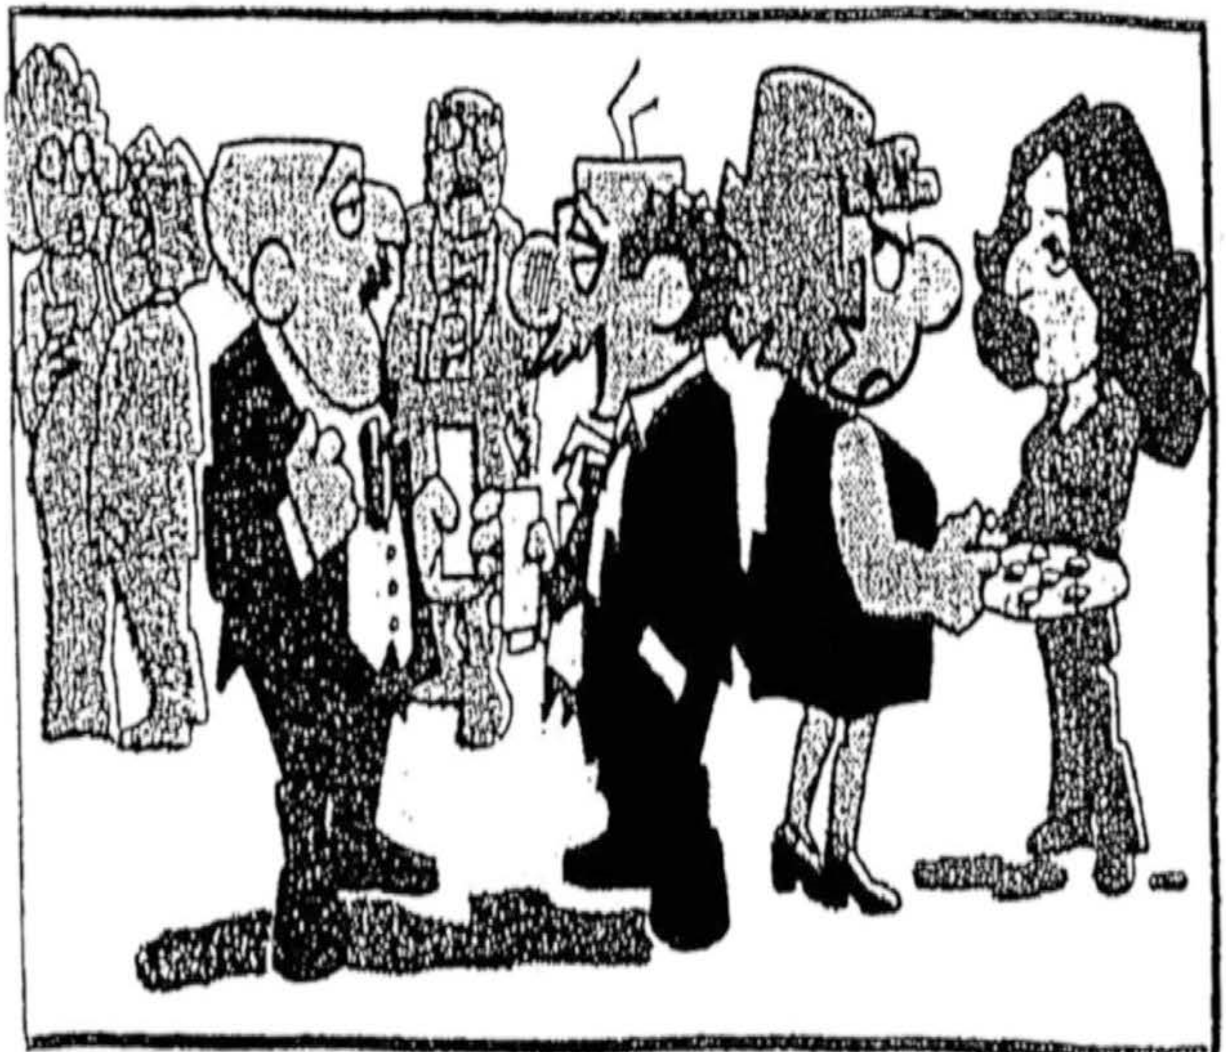

"LORETTA'S DRIVING BECAUSE I'M DRINKING, AND I'M DRINKING BECAUSE SHE'S DRIVING."

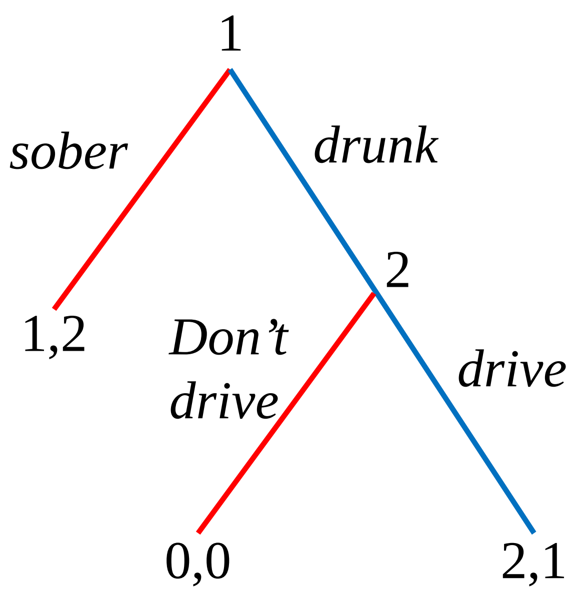
1
sober
drunk
2
1,2
Don't drive
drive
0,0
2,1

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

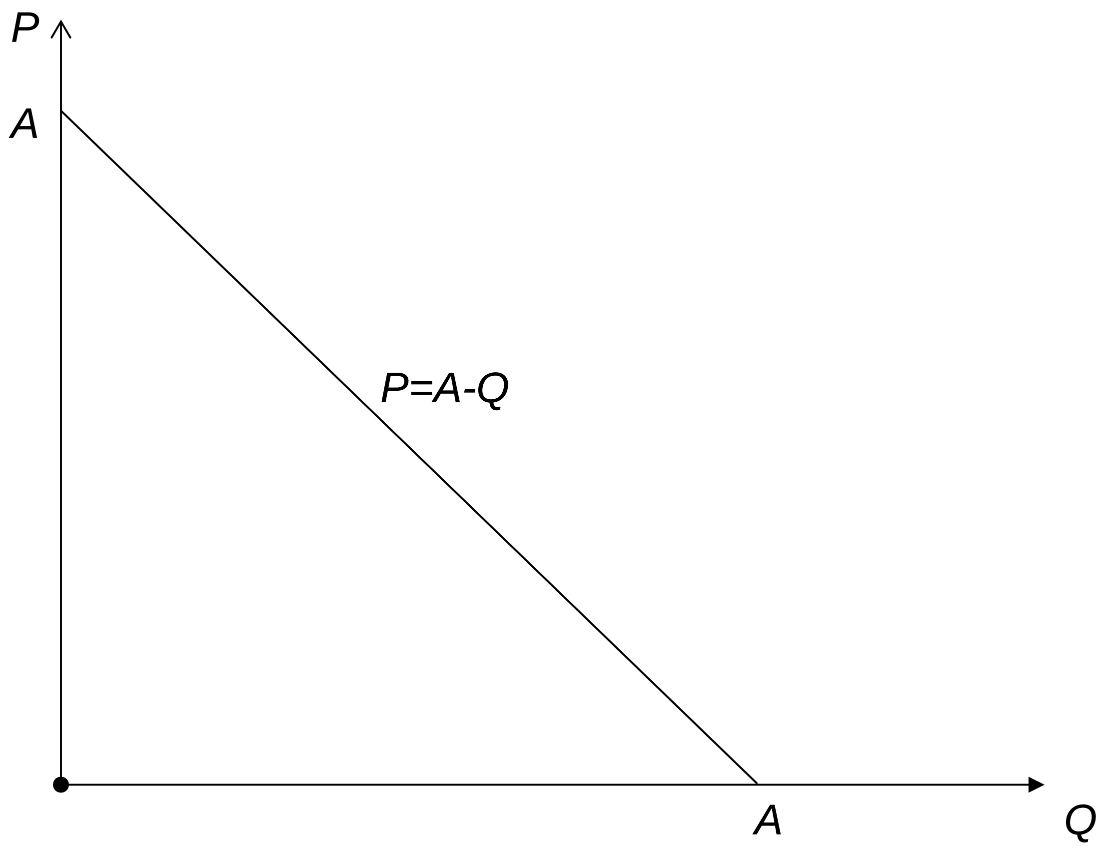

To find the Nash equilibria of the Cournot's game, we can use the procedures based on the firms' best response functions.

