

**UC Berkeley
Haas School of Business
Game Theory
(EMBA 296 & EWMBA 211)
Spring 2023**

**Back to decision theory...
Review of extensive games, and games w/ imperfect information**

**Block III
Mar 8-10, 2023**


Risk preferences

Life is full of lotteries :-)


$$x := \begin{array}{l} \nearrow^p \quad \$A \\ \searrow \quad \$B \\ 1-p \end{array}$$


$$y := \begin{array}{l} \nearrow^p \quad \$A \\ \xrightarrow{q} \quad \$B \\ \searrow \quad \$C \\ 1-p-q \end{array}$$


A risky lottery (left) and an ambiguous lottery (right)



A compounded lottery



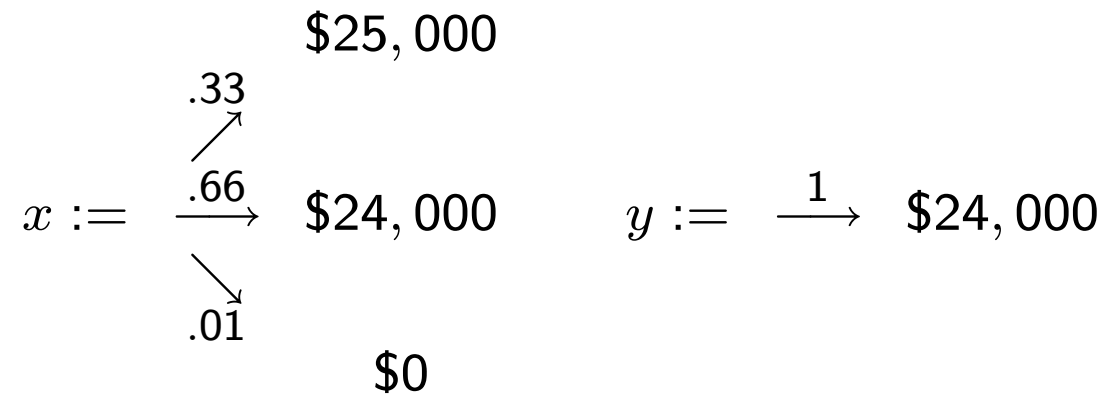




von Neumann and Morgenstern Expected Utility Theory (EUT)



Allais (1953) I

– Choose between the two gambles:




Allais (1953) II

- Choose between the two gambles:




The (Marschak-Machina) probability triangle




Consider three monetary payouts H , M , and L where $H > M > L$

Risk profiling




A “complete” risk profiling requires knowing all possible comparisons like between **A** and **B**.

A topographic map




An indifference map of a loss-neutral (expected utility) individual




Expected Utility Theory (EUT) requires that indifference lines are parallel

A test of Expected Utility Theory (EUT)



EUT requires that indifference lines are parallel so one must choose either **A** and **C**, or **B** and **D**.


Loss neutral and more risk tolerant




Mr. **Green** is more risk tolerant than Mr. **Blue** who is more risk tolerant than Mr. **Red**. The gentlemen are loss neutral.


Putting (risk) preferences under the microscope


The decision problem




Some “fingerprints” of individual behaviors











Who is Homo-Economicus?

The construction of a Homo-Economicus score



Homo Economicus: equiprobable lotteries





Wealth differentials

- ⇒ The heterogeneity in wealth is not well-explained either by standard observables (income, education, family structure) or by standard unobservables (intertemporal substitution, risk tolerance).
- ⇒ If consistency with utility maximization in the experiment is a good proxy for (financial) *DMQ* then the degree to which consistency differ across subjects should help explain wealth differentials.

The relationship between CCEI scores and wealth

	(1)	(2)	(3)
CCEI	1.351** (0.566)	1.109** (0.534)	101888.0* (52691.9)
Log 2008 household income	0.584*** (0.132)	0.606*** (0.126)	
2008 household income			1.776*** (0.4)
Female	-0.313* (0.177)	-0.356** (0.164)	-32484.3* (17523.9)
Partnered	0.652*** (0.181)	0.595*** (0.171)	46201.9*** (17173.7)
# of children	0.090 (0.093)	0.109 (0.086)	14078.6* (8351.5)
Age	Y	Y	Y
Education	Y	Y	Y
Occupation	Y	Y	Y
Constant	6.292 (6.419)	0.469 (3.598)	76214.4 (559677.5)
R^2	0.179	0.217	0.188
# of obs.	517	566	568

Ambiguity Aversion



- I have an urn with 1000 balls in it. Some of the balls are red and some are blue. All the balls are the same size and weight, and they are not distinguishable in any way except in color.
- I will let you reach into this urn without looking and draw out one of these balls. I want to know your preferences between investments (gambles) based on the outcome of this random event.

[1] There are precisely 500 red balls in the urn. Would you rather (fill the blank)

(a) get \$_____ for sure

(b) get \$100,000 if the ball drawn from the urn is red and \$0 if it is blue.

[1] There are precisely 500 red balls in the urn. Would you rather (fill the blank)

(a) get \$_____ for sure

(b) get \$100,000 if the ball drawn from the urn is red and \$0 if it is blue.

[2] There are precisely 500 red balls in the urn. Would you rather (fill the blank)

(a) lose \$_____ for sure

(b) lose \$100,000 if the ball drawn from the urn is red and \$0 if it is blue.

[1] There are precisely 500 red balls in the urn. Would you rather (fill the blank)

(a) get \$_____ for sure

(b) get \$100,000 if the ball drawn from the urn is red and \$0 if it is blue.

[2] There are precisely 500 red balls in the urn. Would you rather (fill the blank)

(a) lose \$_____ for sure

(b) lose \$100,000 if the ball drawn from the urn is red and \$0 if it is blue.

[3] The number of red balls in the urn is unknown. Would you rather (fill the blank)


(a) get \$_____ for sure


(b) get \$100,000 if the ball drawn from the urn is red and \$0 if it is blue.

Review of extensive games w/ 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).

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.



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



Oligopoly

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 $A > c$.

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,

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

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



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

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

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



Games with imperfect (and asymmetric) information

Markets with asymmetric information

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

⇒ This is clearly a drastic simplification!!!

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

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



The Market for Lemons

Example I

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

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

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

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

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

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

Example II

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

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

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

Signaling

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

Example – educational signaling

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

$$L_1 a_1 + L_2 a_2$$

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

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

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

to competent workers, respectively.

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

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

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

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

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

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

$$c_1e \text{ and } c_2e$$

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.

Social Learning
Herd behavior and informational cascades

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

Examples

Business strategy

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

Finance

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

Politics

- The solid New 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.

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

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


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*

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

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

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

which simplifies to

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

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.