BELIEFS & EVOLUTIONARY GAME THEORY

Heinrich H. Nax & Bary S. R. Pradelski

hnax@ethz.ch bpradelski@ethz.ch

May 15, 2017: Lecture 10

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Plan

- **Normal form games**
  - Equilibrium invariance
  - Equilibrium refinements
- **Dynamic games**
  - Extensive form games
  - Incomplete information
  - Sub-game perfection
  - Beliefs and Bayes
  - Sequential and perfect Bayesian equilibria
- Evolutionary game theory

Selten's horse

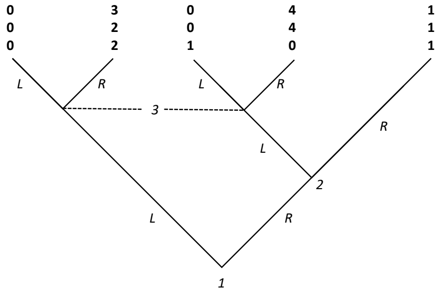

$s = (L, R, R)$ is a subgame perfect equilibrium.

Player 2's move is inconsequential and hence does not need to be examined.

- But why does player 2 play R if his best reply to s is L?
- Moreover if player 1 expects this then he should play R.
- But if player 3 realizes all that he should deviate and play L.

The subgame perfect equilibrium is not "self-enforcing".

Beliefs

Definition: Beliefs (Kreps and Wilson 1982)

A belief system in an extensive form game is a function μ that maps all actions in each information set to a probability distribution, that is:

$$\sum_{a \in D} \mu(a) = 1 \;\; \forall D \in \mathcal{D}$$

A belief system thus encodes players' expectations of each others' play in the whole game (not just along one path).

Weakly consistent and sequentially rational

Definition: weakly consistent

A belief system μ is *weakly consistent* with a behavior strategy profile y if μ agrees with the conditional probability distribution $\mu(\cdot|y)$ induced by y over the information sets on its path.
Suppose we are in information set I, then:

$$\mu(a|y) = \frac{\mathbb{P}[a, y]}{\mathbb{P}[I, y]} \quad \forall a \in I$$

That is, the beliefs are computed via Bayes' rule along the path of play.

We shall assume that all other probabilities are chosen freely.

Perfect Bayes equilibrium

Definition: sequentially rational

A behavior strategy y^* is *sequentially rational* under a belief system μ if for every player $i \in I$ and information set $D \in \mathcal{D}_i$:

$$y_i^* \in \arg\max_{y_i \in Y_i} \sum_{a \in D} \mu(a) U_{ia}(y_i, y_{-i}^*)$$

That is, each player chooses optimally given his beliefs at each information set and the others' equilibrium strategies.

Definition: Perfect Bayesian equilibrium

A behavior strategy profile y^* is a *perfect Bayesian equilibrium* if there exists a weakly consistent belief system under which y^* is sequentially rational for every player.

Back to Selten's horse

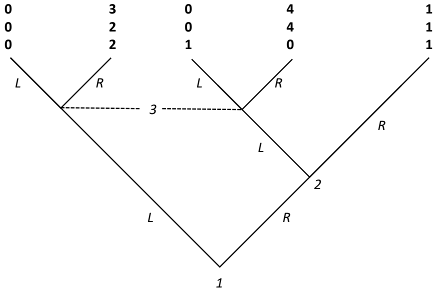

Observation 1: (L, R, R) is not a perfect Bayes equilibrium:

If player 2 gets to move then by playing R he is not sequentially rational.

Back to Selten's horse

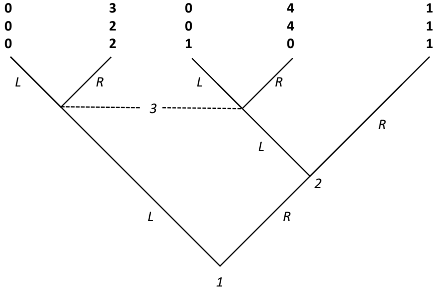

Observation 2: A perfect Bayes Nash equilibrium is given by (R, R, L).

