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

Social learning and bargaining (axiomatic approach)

Block 4
Jul 31 and Aug 1, 2015

Auction results

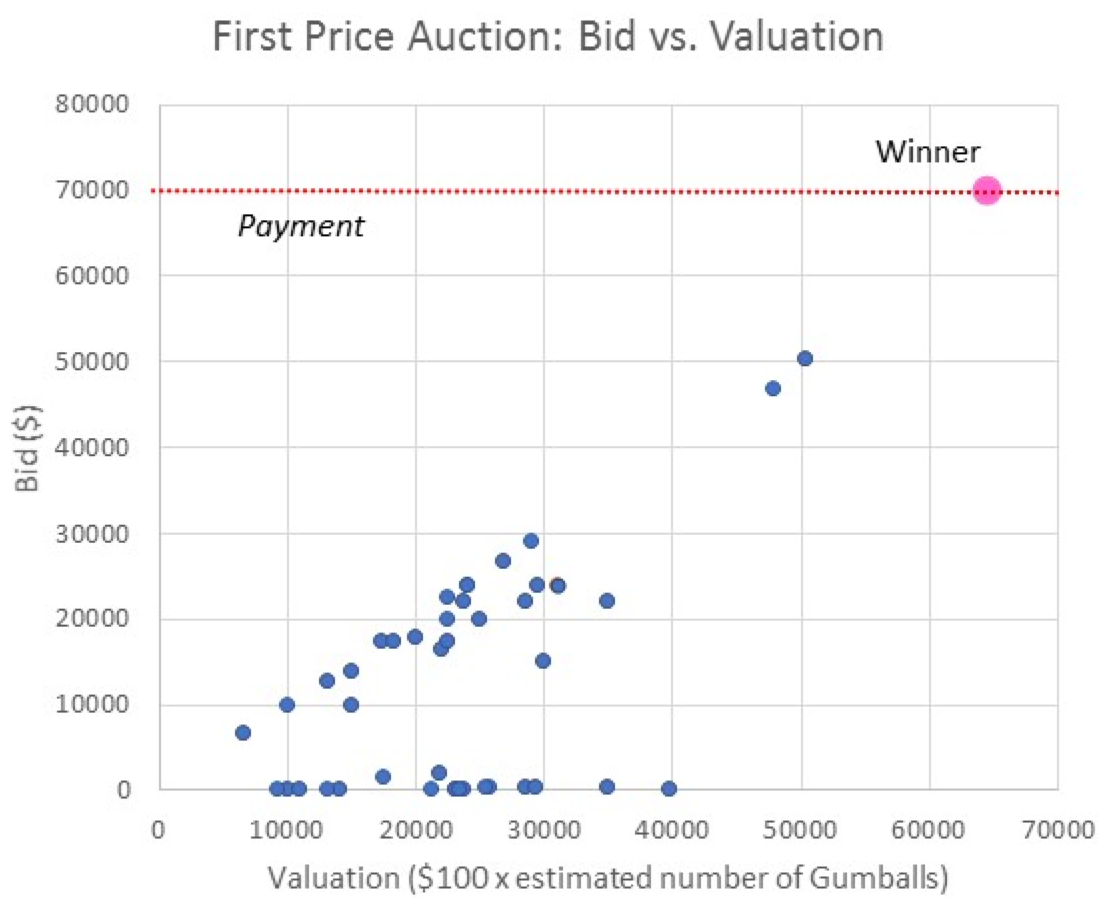
First Price Auction: Bid vs. Valuation
Winner
Payment
Bid ($)
Valuation ($100 x estimated number of Gumballs)
80000
70000
60000
50000
40000
30000
20000
10000
0
0
10000
20000
30000
40000
50000
60000
70000

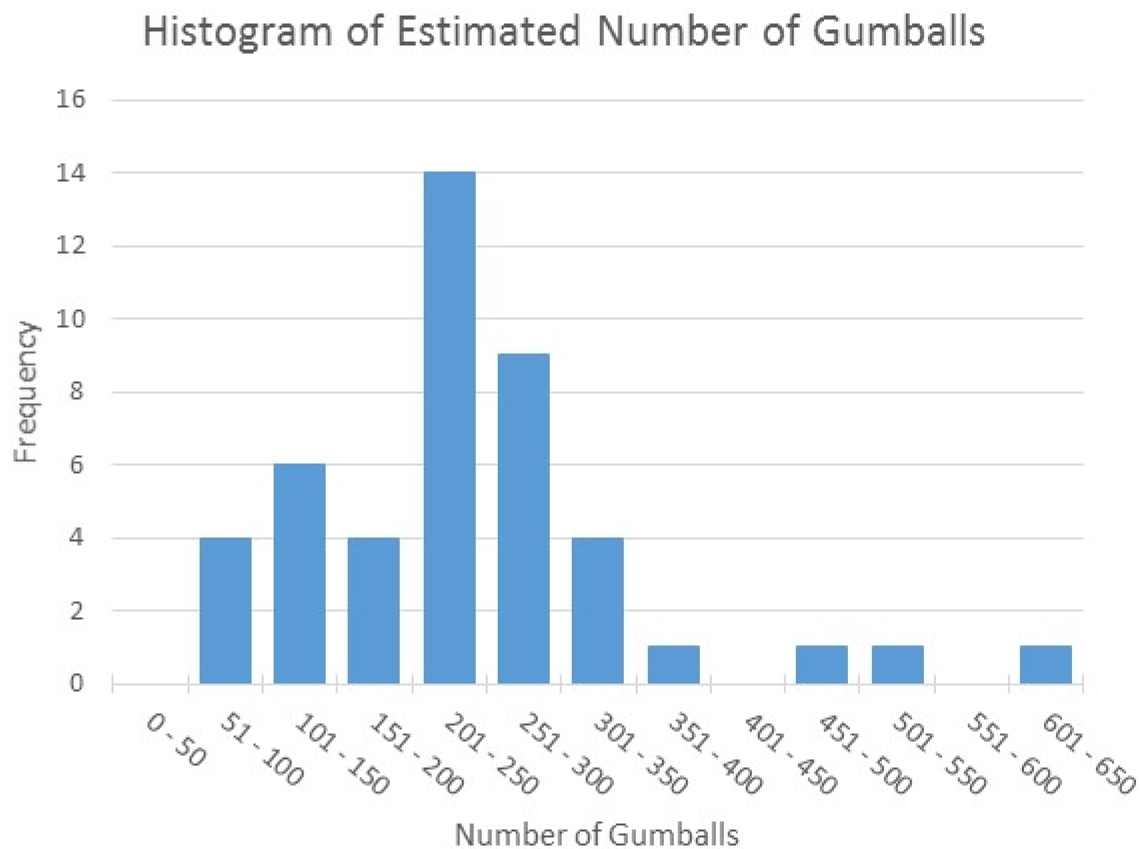
Histogram of Estimated Number of Gumballs
Frequency
0
2
4
6
8
10
12
14
16
0 - 50
51 - 100
101 - 150
151 - 200
201 - 250
251 - 300
301 - 350
351 - 400
401 - 450
451 - 500
501 - 550
551 - 600
601 - 650
Number of Gumballs


