

**UC Berkeley
Haas School of Business
Game Theory
(EMBA 296 & EWMBA 211)
Summer 2017**

Signaling, herding and bargaining

**Block 4
Jul 20-22, 2017**

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.”

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

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.

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.

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.

**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 .

SYM and *WPO*

restrict the solution on single bargaining problems.

INV 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 $f^N : B \rightarrow \mathbb{R}^2$ satisfying *SYM*, *WPO*, *INV* and *IIA* is given by

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

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-concave 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

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

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.