But first we need the firms payoffs (profits):

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

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)

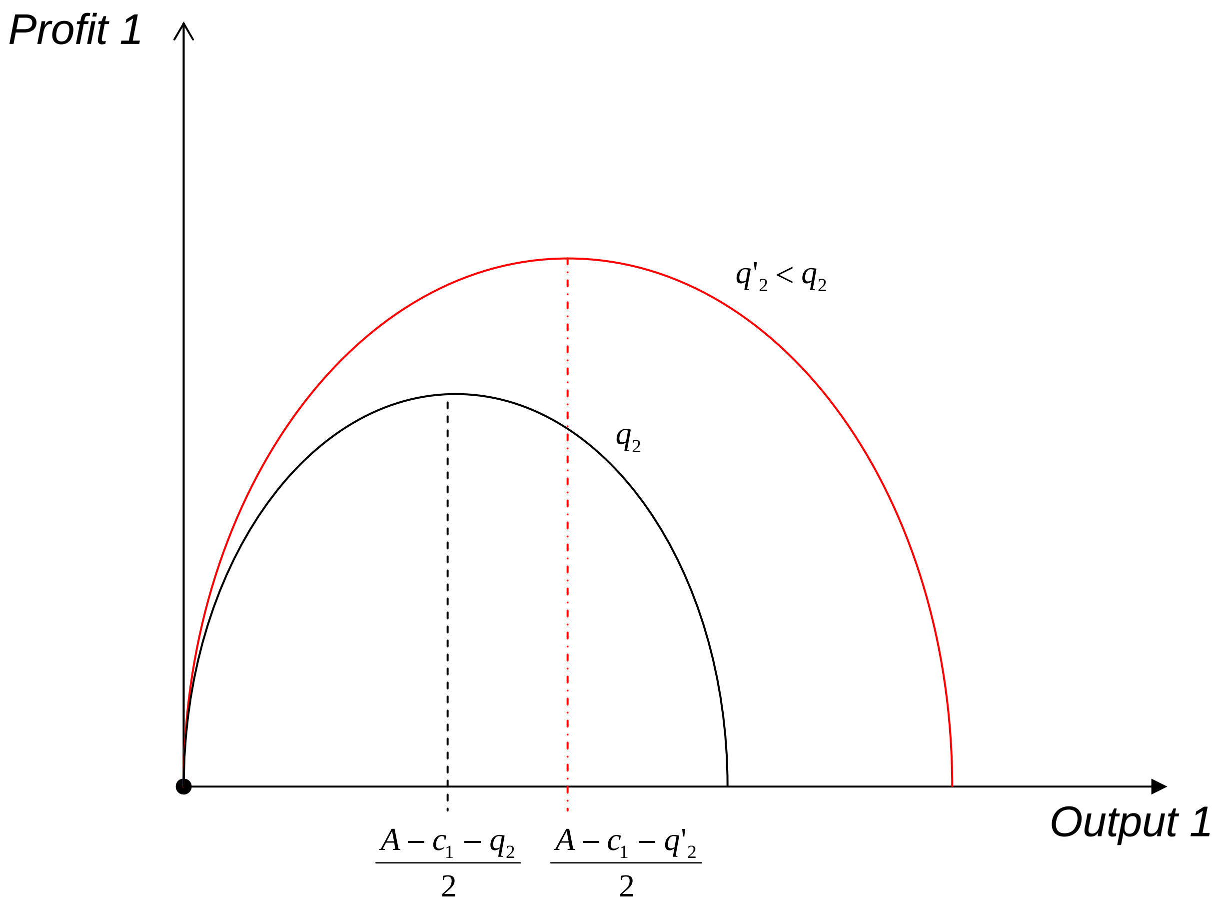

To find firm 1's best response to any given output q_2 of firm 2, we need to study firm 1's profit as a function of its output q_1 for given values of q_2.

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

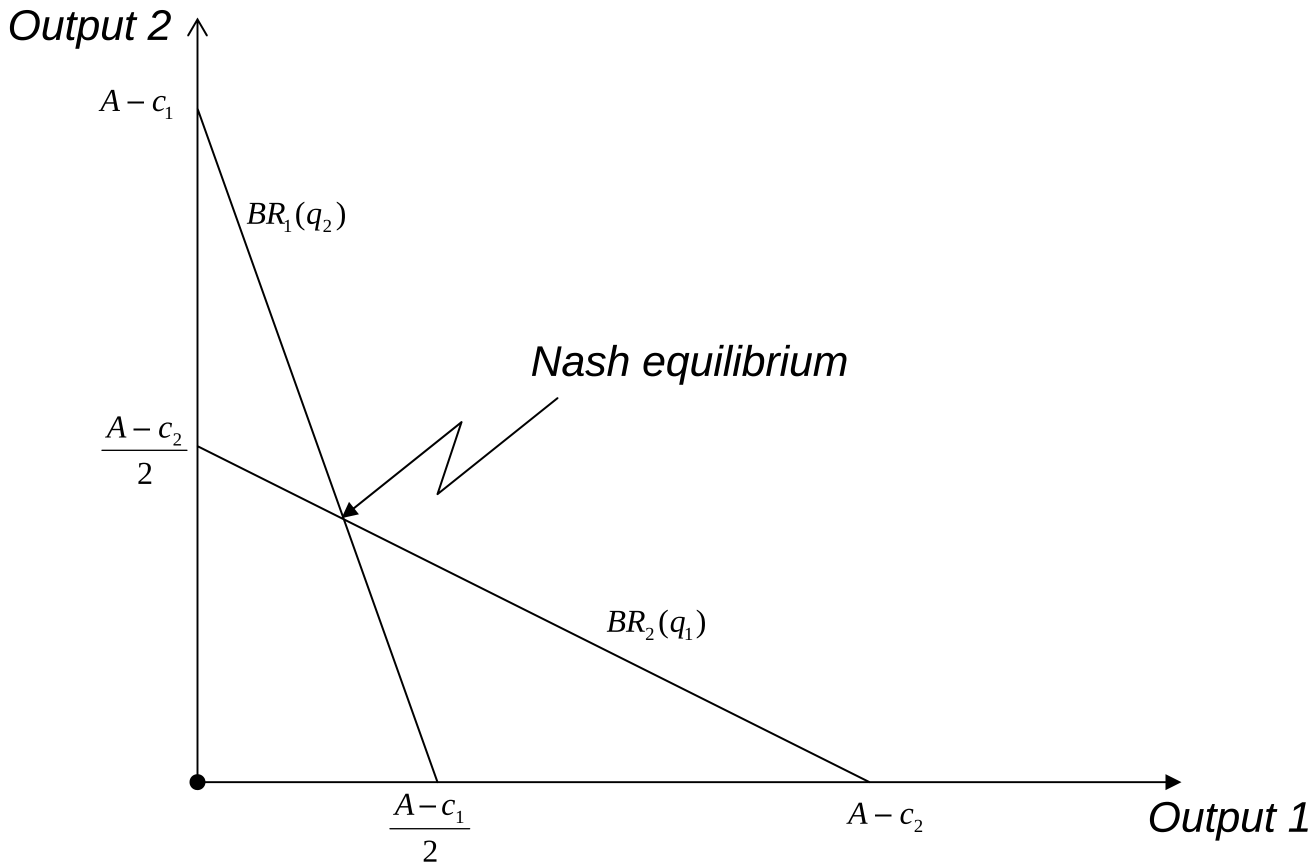

Nash equilibrium comparative statics
(a decrease in the cost of firm 2)

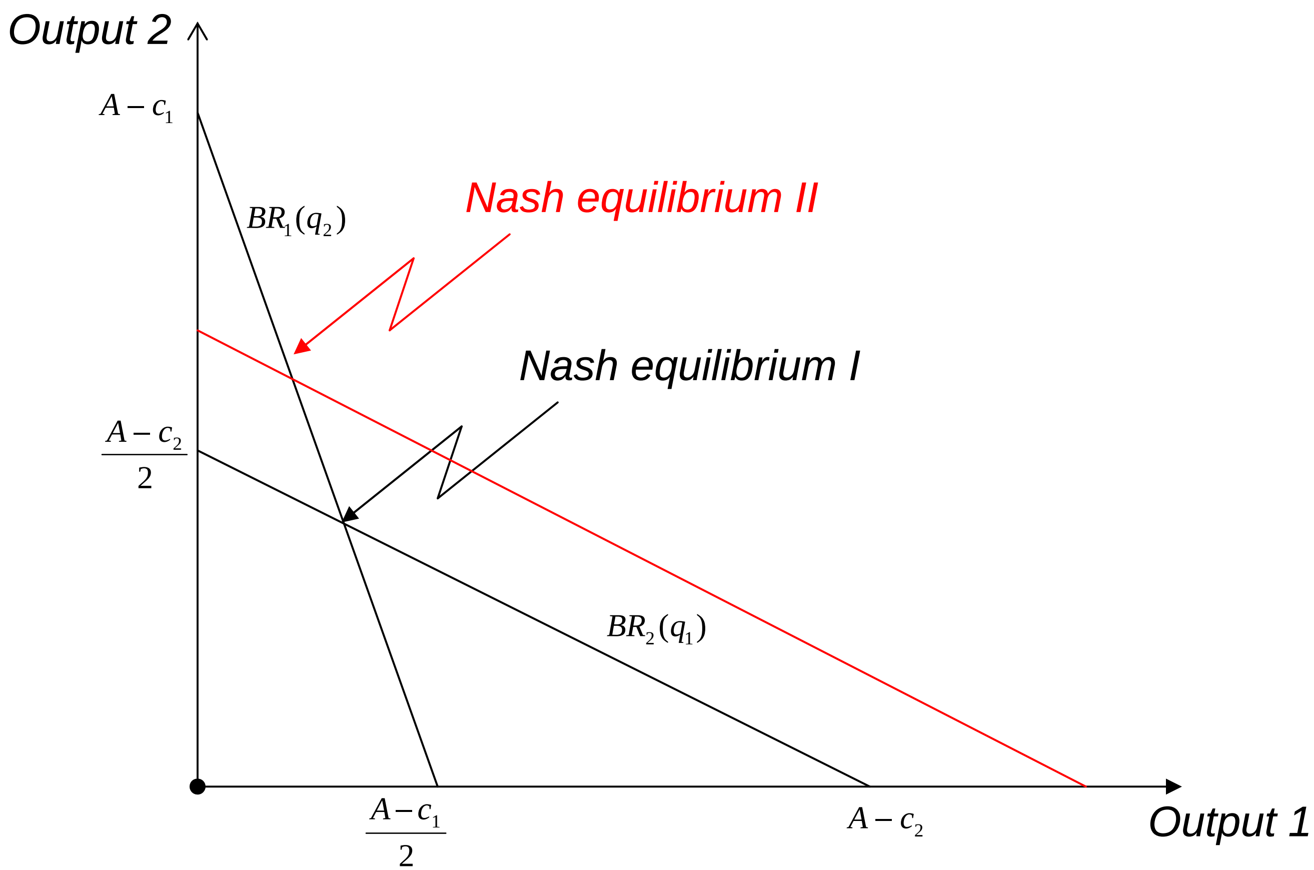

A question: what happens when consumers are willing to pay more (*A* increases)?