First vs. Second Price Auction Bids
Bid in Second Price Auction
Bid in First Price Auction
0
10000
20000
30000
40000
50000
60000
70000
80000

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

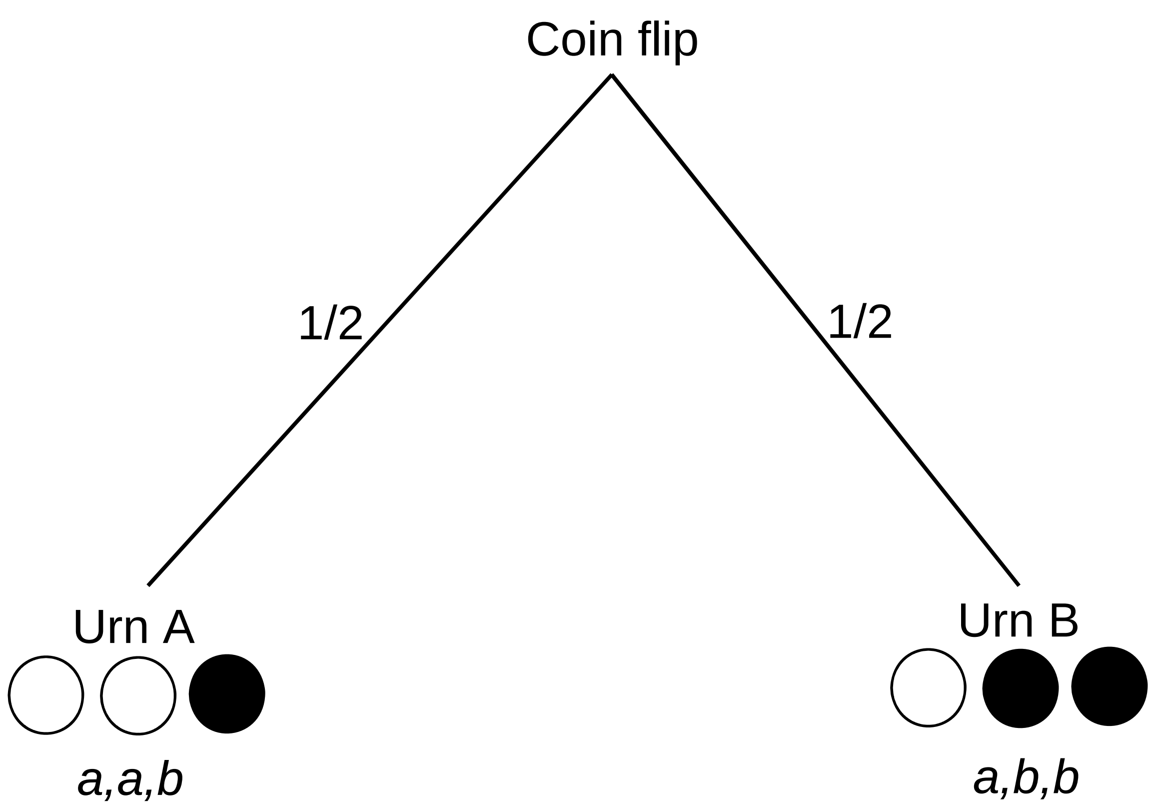
Coin flip
1/2
1/2
Urn A
a,a,b
Urn B
a,b,b

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:

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

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?