Suppose that player 3 holds belief $\tilde{\mu} = \mu_{\text{player 1 plays L}} \leq 1/3$ then his expected payoff from R is less than that from L:

$$\tilde{\mu} \cdot 2 + (1 - \tilde{\mu}) \cdot 0 \leq \frac{2}{3} \leq \tilde{\mu} \cdot 0 + (1 - \tilde{\mu}) \cdot 1$$

$\Rightarrow$ As player 3 chooses L, both other players are better off playing R.

Quick aside on Bayes' rule

Bayes' rule is central to game theory in its treatment of imperfect information.

Consider the following example:

- The prior probability of a disease is 0.9%.
- If the disease is present then the test is positive with 90% chance.
- The test also gives a false positive with 7% probability.

How likely is it that the patient is ill if the test is positive?

In a study by Gigerenzer (Simon & Schuster 2002), 95% of American medical doctors guessed the answer to be ca. 75%.

Quick aside on Bayes' rule

Bayes' Theorem

Let A and B be events and $\mathbb{P}(B) \neq 0$. Then

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

- $\mathbb{P}(\text{disease}) =: \mathbb{P}(A) = 0.9\%$
- $\mathbb{P}(\text{test positive}) =: \mathbb{P}(B)$
- $\mathbb{P}(B|A) = 90\%$
- $\mathbb{P}(B) = \mathbb{P}(B|A) \cdot \mathbb{P}(A) + \mathbb{P}(B|\neg A)\mathbb{P}(\neg A) = 90\% \cdot 0.9\% + 7\% \cdot 99.1 \approx 7.7\%$
- $\mathbb{P}(A|B) = 10.5\%$

Quick aside on Bayes rule: Intuition

A frequencies approach:

- Imagine 1000 patients, 9 sick, 991 healthy
- Of the 9 sick, $90\% \cdot 9 \approx 8$ test positive
- Of the 991 healthy, $7\% \cdot 991 \approx 69$ test positive.
- Of the total 77 with positive tests, $8/77 \approx 10.5\%$ are ill.

Bayes' Rule is a great tool for everyday life!

Entry deterrence game

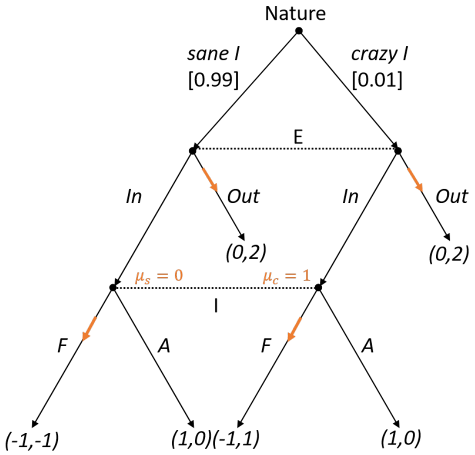

Entry deterrence
***I* unaware if sane/crazy**
The set of PBE includes (Out, F) supported by I's beliefs $\mu_s = 0, \mu_c = 1$.

But E does not observe Nature's move, so how could his action convey information about it and change I's prior beliefs?

Consistency and sequential equilibrium

Definition: consistent

A belief system μ is *consistent* with a behavior strategy profile y^* if there exists a sequence of interior behavior strategy profiles $y^t \to y^*$ such that $\mu(a|y^t) \to \mu^*(a)$ for all a where $\mu(\cdot|y^t)$ is the belief system induced by Bayes' law from y^t.

Definition: Sequential equilibrium

A behavior strategy profile y^* is a *sequential equilibrium* if there exists a consistent belief system under which y^* is sequentially rational for every player.

Existence and inclusions

Theorem (Fudenberg & Tirole 1991)

Every finite game has a perfect Bayesian equilibrium.

Theorem (Kreps & Wilson 1982)

Every finite game has a sequential equilibrium.