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 deceases, 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

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

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

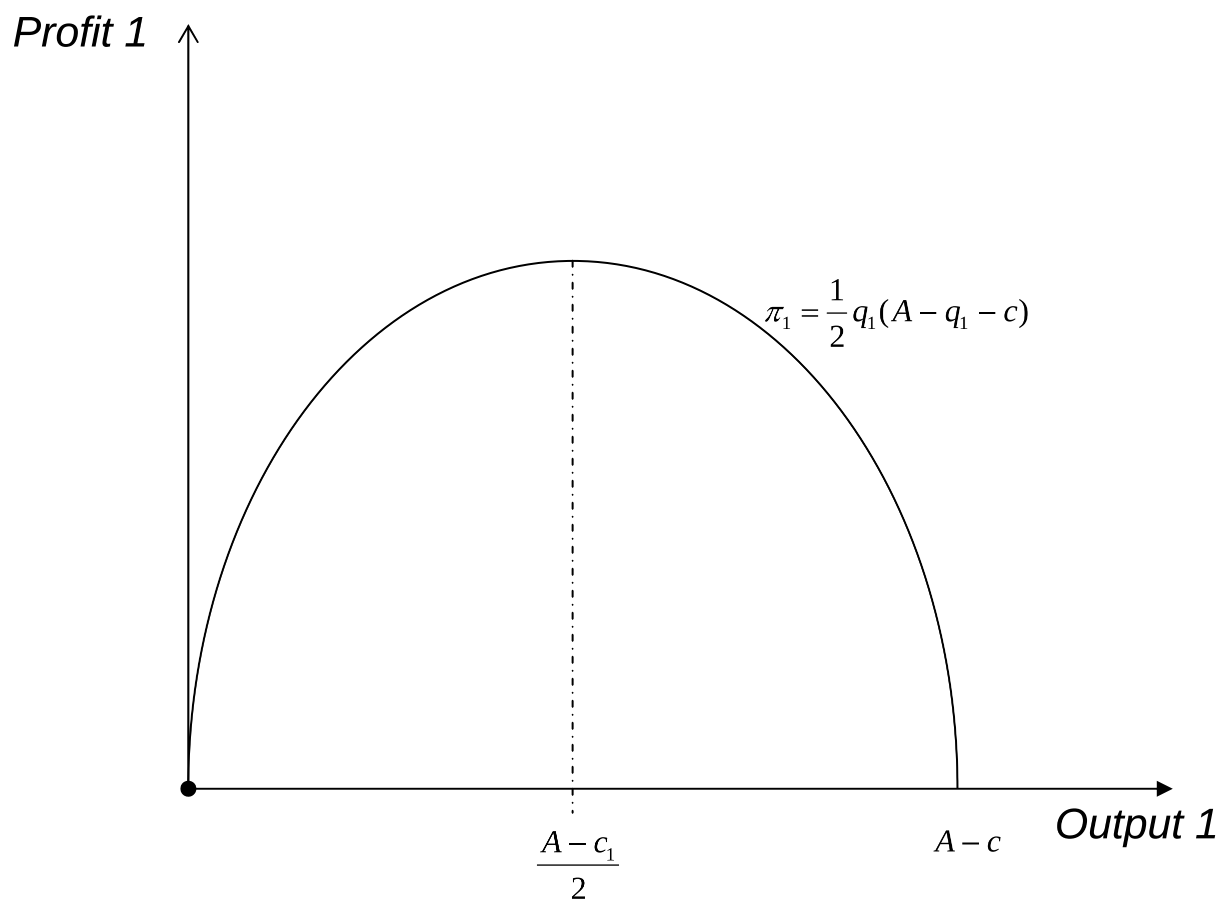

We conclude that Stackelberg's duopoly game has a unique subgame perfect equilibrium, in which firm 1's strategy is the output

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

and firm 2's output is

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

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

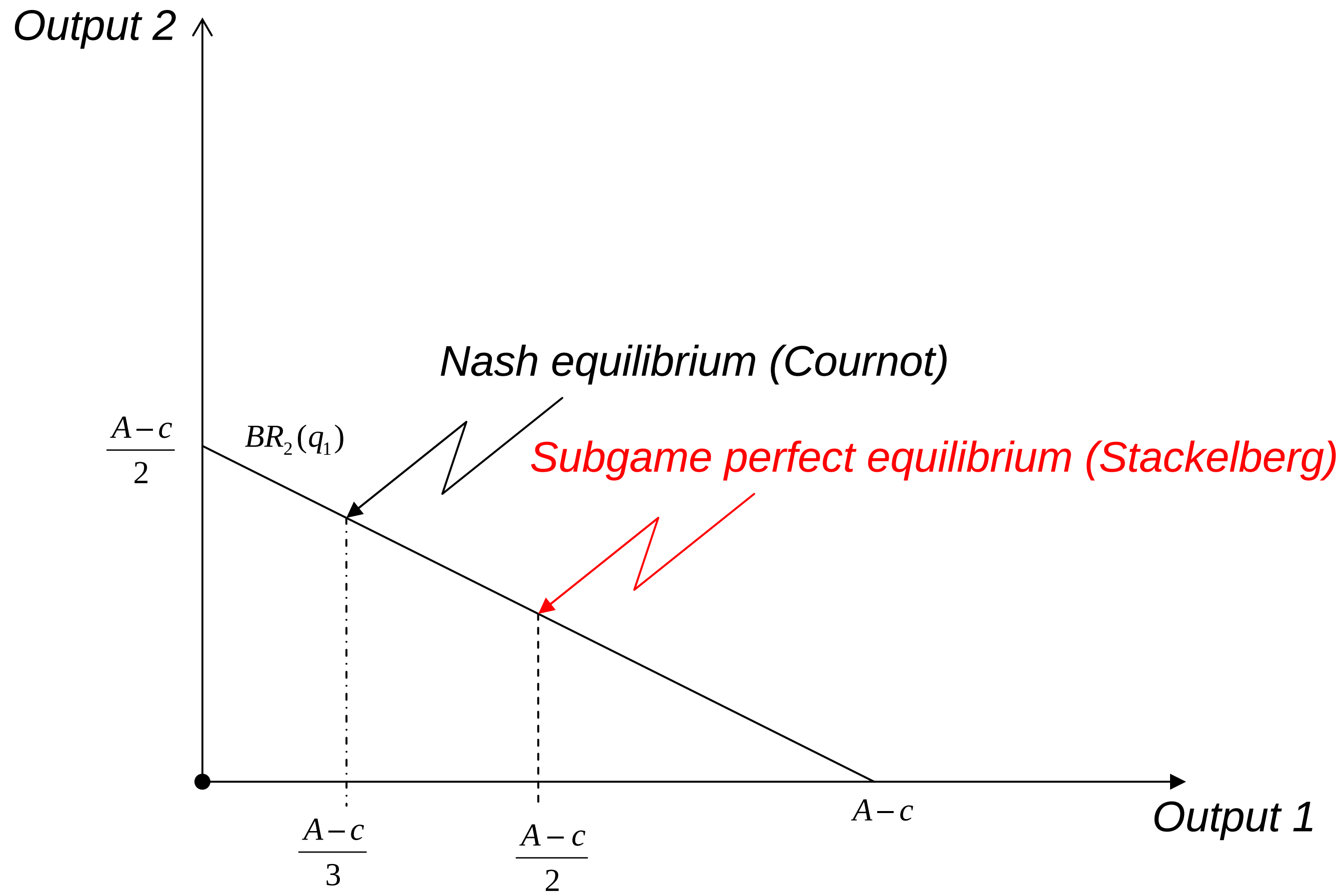

Avoiding the Bertrand trap

If you are in a situation satisfying the following assumptions, then you will end up in a Bertrand trap (zero profits):

[1] Homogenous products

[2] Consumers know all firm prices

[3] No switching costs

[4] No cost advantages

[5] No capacity constraints

[6] No future considerations

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.

$\Longrightarrow$ 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”

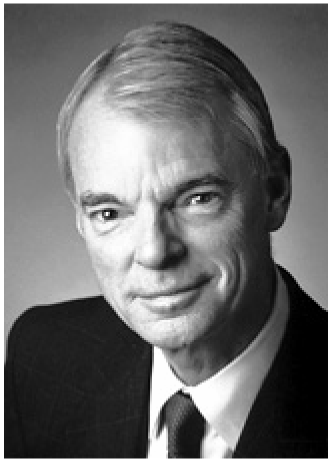

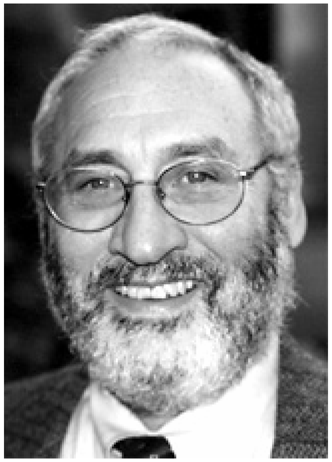

The Market for Lemons

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.

$\Longrightarrow$ 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 <u>signal</u> 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 paid his marginal product and we would have an <u>efficient</u> 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

$$\begin{aligned} s(e_1) &= a_1 \\ s(e_2) &= a_2 \end{aligned}$$

[2] self-selection conditions

$$\begin{aligned} s(e_1) - c_1 e_1 &\geq s(e_2) - c_1 e_2 \\ s(e_2) - c_2 e_2 &\geq s(e_1) - c_2 e_1 \end{aligned}$$

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

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

and

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

$\Longrightarrow$ Signaling can make things better or worse – each case has to 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.

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!

Auctions

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 soled to the highest bidder.
- Dutch – the seller begins by offering units at a "high" price and reduces it until all units are soled.
- 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 (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

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

where $\overline{b}$ is the highest bid.

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

$$\begin{array}{ll} [1] & v_2 \le b_1 \le v_1 \\ [2] & b_j \le b_1 \text{ \underline{for all} } j \ne 1 \\ [3] & b_j = b_1 \text{ \underline{for some} } j \ne 1 \end{array}$$

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.

$\Longrightarrow$ 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

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

  where $\overline{b}$ 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] $b_i = 0$ <u>for all</u> $i \neq \{1, j\}$

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.

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."