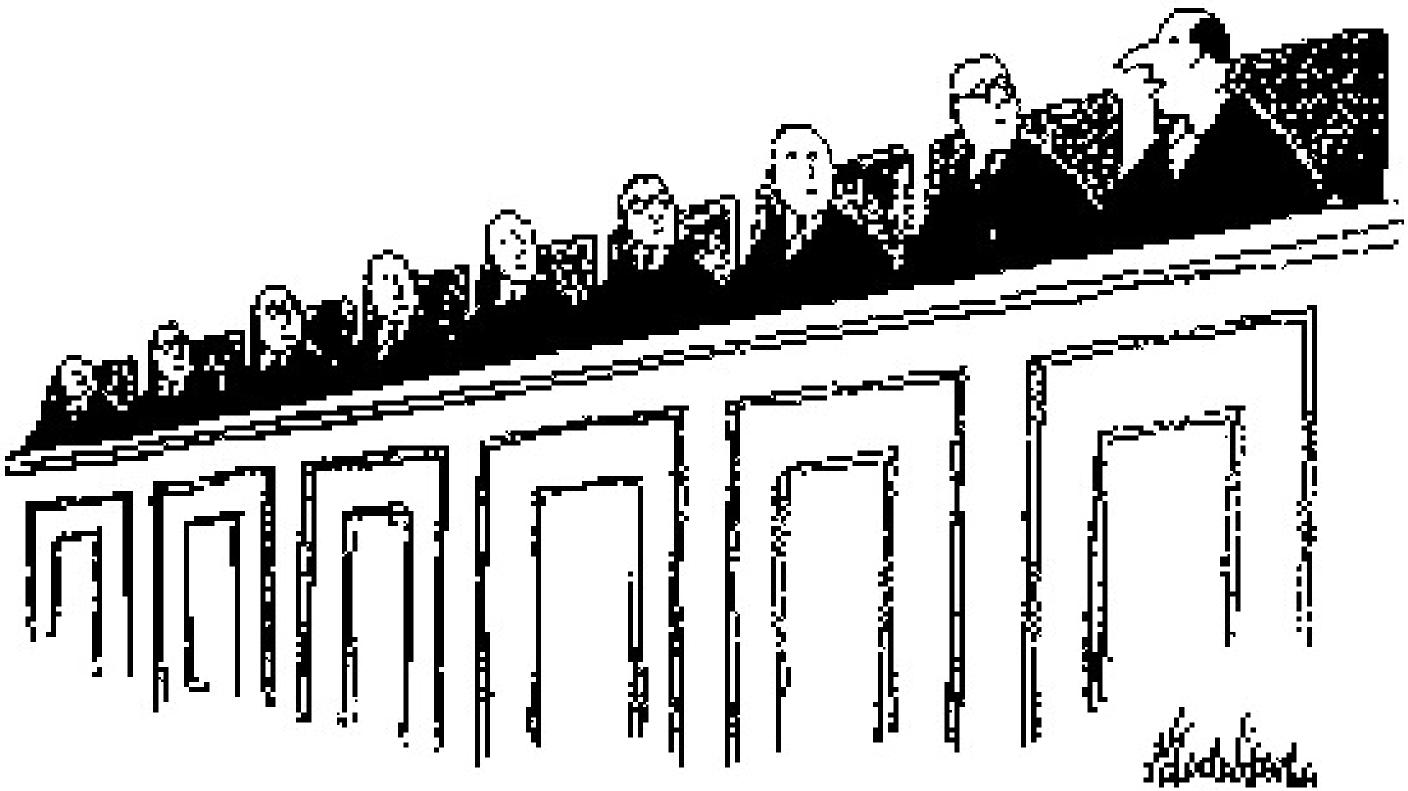

Imperfect information:
Which way is the wind blowing?!