Sequential equilibrium $\subseteq$ perfect Bayesian equilibria $\subseteq$ subgame perfect equilibria

Proposition

Every subgame perfect equilibrium is a Nash equilibrium of the associated strategic form game.

Key learnings

- Dynamic games take timing and information into consideration
- Equilibrium refinements are needed to make “reasonable” predictions
- When converting to reduced normal form, remember that strategies are complete contingent plans
- Perfect Bayesian equilibrium and sequential equilibrium are not only strategy profiles but also (Bayesian or consistent) beliefs at every information set

Common knowledge of rationality and the game

Suppose that players are rational decision makers and that mutual rationality is common knowledge, that is:

- I know that she knows that I will play rational
- She knows that “I know that she knows that I will play rational”
- I know that “She knows that “I know that she knows that I will play rational””
- ...

Further suppose that all players know the game and that again is common knowledge.

Rationality and the “as if” approach

- The rationalistic paradigm in economics (Savage, *The Foundations of Statistics*, 1954)
  - A person’s behavior is based on maximizing some goal function (utility) under given constraints and information
- The “as if” approach (Friedman, *The methodology of positive economics*, 1953)
  - Do not theorize about the intentions of agents’ actions but consider only the outcome (observables)
  - Similar to the natural sciences where a model is seen as an approximation of reality rather than a causal explanation (e.g., Newton’s laws)

But is the claim right? Do people act (as if) they where rational?

Nash's mass-action interpretation (Nash, *PhD thesis*,1950)

*"We shall now take up the **"mass-action" interpretation of equilibrium** points. In this interpretation solutions have no great significance. It is **unnecessary to assume that the participants have full knowledge** of the total structure of the game, or the ability and inclination to go through **any complex reasoning** processes. But the participants are supposed to **accumulate empirical information on the relative advantages** of the various pure strategies at their disposal.*

...

*Thus the assumption we made in this "mass-action" interpretation lead to the conclusion that the **mixed strategies representing the average behavior** in each of the populations form an **equilibrium**."*

(bold text added for this presentation)

Nash's mass-action interpretation (Nash, *PhD thesis*,1950)

- A large population of identical individuals represents each player role in a game
- The game is played recurrently ($t = 0, 1, 2, 3, ...$):
  - In each period one individual from each player population is drawn randomly to play the game
- Individuals observe samples of earlier behaviors in their own population and avoid suboptimal play

Nash's claim: If all individuals avoid suboptimal pure strategies and the population distribution is stationary then it constitutes a [Nash] equilibrium

Almost true! Evolutionary game theory formalizes these questions and provides answers.

The folk theorem of evolutionary game theory

Folk theorem

- If the population process converges from an interior initial state, then for large t (in the limit) the distribution is a Nash equilibrium
- If a stationary population distribution is stable, the it coincides with a Nash equilibrium

Charles Darwin: “Survival of the fittest”
The population which is best adapted to environment (exogenous) will reproduce more

Evolutionary game theory
The population which performs best against other populations (endogenous) will survive/reproduce more

Domain of analysis

Symmetric two-player games

A symmetric two-player normal form game $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ consists of three object:

1. *Players:* $N = \{1, 2\}$, with typical player $i \in N$.
2. *Strategies:* $S_1 = S_2 = S$ with typical strategy $s \in S$.
3. *Payoffs:* A function $u_i : (h, k) \to \mathbb{R}$ mapping strategy profiles to a payoff for each player i such that for all $h, k \in S$:

$$u_2(h, k) = u_1(k, h)$$

Battle of the Sexes

	Cafe	*Pub*
Cafe	4, 3	0, 0
Pub	0, 0	3, 4

Not symmetric since:

$$u_1(Cafe, Cafe) \neq u_2(Cafe, Cafe)$$

Prisoner's dilemma

	Cooperate	*Defect*
Cooperate	$-1, -1$	$-8, \ 0$
Defect	$0, -8$	$-5, -5$

Symmetric since:

$$\begin{aligned}
u_1(Cooperate, Cooperate) &= u_2(Cooperate, Cooperate) &= -1 \\
u_1(Cooperate, Defect) &= u_2(Defect, Cooperate) &= -8 \\
u_1(Defect, Cooperate) &= u_2(Cooperate, Defect) &= 0 \\
u_1(Defect, Defect) &= u_2(Defect, Defect) &= -5
\end{aligned}$$