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

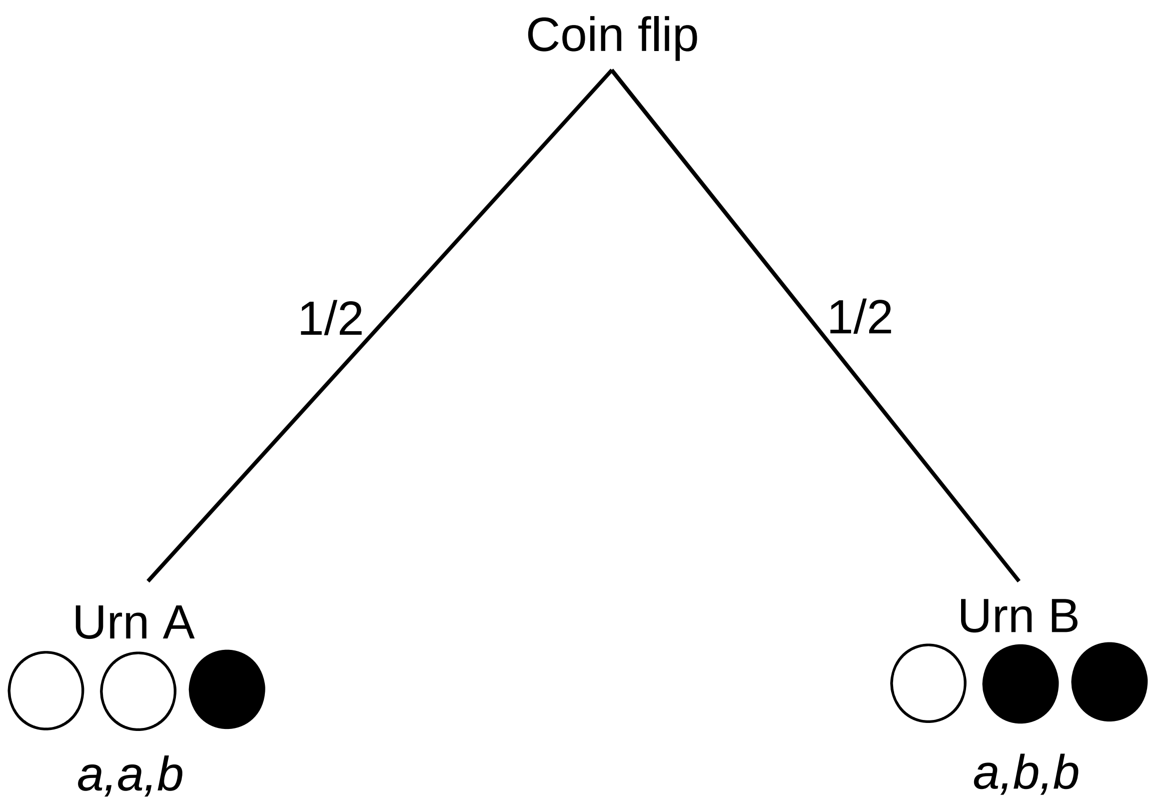
Coin flip
1/2
1/2
Urn A
Urn B
a,a,b
a,b,b

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 $n \in \{1, ..., N\}$ 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

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

- 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?

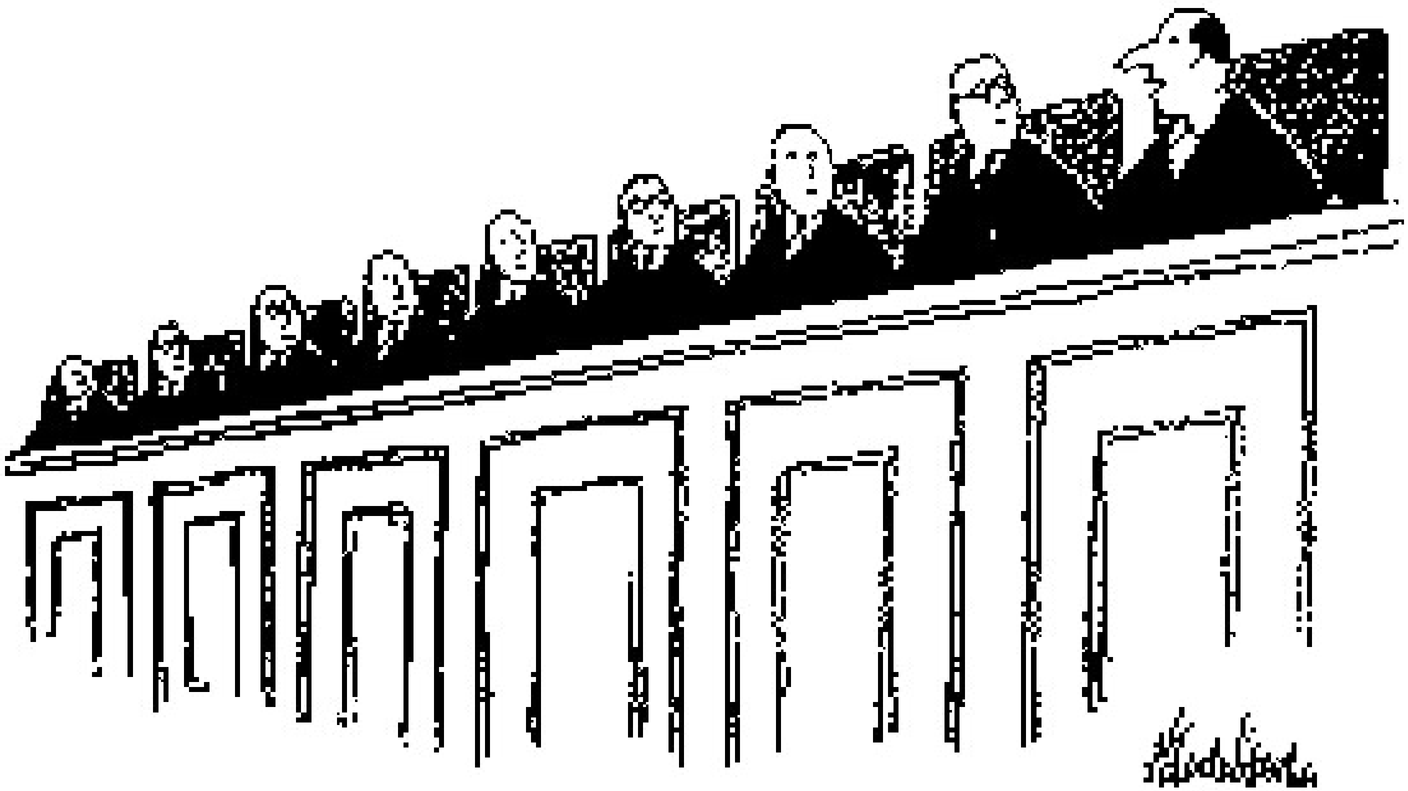

Imperfect information:
Which way is the wind blowing?!