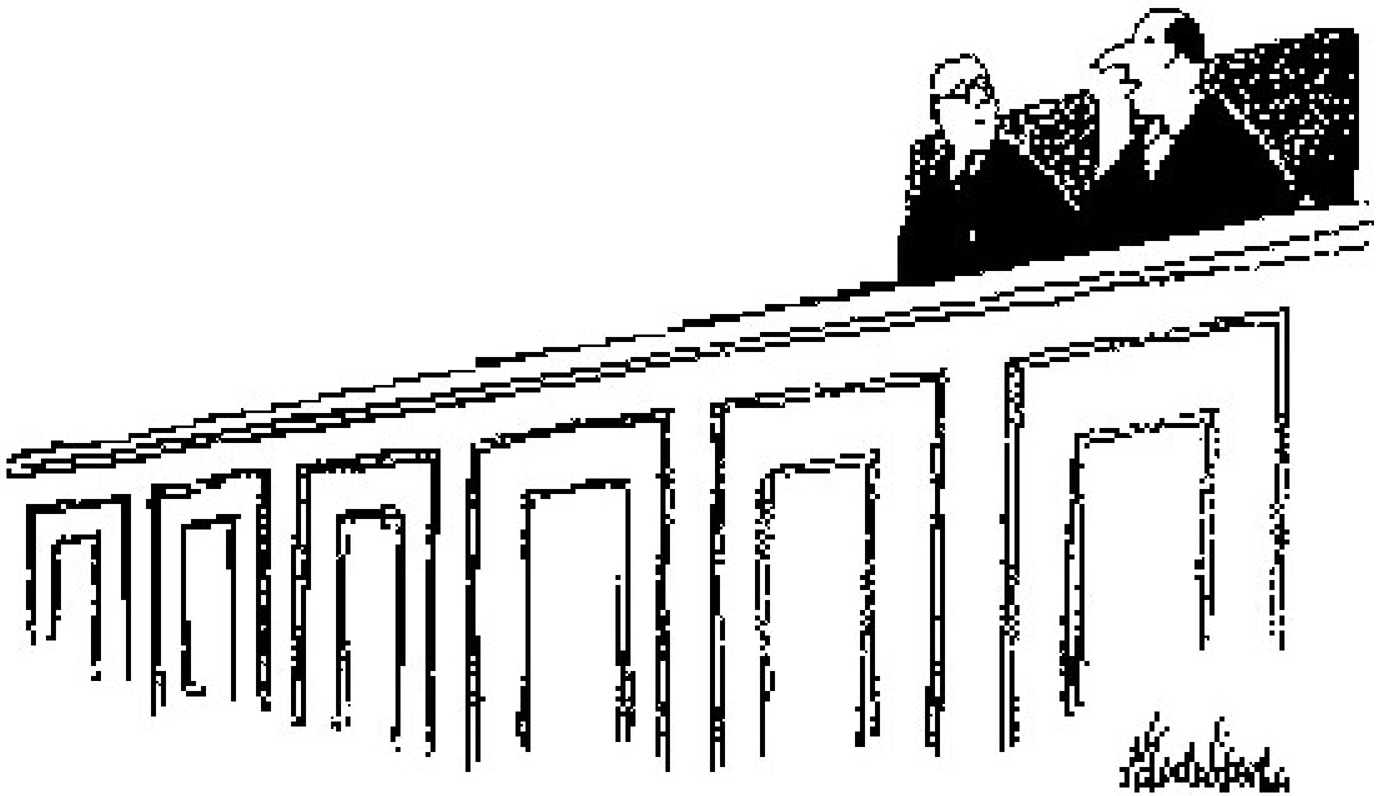

A three-agent example

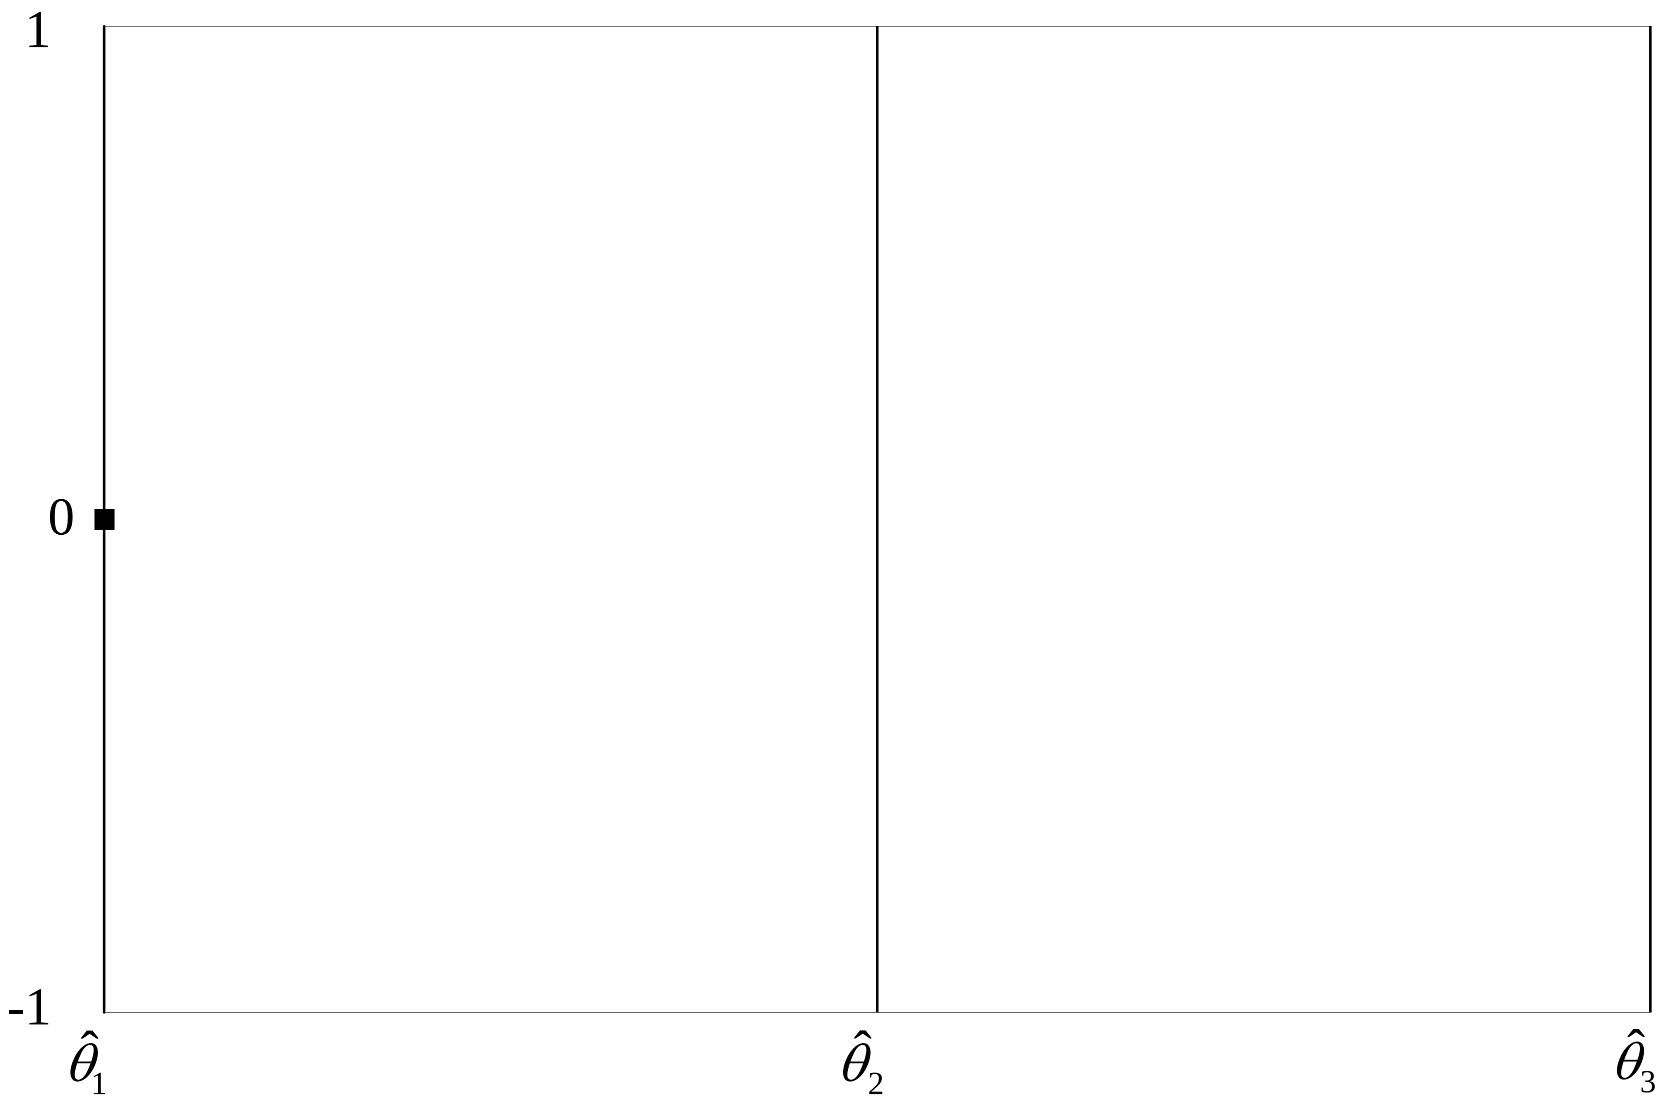

A three-agent example

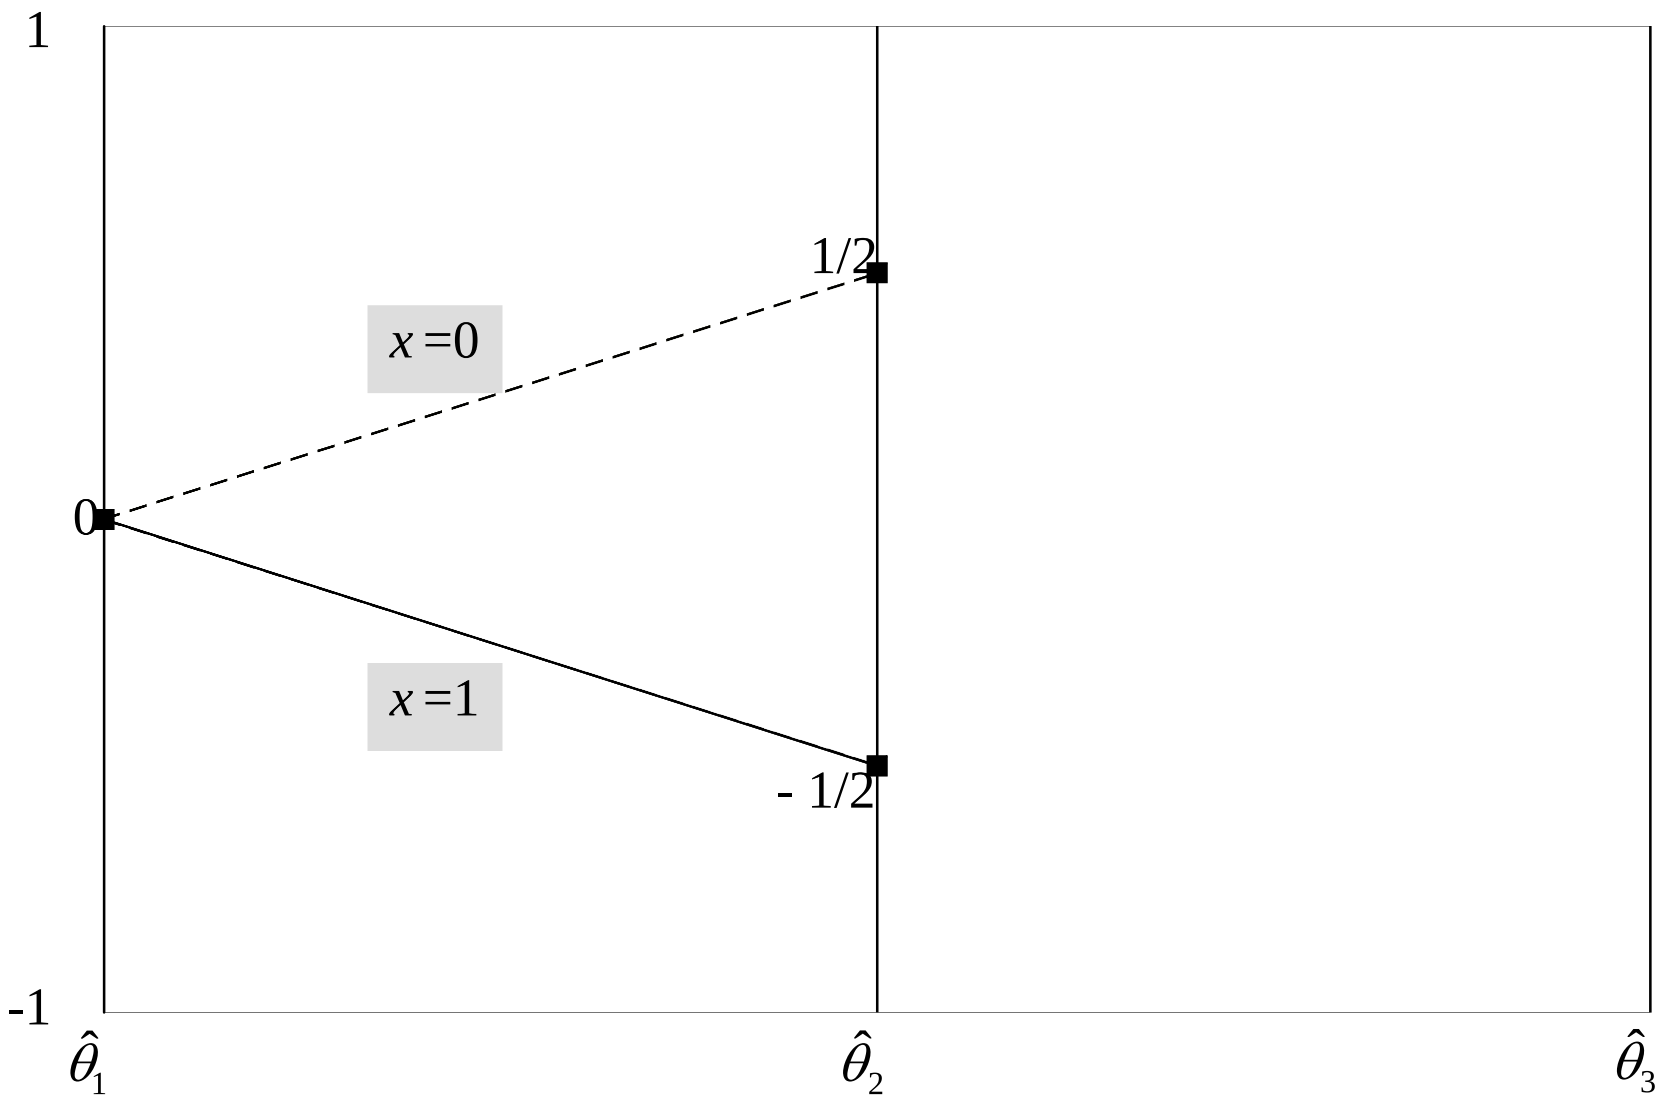

