

The background features a light gray technical drawing of a mechanical assembly, possibly a watch movement, with various gears, levers, and a grid pattern overlaid on the right side.

Economics of Networks

Incomplete Information and Introduction to Social Learning

Evan Sadler

Massachusetts Institute of Technology

Agenda

- Games with incomplete information
- Bayes-Nash Equilibrium
- Extensive form games
- Perfect Bayesian Equilibrium
- Rational Herding

Reading: Osborne Chapter 9; EK Chapter 16

Incomplete Information

Strategic situations often involve uncertainty

- Uncertainty about others' preferences
- Uncertainty about others' available strategies
- Uncertainty about others' information

Possibility of learning affects incentives

Examples:

- Bargaining (how much is opponent willing to pay?)
- Auctions (how much do others value the object?)
- Market competition (what costs do my competitors face?)
- Social learning (What can I infer from others' choices?)

An Example

You (player 1) and a friend are trying to coordinate a meeting place (say, the mall or the library)

- Different preferences over the two options

	M	L
M	(2, 1)	(0, 0)
L	(0, 0)	(1, 2)

Two pure strategy equilibria, one mixed

Now, suppose you are unsure whether your friend wants to meet or avoid you

Example, continued

To model this, we assume your friend (player 2) has one of two possible *types*

- One type wants to meet, the other wants to avoid

Suppose the two types are equally likely, two possible payoff matrices

	M	L
M	(2, 1)	(0, 0)
L	(0, 0)	(1, 2)

	M	L
M	(2, 0)	(0, 2)
L	(0, 1)	(1, 0)

Your friend knows which game is played, but you do not

- What are the strategies?

Example, continued

Idea: Use Nash Equilibrium in an expanded game

- Think of each type of player 2 as a separate player

Equivalently, form conjecture about behavior in each possible state, and maximize *expected utility* given the conjecture

Consider the profile $(M, (M, L))$

- Player 1 goes to the mall
- Player 2 goes to the mall in state 1 and the library in state 2

Clearly a best response for player 2

Example, continued

Now check for player 1

Since both states are equally likely,

$$\mathbb{E}[u_1(M, (M, L))] = \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$$

If player 1 chose L instead, the expected payoff is

$$\mathbb{E}[u_1(L, (M, L))] = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 1 = \frac{1}{2}$$

Hence, the profile $(M, (M, L))$ is a (Bayes) Nash Equilibrium

Example, continued

Note meeting at the library, player 2's preferred outcome, is no longer part of an equilibrium

- Consider the profile $(L, (L, M))$

Again, clearly a best response for player 2

A similar calculation shows player 1 earns $\frac{1}{2}$ by playing L , but earns 1 by deviating to M

The profile $(L, (L, M))$ is not a Bayes Nash Equilibrium

Bayesian Games

More formally...

Definition

A Bayesian game consists of

- A set of players N
- A set of actions (pure strategies) S_i for each player i
- A set of types Θ_i for each player i
- A payoff function $u_i(s, \theta)$ for each player i
- A (joint) probability distribution $p(\theta_1, \theta_2, \dots, \theta_n)$ over types

Note payoffs depend on vector of actions *and* vector of types

Bayesian Games, continued

We maintain the assumption that the probability distribution p is common knowledge

- Players agree on the prior probability of different type vectors

Very strong assumption, but very convenient

- Avoid having to deal with hierarchies of beliefs

Definition

A pure strategy for player i is a map $s_i : \Theta_i \rightarrow S_i$ prescribing an action for each type of player i .

Bayes' Rule

Recall types are drawn from the prior distribution $p(\theta_1, \theta_2, \dots, \theta_n)$

On observing one's own type θ_i , can compute conditional distribution $p(\theta_{-i} | \theta_i)$ via Bayes' rule

Player i evaluates expected payoff according to the conditional distribution:

$$U_i(s'_i, s_{-i}, \theta_i) = \int_{\Theta_{-i}} u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) dp(\theta_{-i} | \theta_i)$$

Bayes' Rule, continued

Quick review, basic definitions, events A and B :

- Probability of events $\mathbb{P}(A)$ and $\mathbb{P}(B)$
- Conditional probabilities $\mathbb{P}(A | B)$ and $\mathbb{P}(B | A)$
- Joint probability $\mathbb{P}(A \cap B)$

Definition of conditional probability:

$$\mathbb{P}(A | B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)}$$

Probability that A is true given B

If events are independent, then $\mathbb{P}(A \cap B) = \mathbb{P}(A) \cdot \mathbb{P}(B)$ and $\mathbb{P}(A | B) = \mathbb{P}(A)$