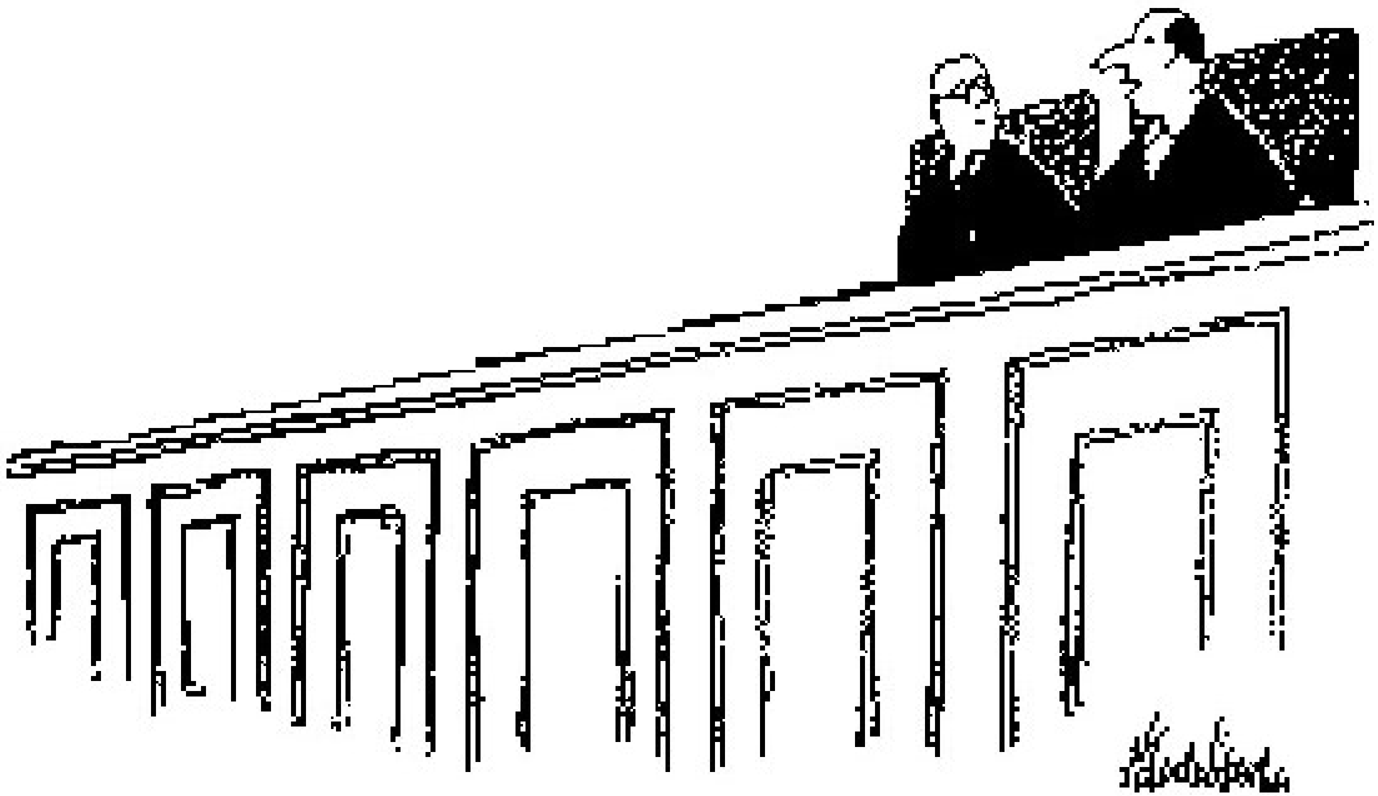

A three-agent example

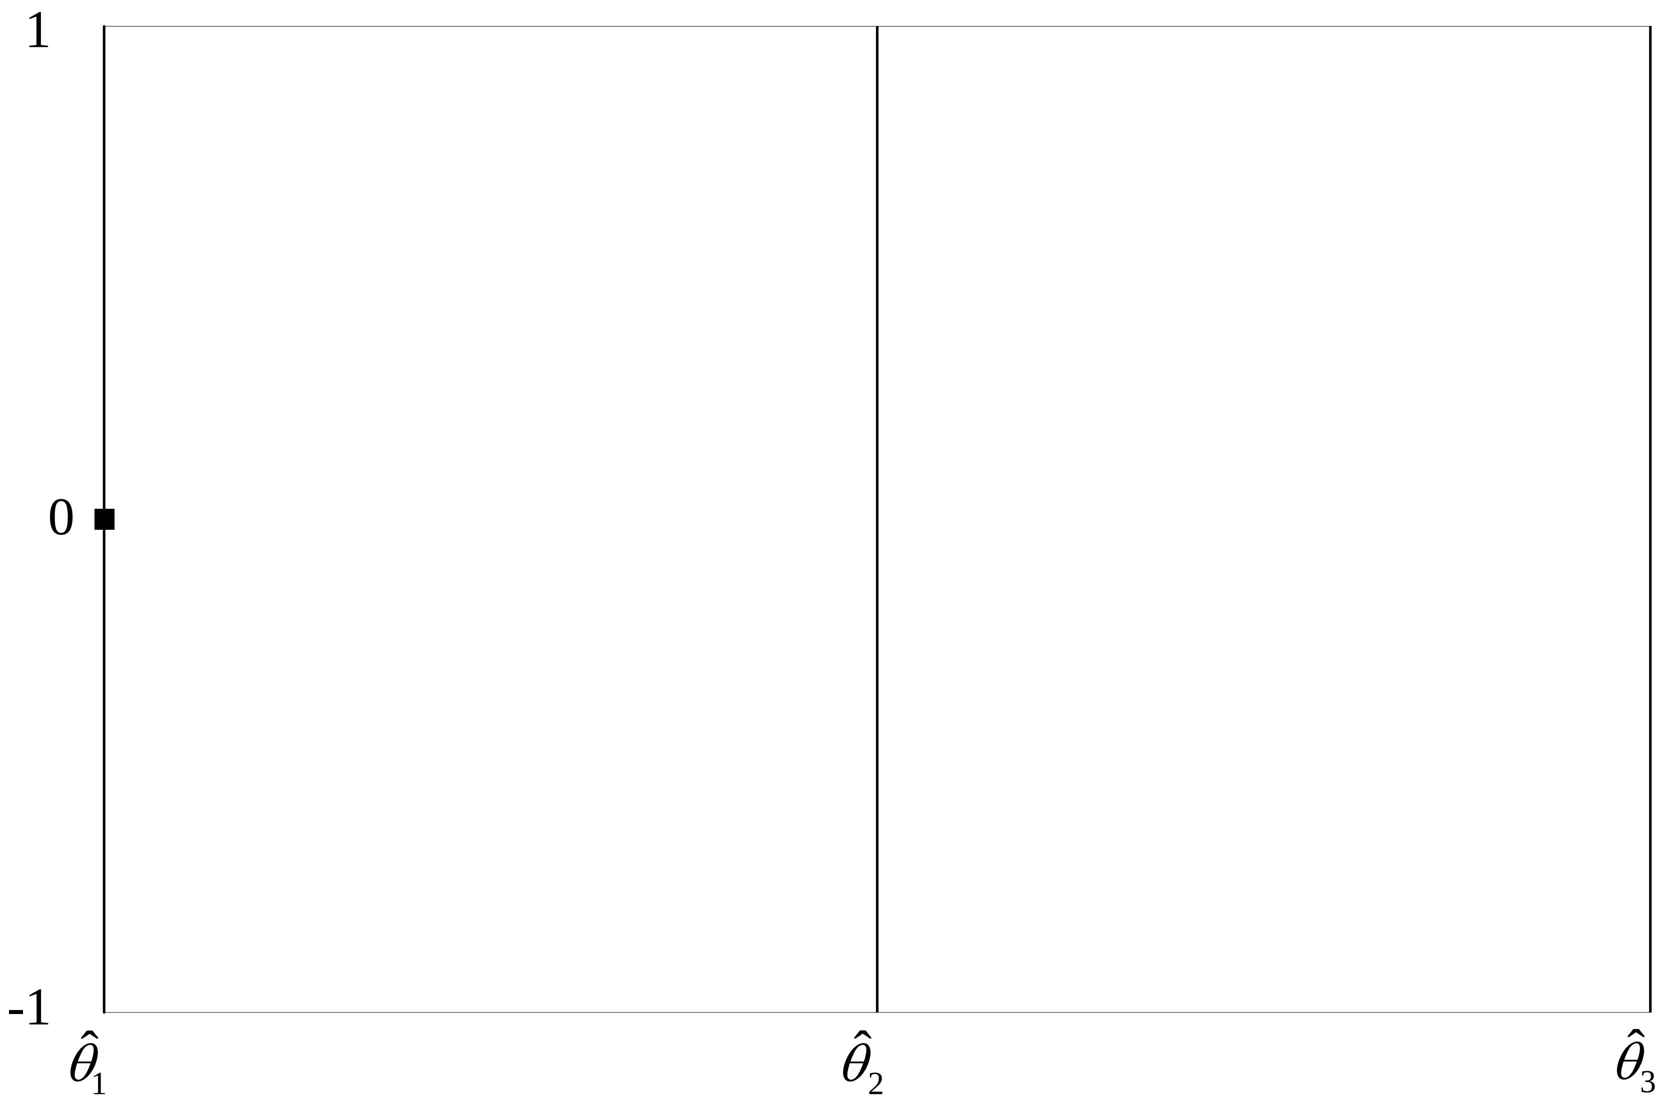

A three-agent example

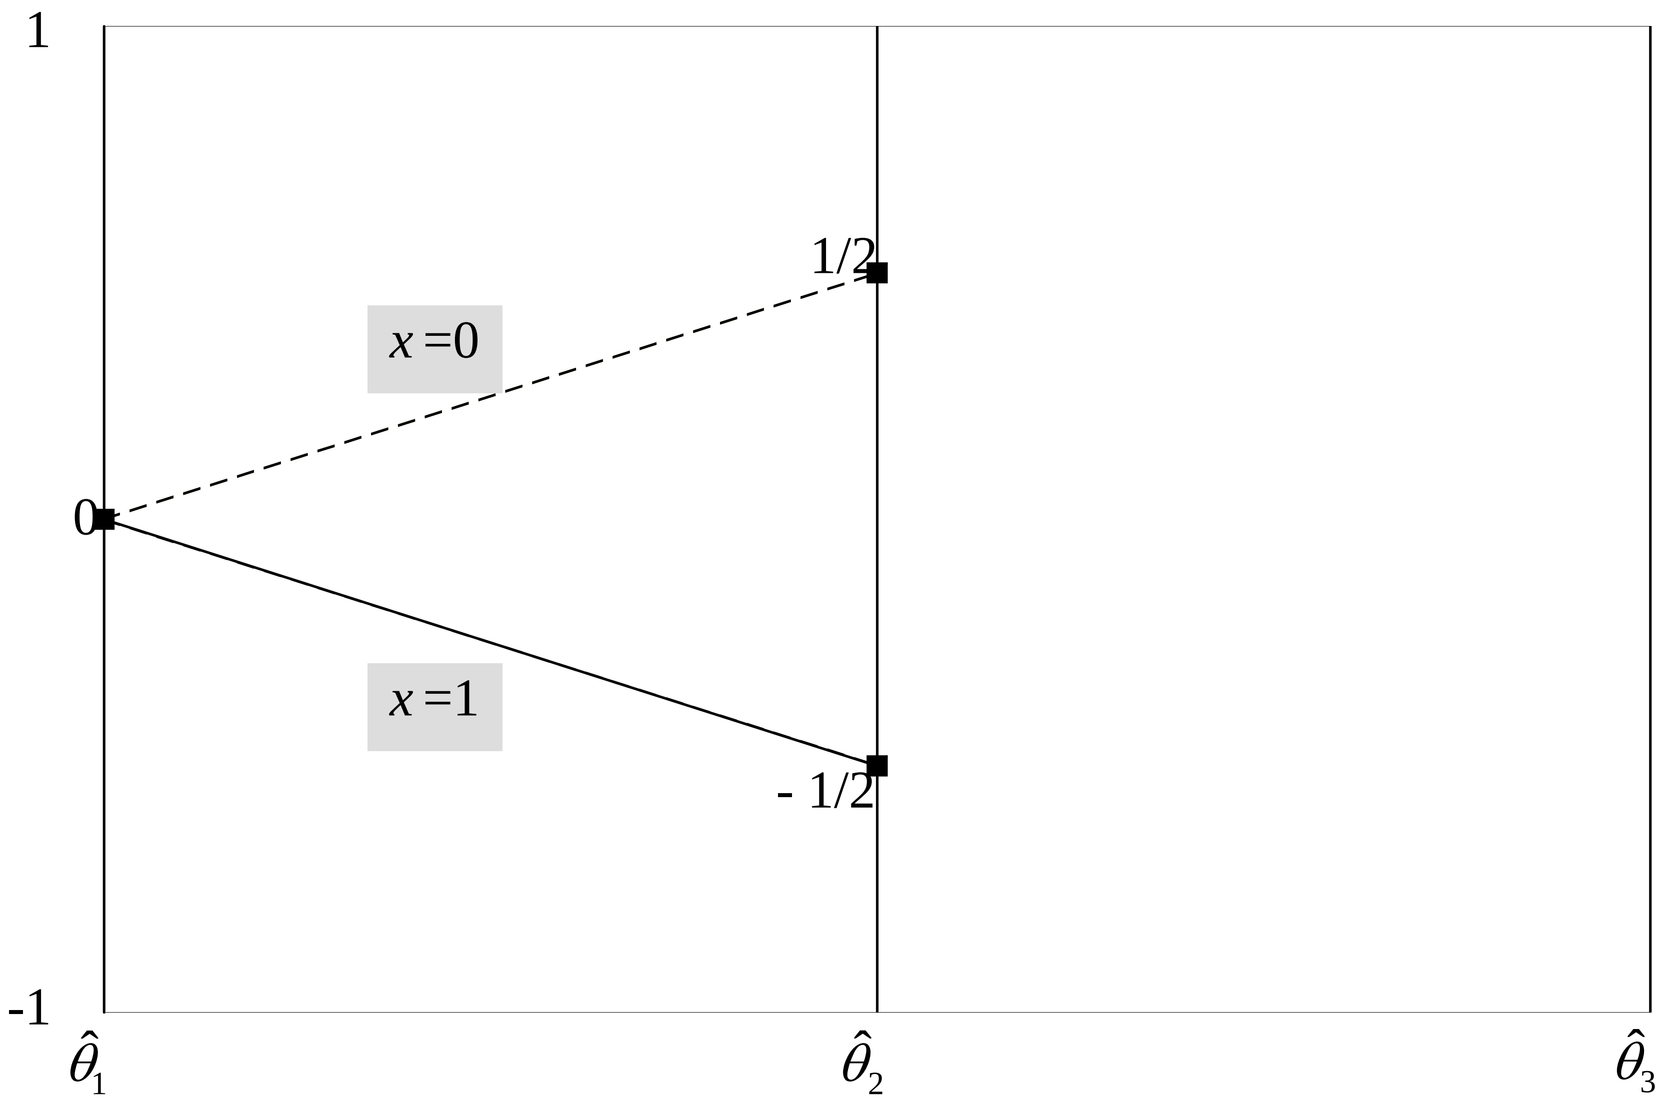

A three-agent example under perfect information

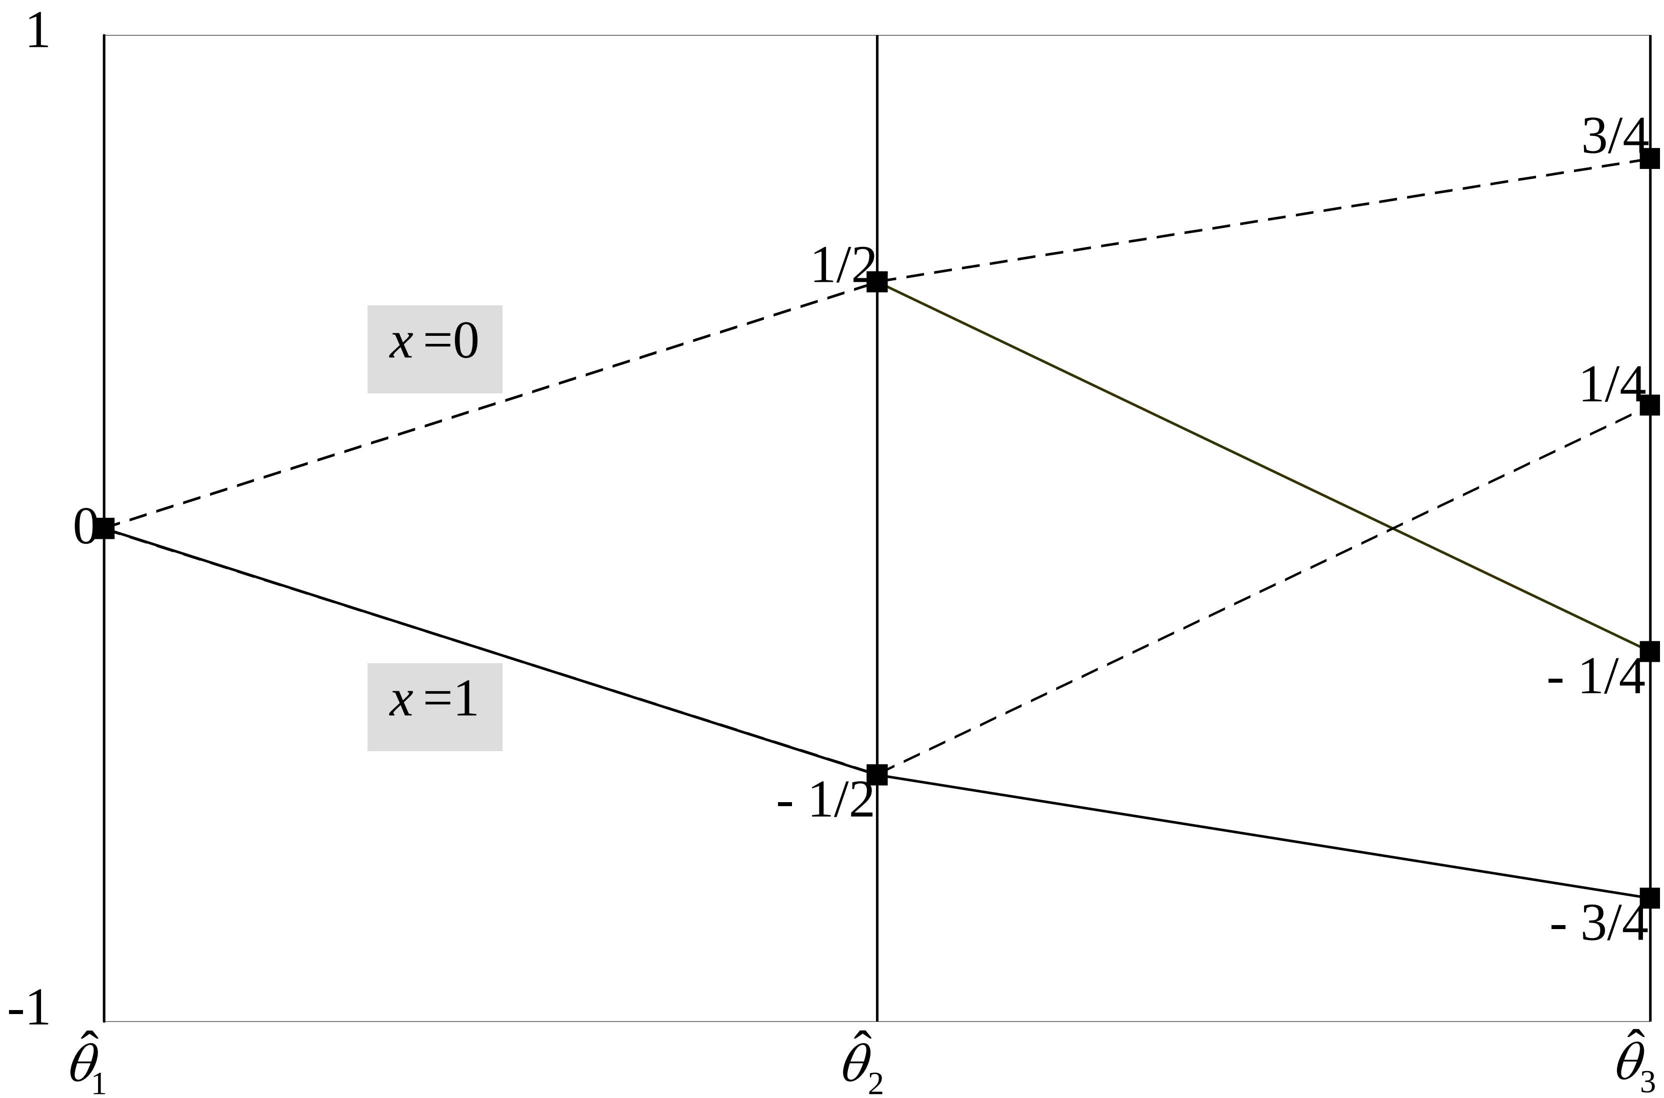

A three-agent example under imperfect information

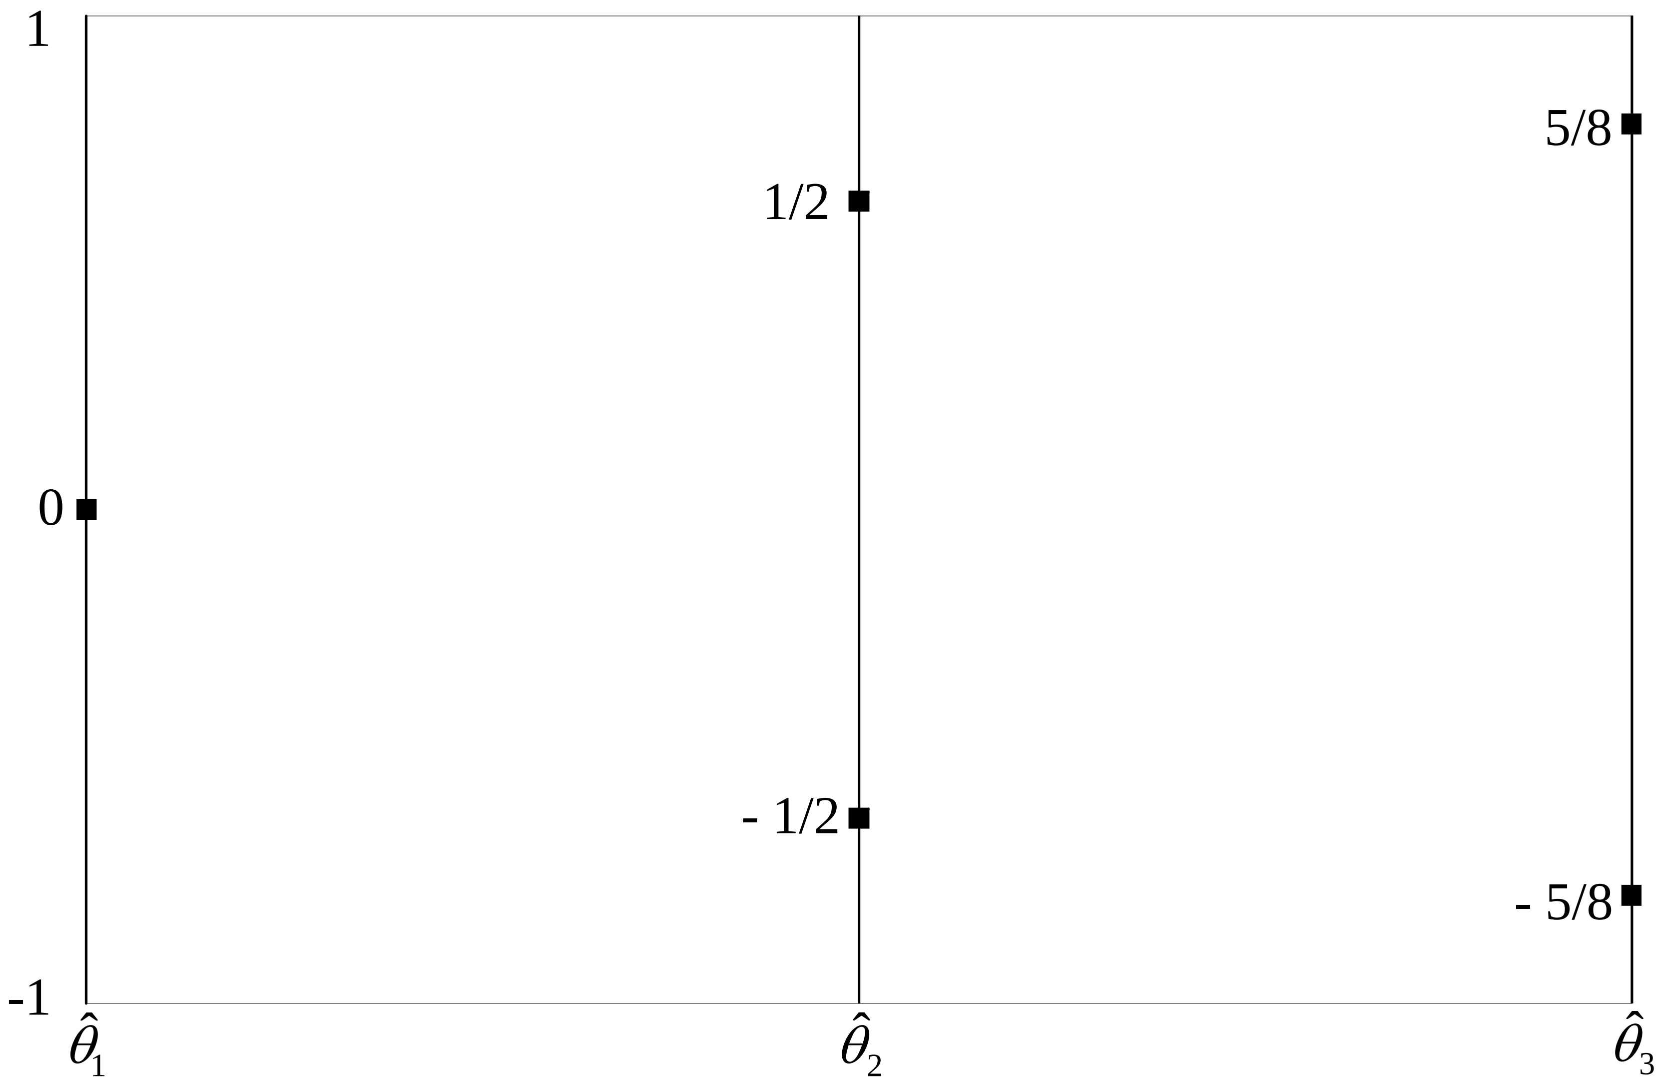

A sequence of cutoffs under imperfect and perfect information

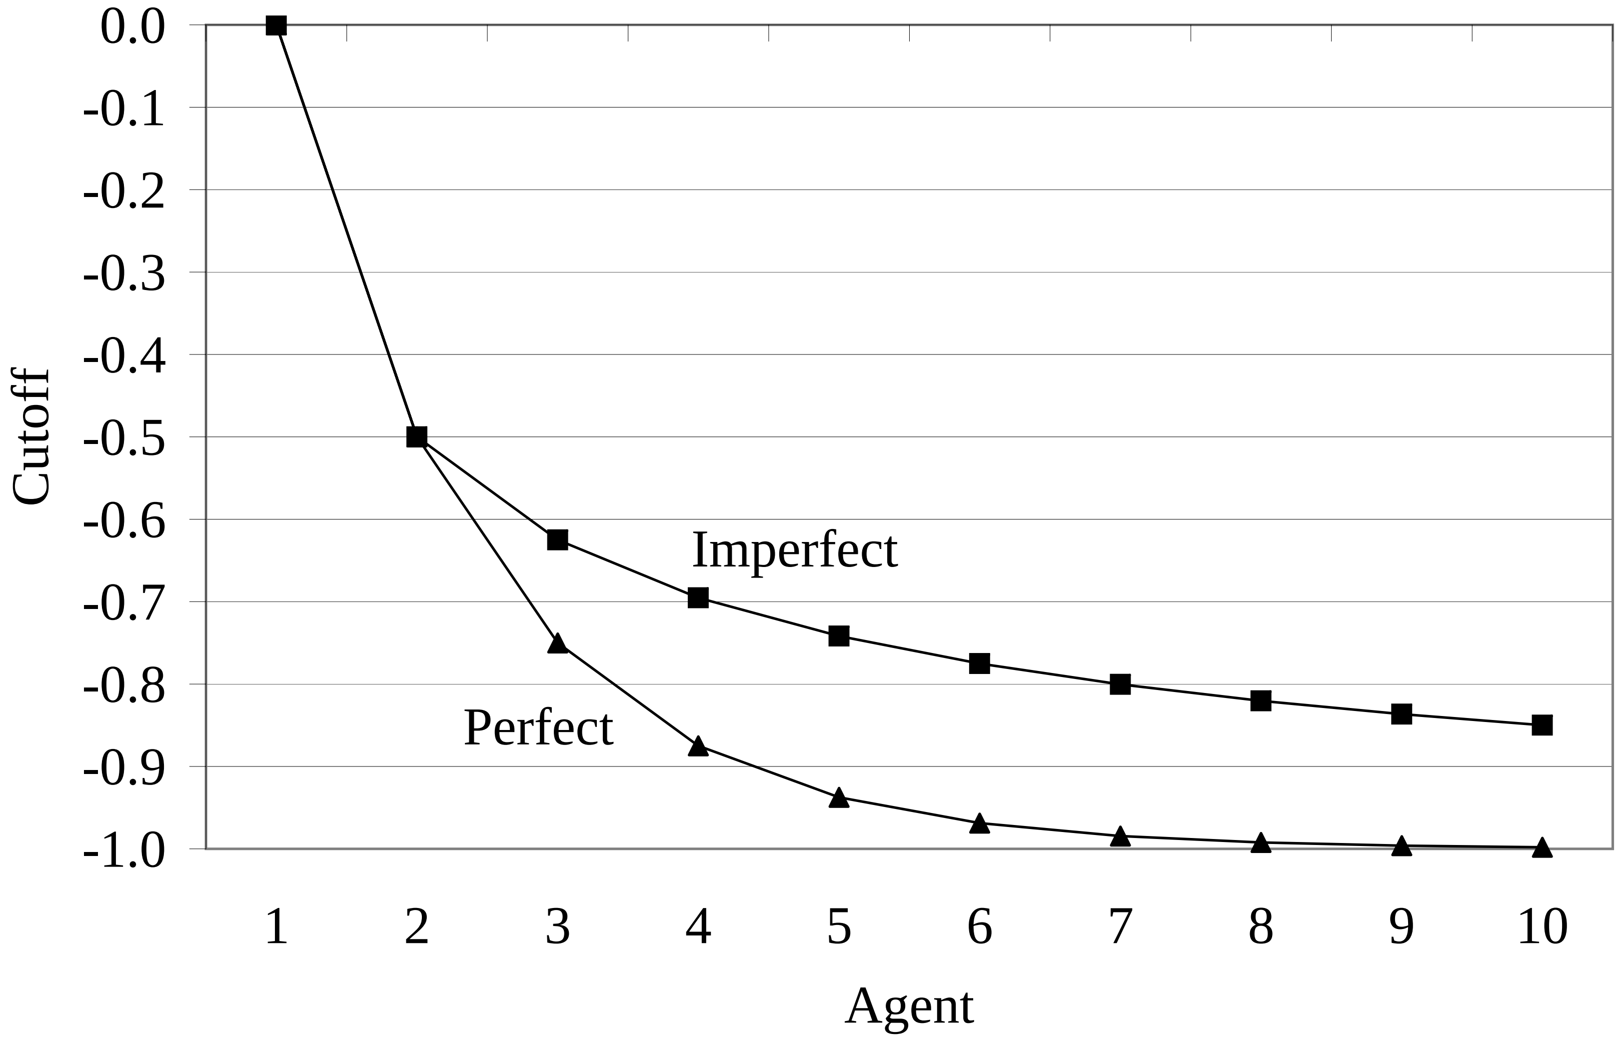

A sequence of cutoffs under imperfect and perfect information

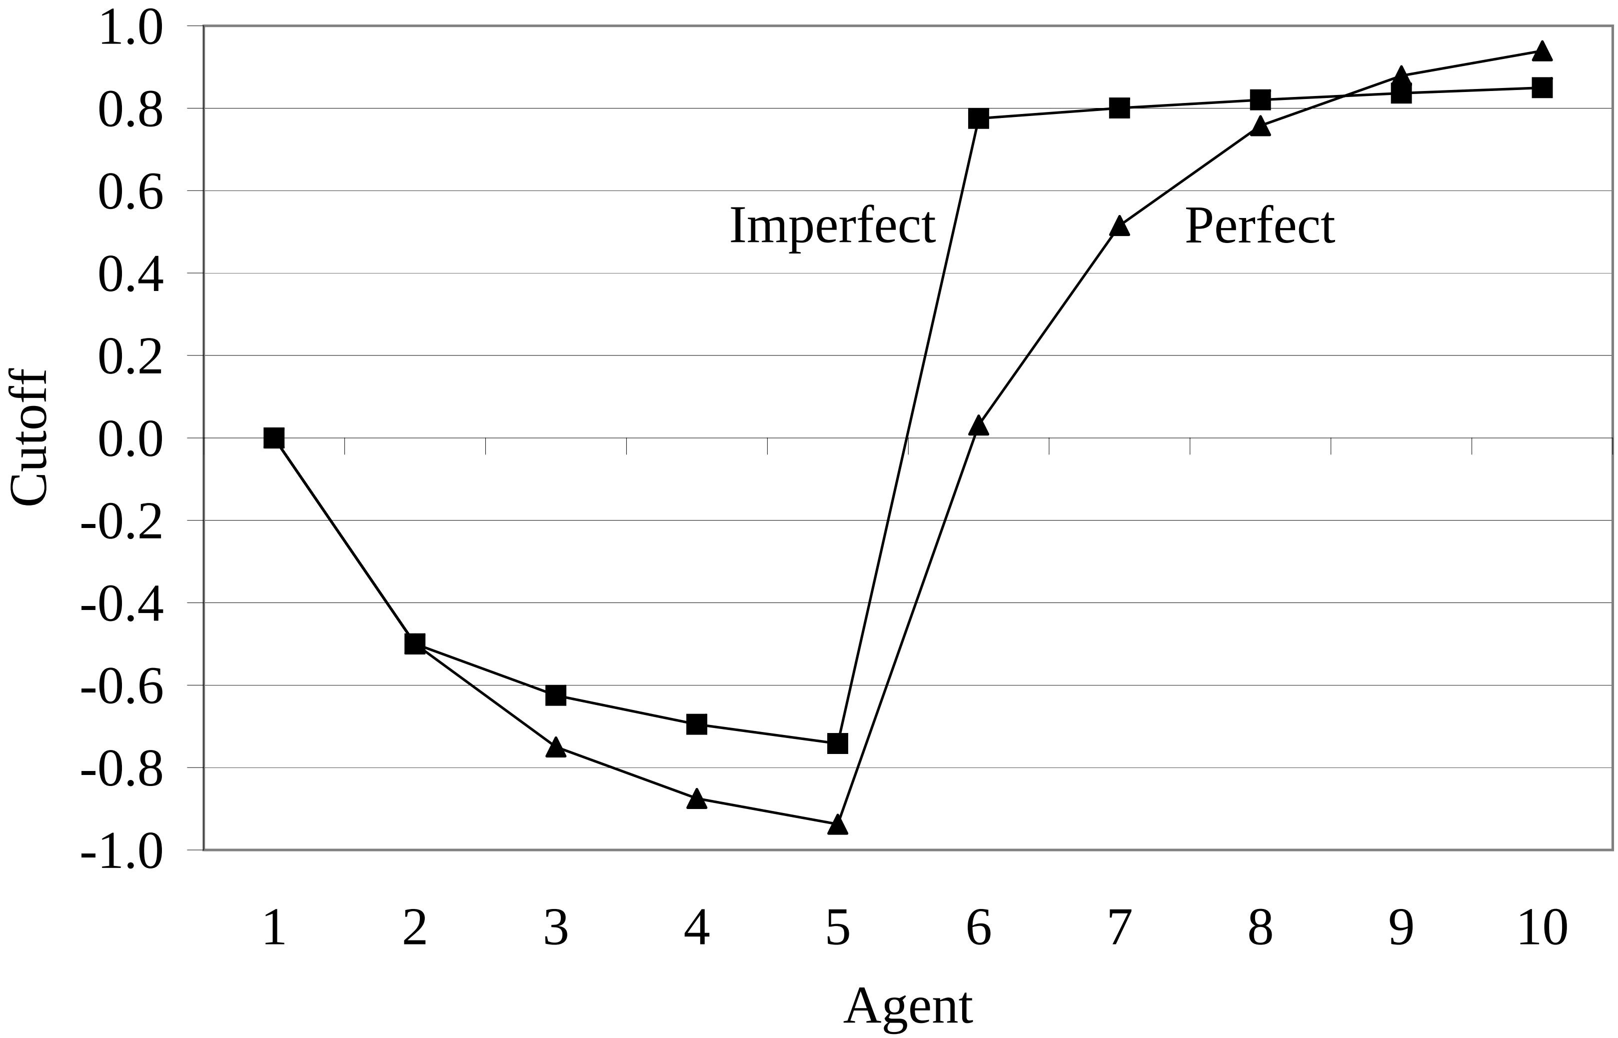

The decision problem

– 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*

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

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

<u>Perfect information</u>

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

<u>Imperfect information</u>

– 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

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

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

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

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 $x_{n-1} = 0$

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

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

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

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

which simplifies to

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

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

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

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.

Nash bargaining
(the axiomatic approach)

Bargaining

Nash's (1950) work is the starting point for formal bargaining theory.

The bargaining problem consists of