A three-agent example under perfect information

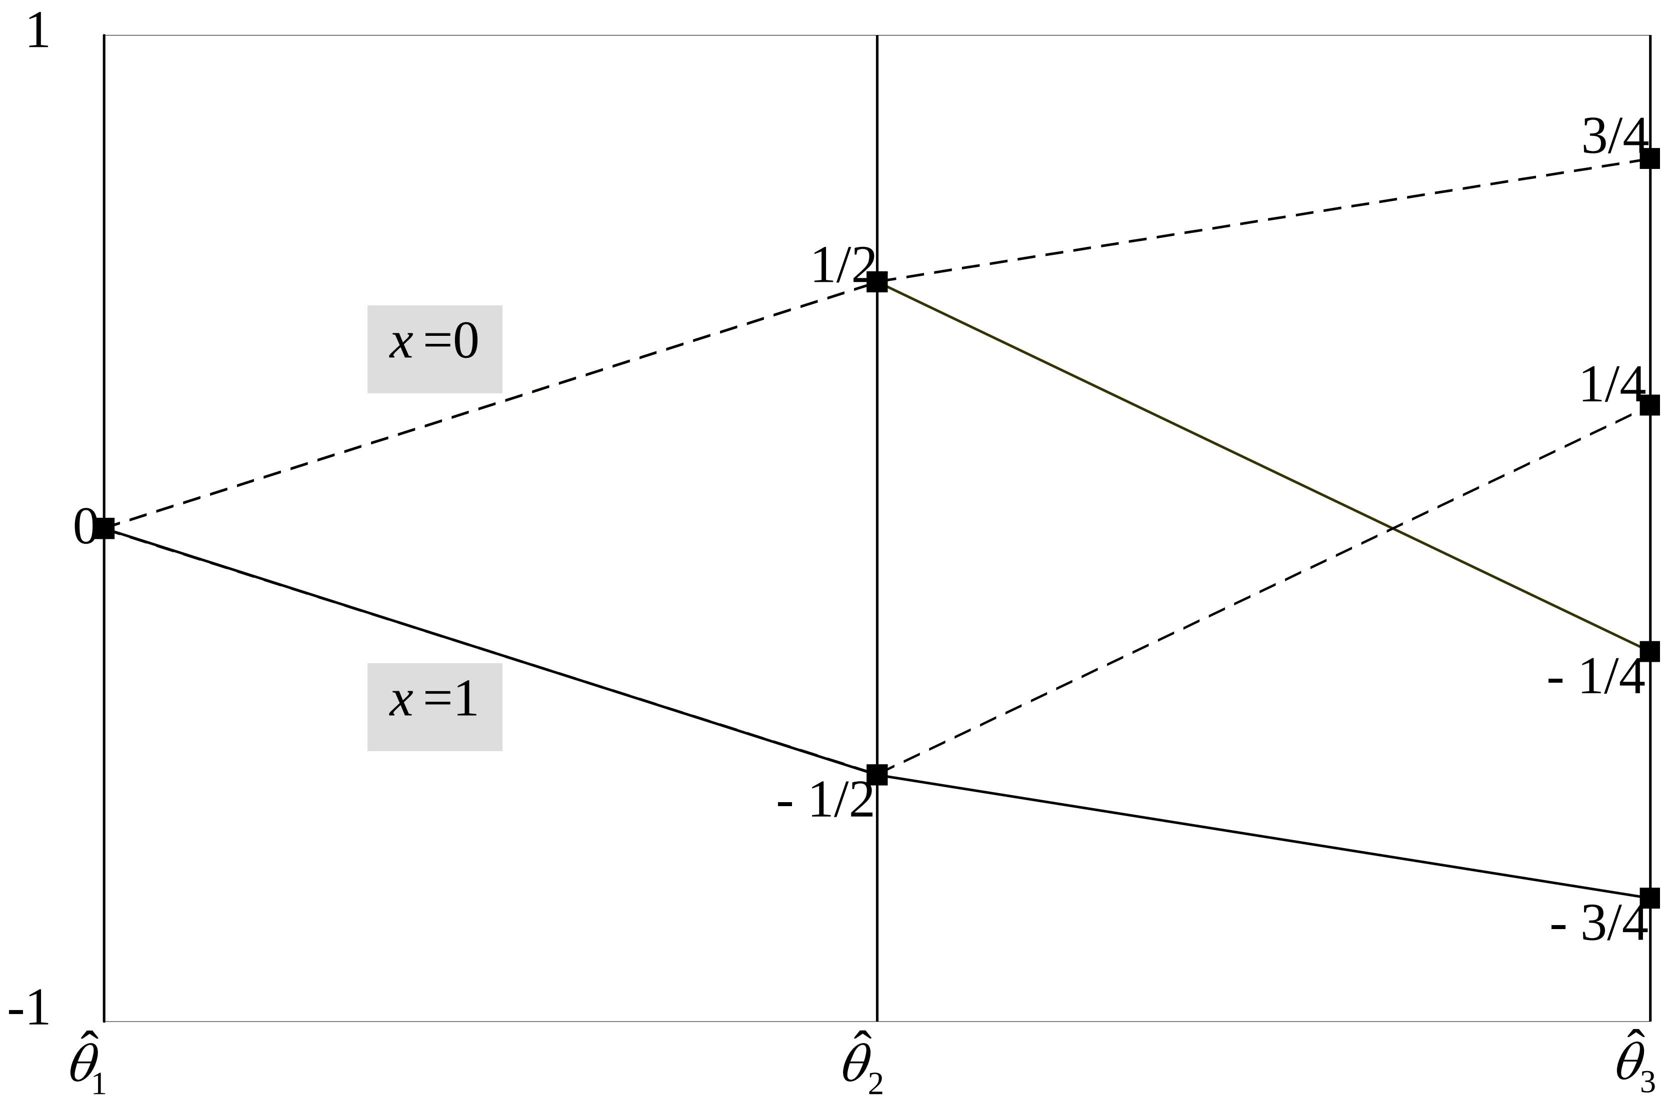

A three-agent example under imperfect information

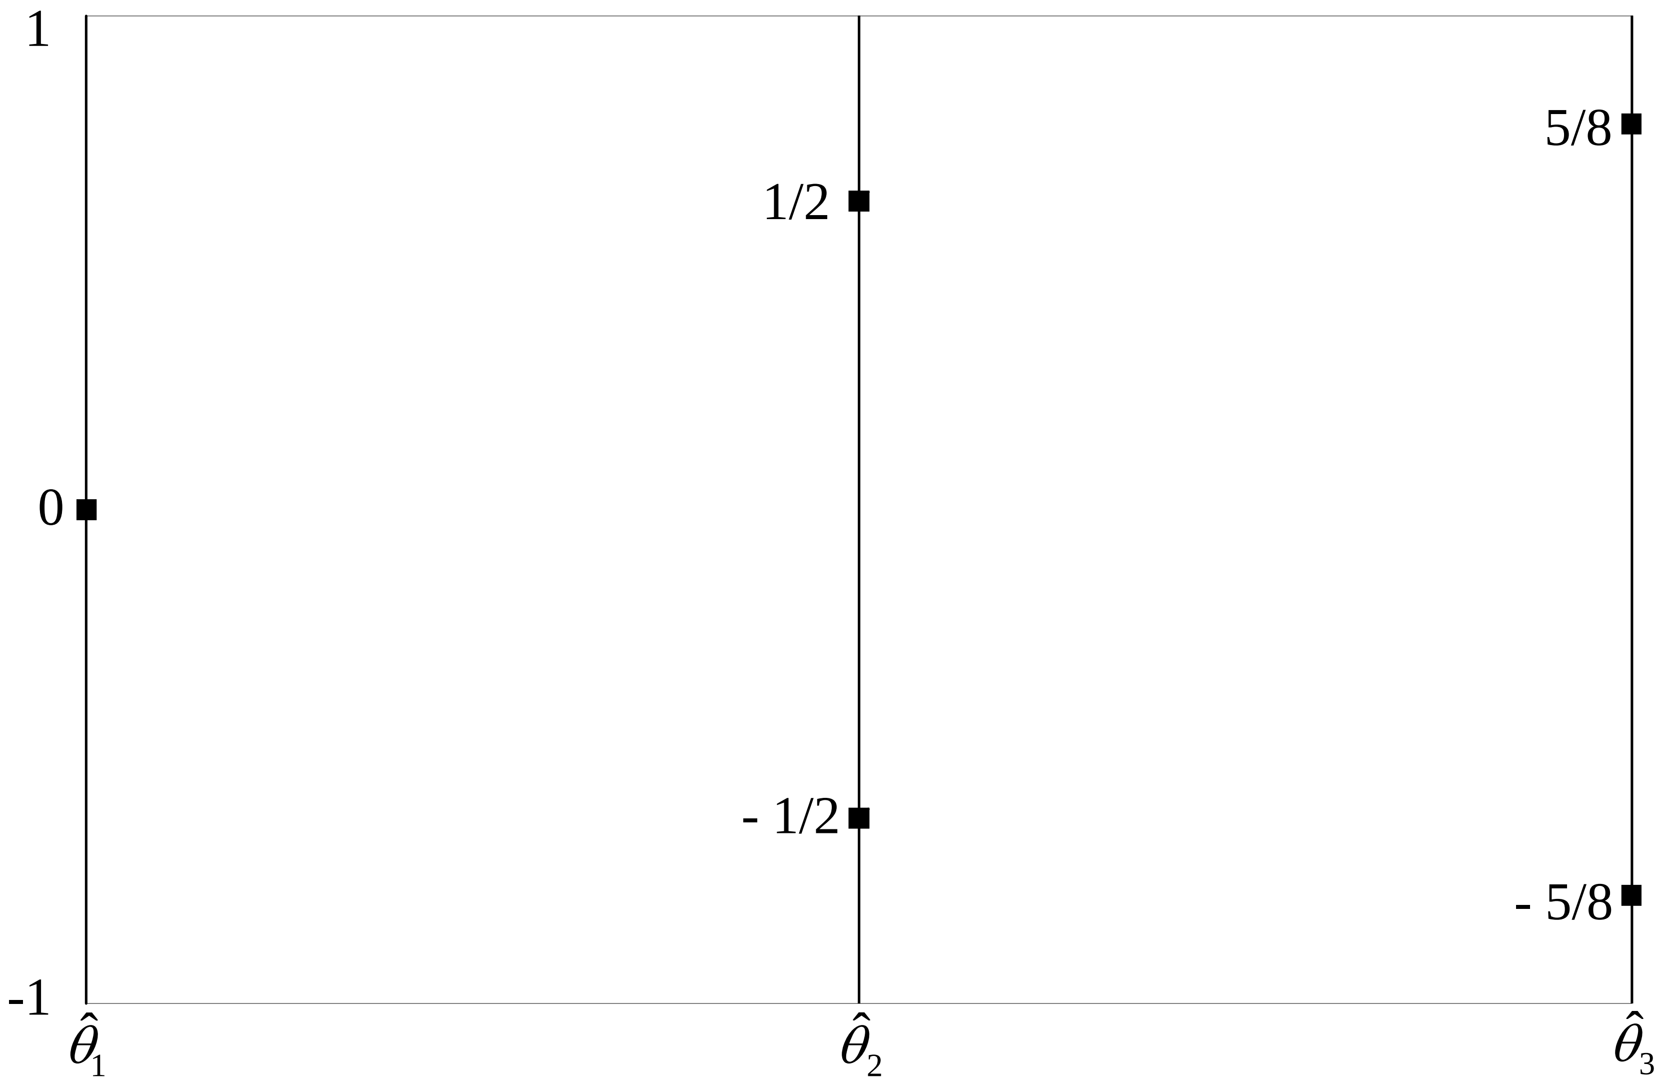