Bayes' Rule, continued

Can also express conditional probabilities in terms of one another

- Write A^c for the complement of A
- Have $\mathbb{P}(A^c) = 1 - \mathbb{P}(A)$

Using $\mathbb{P}(B) = \mathbb{P}(A) \cdot \mathbb{P}(B | A) + \mathbb{P}(A^c) \cdot \mathbb{P}(B | A^c)$, we have

$$\mathbb{P}(A | B) = \frac{\mathbb{P}(A) \cdot \mathbb{P}(B | A)}{\mathbb{P}(A) \cdot \mathbb{P}(B | A) + \mathbb{P}(A^c) \cdot \mathbb{P}(B | A^c)}$$

More generally, for any countable partition $\{A_i\}_{i=1}^n$, we have

$$\mathbb{P}(A_i | B) = \frac{\mathbb{P}(A_i) \cdot \mathbb{P}(B | A_i)}{\sum_{j=1}^n \mathbb{P}(A_j) \cdot \mathbb{P}(B | A_j)}$$

Bayes Nash Equilibrium

Definition (Bayes Nash Equilibrium)

The profile σ is a pure strategy Bayes Nash Equilibrium if for all $i \in N$ and all $\theta_i \in \Theta$, we have

$$\sigma_i(\theta_i) \in \arg \max_{s'_i \in S_i} \int_{\Theta_{-i}} u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) dp(\theta_{-i} | \theta_i)$$

Bayes Nash equilibrium is a Nash equilibrium of the expanded game in which player i 's pure strategies are maps from Θ_i to S_i

Existence of Bayes-Nash Equilibrium

Theorem

In any finite Bayesian game, a mixed strategy Bayes Nash equilibrium exists

Theorem

Consider a Bayesian game with continuous strategy spaces and types. If the strategy and type sets are compact, and payoff functions are continuous and concave in own strategies, then a pure strategy Bayes Nash equilibrium exists.

Proofs based on Kakutani's fixed point theorem are essentially identical to what we did a few weeks ago.

Example: Incomplete Information Cournot

Two firms produce an identical good at constant marginal cost

Inverse demand is $P(Q)$, where Q is the total quantity

Firm 1 has marginal cost c that is common knowledge

Firm 2's marginal cost is private information

- Cost c_l with probability θ and c_h with probability $1 - \theta$, $c_l < c_h$

Game has two players and two states (l and h), actions $q_i \in [0, \infty)$

Example: Incomplete Information Cournot

Payoffs for the two players are

$$u_1(q_1, q_2, t) = q_1(P(q_1 + q_2) - c)$$

$$u_2(q_1, q_2, t) = q_2(P(q_1 + q_2) - c_t),$$

where $t \in \{l, h\}$ is firm 2's type

Can think of strategy profile as a triple (q_1^*, q_l^*, q_h^*)

Three best response functions

$$B_1(q_l, q_h) = \arg \max_{q_1 \geq 0} q_1 (\theta P(q_1 + q_l) + (1 - \theta)P(q_1 + q_h) - c)$$

$$B_l(q_1) = \arg \max_{q_l \geq 0} q_l (P(q_1 + q_l) - c_l)$$

$$B_h(q_1) = \arg \max_{q_h \geq 0} q_h (P(q_1 + q_h) - c_h)$$

Example: Incomplete Information Cournot

Bayes Nash equilibria are triples (q_1^*, q_l^*, q_h^*) such that

$$B_1(q_l^*, q_h^*) = q_1^*, \quad B_l(q_1^*) = q_l^*, \quad B_h(q_1^*) = q_h^*$$

If we take $P(Q) = \alpha - Q$, then the solution is

$$q_1^* = \frac{1}{3} (\alpha - 2c + \theta c_l + (1 - \theta)c_h)$$

$$q_l^* = \frac{1}{3} (\alpha - 2c_l + c) - \frac{1 - \theta}{6} (c_h - c_l)$$

$$q_h^* = \frac{1}{3} (\alpha - 2c_h + c) + \frac{\theta}{6} (c_h - c_l)$$

Note $q_l^* > q_h^*$, type with lower cost produces more

Example: Incomplete Information Cournot

In the game with complete information, the Nash equilibrium involves players producing

$$q_i = \frac{1}{3}(\alpha - 2c_i + c_j)$$

With incomplete information, firm 2 produces more than this when its cost is c_h and less when its cost is c_l

This is because firm 1 produces a moderated output

- When firm 2 has cost c_h , firm 1 produces less than it would if it knew c_h , so firm 2 gets to produce a bit more
- When firm 2 has cost c_l , firm 1 produces more than it would if it knew c_l , so firm 2 gets to produce a bit less

Dynamic Games with Incomplete Information

We often need to think about information in dynamic settings in which players learn about the environment over time

- Look at extensive form games, explicit order of moves


As before, we use information sets to represent what players know at each stage

Will also refine away non-credible threats, as in subgame perfect equilibria

- New solution concept: Perfect Bayesian Equilibrium

Example

Selten's Horse:



Dynamic Games with Incomplete Information

Definition

A dynamic game of incomplete information consists of

- A set of players N
- A sequence of histories $\{h^t\}$, each assigned to a player or nature
- An information partition (which histories are in an information set)
- A set of pure strategies S_i for each player i (must include action for each information set)
- A set of types Θ_i for each player i
- A payoff function $u_i(s, \theta)$ for each player i
- A joint probability distribution $p(\theta_1, \dots, \theta_n)$ over types

Strategies and Beliefs

A belief system μ gives a probability distribution over nodes in each information set

- In Selten's horse, player 3 needs beliefs about whether the left or right node was reached

A strategy is *sequentially rational* if, given beliefs, no player can improve her payoff by deviating at any stage of the game

A belief system is *consistent* if it is derived from equilibrium strategies using Bayes' rule

Strategies and Beliefs

For instance, in Selten's horse, if player 1's strategy is D , then we must have $\mu_3(left) = 1$

If player 1 chooses D with probability p and player 2 chooses d with probability q , then Bayes' rule implies

$$\mu_3(left) = \frac{p}{p + (1 - p)q}$$

What if $p = q = 0$? The consistency requirement has no bite, any belief is valid

- If an information set is not reached on the equilibrium path, beliefs are unrestricted

Perfect Bayesian Equilibrium

Definition

In a dynamic game of incomplete information, a perfect Bayesian equilibrium is a strategy profile σ and a belief system μ such that

- The profile σ is sequentially rational given μ
- The belief system μ is consistent given σ

Relatively weak solution concept, often refined by restricting off-path beliefs

Theorem

In any finite dynamic game of incomplete information, a (possibly mixed) perfect Bayesian equilibrium exists.

Social Learning

An important set of questions in network economics:

- How much do people learn from social connections?
- Can interactions aggregate dispersed information?
- How does network structure influence what people learn?

One approach: observational learning

- Often observe choices that other people make
- People base choices on information
- We might infer information from these choices

A simple Bayesian model gives rise to rational herding

Observational Learning



This image is in the public domain.

What can we infer about the deli?

The Classic Herding Model

Two equally likely states of the world $\theta \in \{0, 1\}$

Agents $n = 1, 2, \dots$ sequentially make binary decisions $x_n \in \{0, 1\}$

Earn payoff 1 for matching the state, payoff 0 otherwise

Each agent receives a binary signal $s_n \in \{0, 1\}$

Signals i.i.d. conditional on the state:

$$\mathbb{P}(s_n = 0 \mid \theta = 0) = \mathbb{P}(s_n = 1 \mid \theta = 1) = g > \frac{1}{2}$$

How a Herd Forms

Suppose agents observe the **actions** of all who move earlier

First agent chooses $x_1 = s_1$

Second agent can always justify following signal

If first two match, third agent copies...

- ...and so does everyone else

We end up in a herd

Some Observations

Action produces an informational externality

- Agents do not internalize the value of the information they provide to others

Coarse information is key

- Observing beliefs instead of actions leads to very different result

Herding is not sensitive to the signal structure

- Define public belief $q_n = \mathbb{P}_\sigma(\theta = 1 \mid x_1, x_2, \dots, x_{n-1})$
- The public belief is a martingale and must converge

Rational Herding

Theorem

In any PBE of the social learning game, a herd forms with probability one. The herd is incorrect with strictly positive probability.

We require some basic results on martingales

- A martingale is a stochastic process $X_1, X_2, \dots, X_t, \dots$ such that $\mathbb{E}[X_t | X_1, \dots, X_{t-1}] = X_{t-1}$
- Belief processes are always martingales

Theorem (Martingale Convergence Theorem)

Suppose $\{X_t\}_{t \in \mathbb{N}}$ is a real-valued martingale and there exists a, b such that $a \leq X_t \leq b$ for all t . Then $\lim_{t \rightarrow \infty} X_t$ exists almost surely.

Proof of Theorem

Recall the public belief process $q_n = \mathbb{P}_\sigma(\theta = 1 \mid x_1, x_2, \dots, x_{n-1})$

- This is a martingale

By the martingale convergence theorem, $\lim_{n \rightarrow \infty} q_n$ exists almost surely

If player n chooses 1 (0), then $q_n \geq (\leq) \frac{1}{2}$

- If there are infinitely many switches, only possible limit point is $\frac{1}{2}$

If q_{n-1} is close to $\frac{1}{2}$, then n will follow private signal

- This implies q_n will reflect a new signal
- Inconsistent with convergence to $\frac{1}{2}$

Next Time

Next time we will enrich the model

- Instead of observing entire history, player n observes some *neighborhood* $B(n)$
- Generalize the signal structure

Will also look at a non-Bayesian approach to learning and opinion dynamics

- Based on belief-averaging procedure from DeGroot (1974)

MIT OpenCourseWare
<https://ocw.mit.edu>

14.15J/6.207J Networks
Spring 2018

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.