- a set of utility pairs that can be derived from possible agreements, and
- a pair of utilities which is designated to be a disagreement point.

Bargaining solution

The bargaining solution is a function that assigns a unique outcome to every bargaining problem.

Nash's bargaining solution is the first solution that

- satisfies four plausible conditions, and
- has a simple functional form, which make it convenient to apply.

A bargaining situation

A bargaining situation:

- N is a set of players or bargainers,
- A is a set of agreements/outcomes,
- D is a disagreement outcome, and

$\langle S, d \rangle$ is the primitive of Nash's bargaining problem where

- $S = (u_1(a), u_2(a))$ for $a \in A$ the set of all utility pairs, and $d = (u_1(D), u_2(D))$.

A bargaining problem is a pair $\langle S, d\rangle$ where $S \subset \mathbb{R}^2$ is compact and convex, $d \in S$ and there exists $s \in S$ such that $s_i > d_i$ for $i = 1, 2$. The set of all bargaining problems $\langle S, d\rangle$ is denoted by B.

A bargaining solution is a function $f : B \rightarrow \mathbb{R}^2$ such that f assigns to each bargaining problem $\langle S, d\rangle \in B$ a unique element in S.

Nash's axioms

One states as axioms several properties that it would seem natural for the solution to have and then one discovers that the axioms actually determine the solution uniquely - Nash 1953 -

Does not capture the details of a specific bargaining problem (e.g. alternating or simultaneous offers).

Rather, the approach consists of the following four axioms: invariance to equivalent utility representations, symmetry, independence of irrelevant alternatives, and (weak) Pareto efficiency.

Invariance to equivalent utility representations (INV)