A sequence of cutoffs under imperfect and perfect information

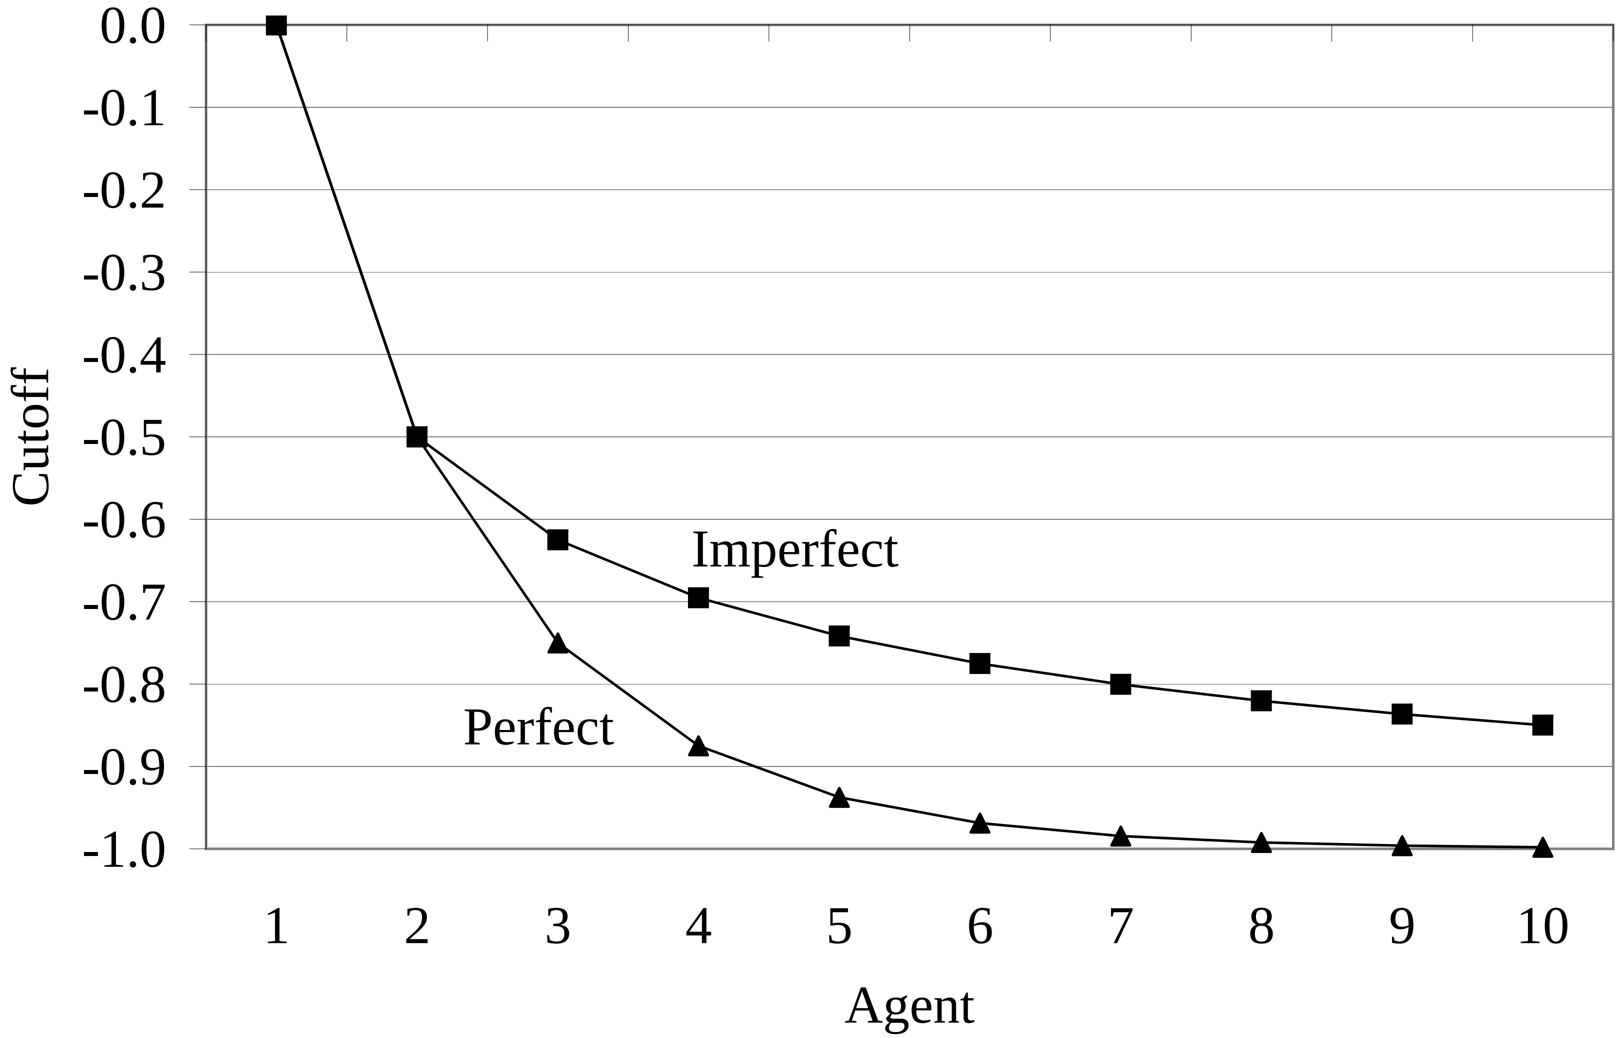

A sequence of cutoffs under imperfect and perfect information

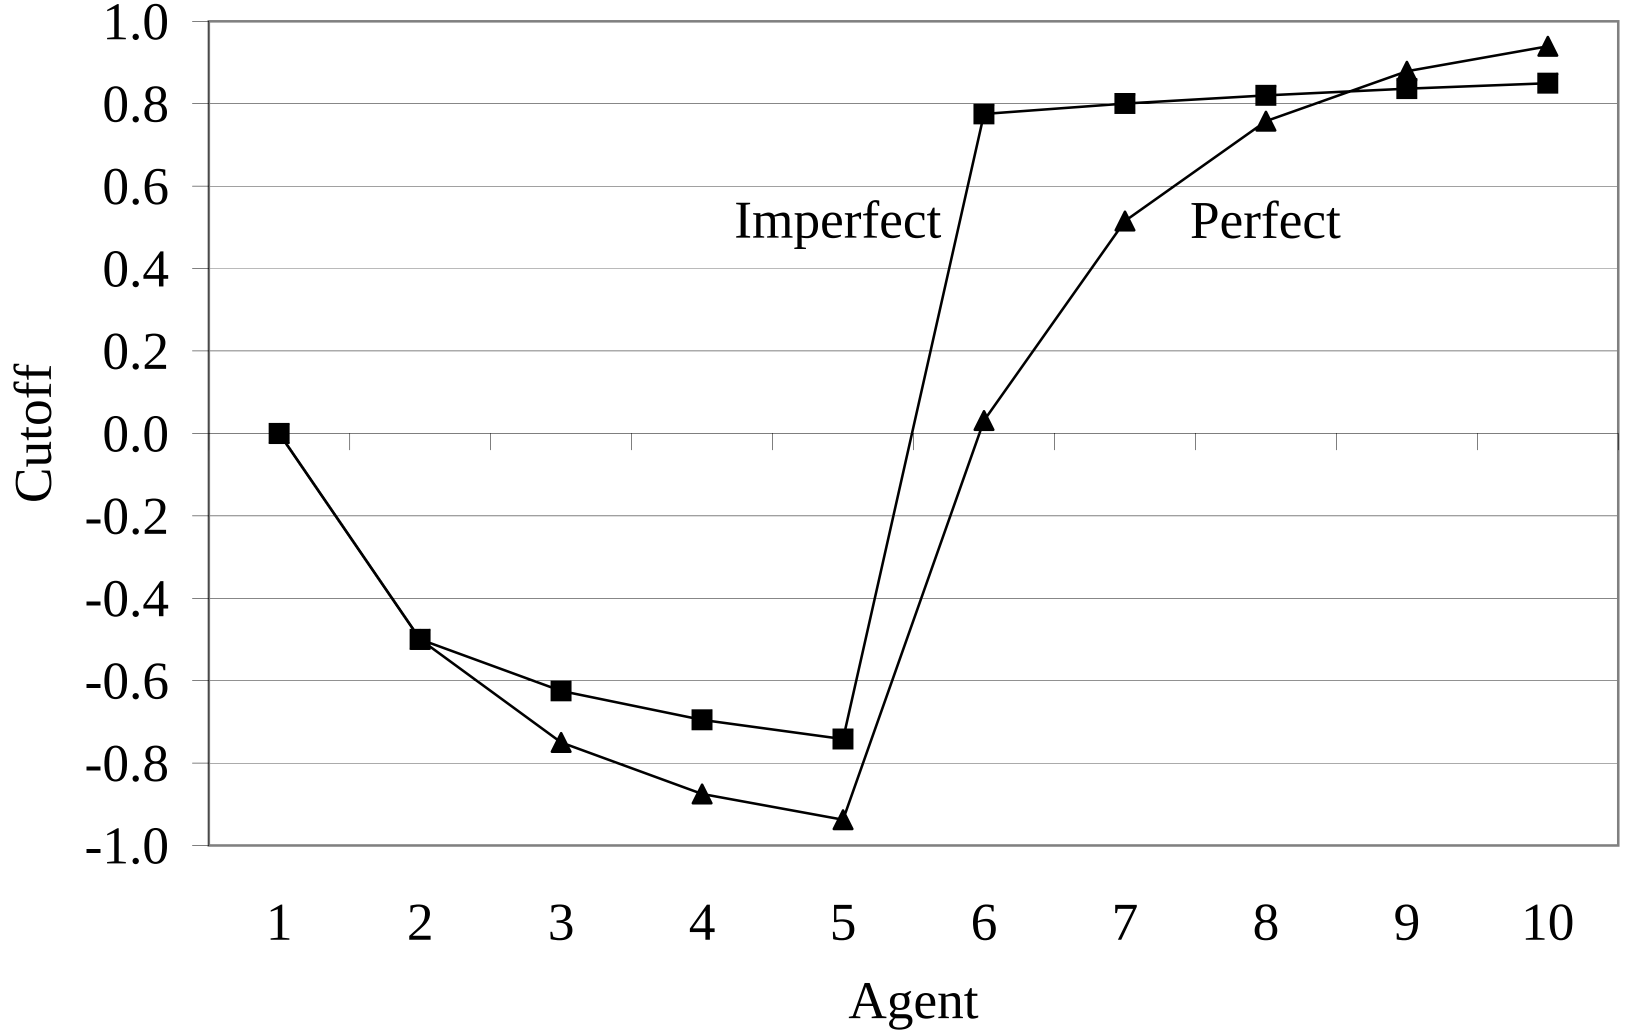

The decision problem

– 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

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

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 <u>bargaining problem</u> 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 <u>bargaining solution</u> 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.