$\langle S', d' \rangle$ is obtained from $\langle S, d \rangle$ by the transformations

$$s_i' \mapsto \alpha_i s_i + \beta_i$$

for $i = 1, 2$ if

$$d_i' = \alpha_i d_i + \beta_i$$

and

$$S' = \{(\alpha_1 s_1 + \beta_1, \alpha_2 s_2 + \beta_2) \in \mathbb{R}^2 : (s_1, s_2) \in S\}.$$

Note that if $\alpha_i > 0$ for $i = 1, 2$ then $\langle S', d' \rangle$ is itself a bargaining problem.

If $\langle S', d' \rangle$ is obtained from $\langle S, d \rangle$ by the transformations

$$s_i \mapsto \alpha_i s_i + \beta_i$$

for $i = 1, 2$ where $\alpha_i > 0$ for each i, then

$$f_i(S', d') = \alpha_i f_i(S, d) + \beta_i$$

for $i = 1, 2$. Hence, $\langle S', d' \rangle$ and $\langle S, d \rangle$ represent the same situation.

Symmetry (SYM)

A bargaining problem $\langle S, d\rangle$ is symmetric if $d_1 = d_2$ and $(s_1, s_2) \in S$ if and only if $(s_2, s_1) \in S$. If the bargaining problem $\langle S, d\rangle$ is symmetric then

$$f_1(S, d) = f_2(S, d)$$

Nash does not describe differences between the players. All asymmetries (in the bargaining abilities) must be captured by $\langle S, d\rangle$.

Hence, if players are the same the bargaining solution must assign the same utility to each player.

Independence of irrelevant alternatives (IIA)

If $\langle S, d\rangle$ and $\langle T, d\rangle$ are bargaining problems with $S \subset T$ and $f(T, d) \in S$ then

$$f(S, d) = f(T, d)$$

If T is available and players agree on $s \in S \subset T$ then they agree on the same s if only S is available.

IIA excludes situations in which the fact that a certain agreement is available influences the outcome.

Weak Pareto efficiency (WPO)

If $\langle S, d\rangle$ is a bargaining problem where $s \in S$ and $t \in S$, and $t_i > s_i$ for $i = 1, 2$ then $f(S, d) \neq s$.

In words, players never agree on an outcome s when there is an outcome t in which both are better off.

Hence, players never disagree since by assumption there is an outcome s such that $s_i > d_i$ for each i.

$\underline{SYM \text{ and } WPO}$

restrict the solution on single bargaining problems.

$\underline{INV \text{ and } IIA}$

requires the solution to exhibit some consistency across bargaining problems.

Nash 1953: there is precisely one bargaining solution, denoted by $f^N(S, d)$, satisfying SYM, WPO, INV and IIA.

Nash's solution

The unique bargaining solution $f^N : B \to \mathbb{R}^2$ satisfying SYM, WPO, INV and IIA is given by

$$f^N(S, 0) = \underset{(s_1,s_2)\in S}{\arg\max} s_1 s_2$$

The solution is the utility pair that maximizes the product of the players' utilities.

Proof

Pick a compact and convex set $S \subset \mathbb{R}^2_+$ where $S \cap \mathbb{R}^2_{++} \neq \emptyset$.

Step 1: f^N is well defined.

- Existence: the set S is compact and the function $f = s_1 s_2$ is continuous.
- Uniqueness: f is strictly quasi-conacave on S and the set S is convex.

Step 2: f^N is the only solution that satisfies SYM, WPO, INV and IIA.

Suppose there is another solution f that satisfies SYM, WPO, INV and IIA.

Let

$$S' = \{(\frac{s_1}{f_1^N(S)}, \frac{s_2}{f_2^N(S)}) : (s_1, s_2) \in S\}$$

and note that $s_1' s_2' \leq 1$ for any $s' \in S'$, and thus $f^N(S', 0) = (1, 1)$.

Since S' is bounded we can construct a set T that is symmetric about the 45° line and contains S'

$$T = \{(a, b) : a + b \leq 2\}$$

By WPO and SYM we have $f(T, 0) = (1, 1)$, and by IIA we have $f(S', 0) = f(T, 0) = (1, 1)$.

By INV we have that $f(S', 0) = f^N(S', 0)$ if and only if $f(S, 0) = f^N(S, 0)$ which completes the proof.