Symmetric Nash equilibrium

Definition: Symmetric Nash Equilibrium

A *symmetric Nash equilibrium* is a strategy profiles σ^* such that for every player i,

$$u_i(\sigma^*, \sigma^*) \geq u_i(\sigma, \sigma^*) \text{ for all } \sigma$$

In words: If no player has an incentive to deviate from their part in a particular strategy profile, then it is Nash equilibrium.

Proposition

In a symmetric normal form game there always exists a symmetric Nash equilibrium.

Note: Not all Nash equilibria of a symmetric game need to be symmetric.

Evolutionarily stable strategy (Maynard Smith and Price, 1972)

Definition: Evolutionarily stable strategy (ESS)

A mixed strategy $\sigma \in \Delta(S)$ is an *evolutionarily stable strategy (ESS)* if for every strategy $\tau \neq \sigma$ there exists $\varepsilon(\tau) \in (0, 1)$ such that for all $\varepsilon \in (0, \varepsilon(\tau))$:

$$U(\sigma, \varepsilon\tau + (1 - \varepsilon)\sigma) > U(\tau, \varepsilon\tau + (1 - \varepsilon)\sigma)$$

Let Δ^{ESS} be the set of evolutionarily stable strategies.

Alternative representation

Note that an ESS needs to be a best reply to itself, thus Δ^{ESS} is a subset of the set of Nash equilibria.

Proposition

A mixed strategy $\sigma \in \Delta(S)$ is an *evolutionarily stable strategy (ESS)* if:

$$\begin{aligned} U(\tau,\sigma) &\leq U(\sigma,\sigma) \ \ \forall\tau \\ U(\tau,\sigma) = U(\sigma,\sigma) &\Rightarrow U(\tau,\tau) < U(\sigma,\tau) \ \ \forall\tau \neq \sigma \end{aligned}$$

Prisoner's dilemma

	Cooperate	*Defect*
Cooperate	$-1, -1$	$-8, \quad 0$
Defect	$0, -8$	$-5, -5$

$$\Delta^{ESS} = \{\text{Defect}\}$$

Coordination game

	A	B
A	$4,4$	$0,0$
B	$0,0$	$1,1$

- Nash equilibria:
  $(A,A),(B,B),(0.2\cdot A+0.8\cdot B, 0.2\cdot A+0.8\cdot B)$
- All Nash equilibra are symmetric.
- But the mixed Nash equilibrium is not ESS:
  - L performs better against it!
- Note that the mixed Nash equilibrium is trembling-hand perfect.

Existence of ESS not guaranteed

Example: Rock, paper, scissors

	R	P	S
R	$0,\ \ 0$	$-1,\ \ 1$	$1, -1$
P	$1, -1$	$0,\ \ 0$	$-1,\ \ 1$
S	$-1,\ \ 1$	$1, -1$	$0,\ \ 0$

- Unique Nash equilibrium and thus symmetric: $\sigma = (\frac{1}{3}R, \frac{1}{3}P, \frac{1}{3}S)$
- All pure strategies are best replies and do as well against themselves as σ does against them $\Rightarrow$ Not an ESS!

Relations to normal form refinements

Propositions

- If $\sigma \in \Delta(S)$ is weakly dominated, then it is not evolutionarily stable.
- If $\sigma \in \Delta^{ESS}$, then (σ, σ) is a perfect equilibrium.
- If (σ, σ) is a strict Nash equilibrium, then σ is evolutionarily stable.

Summary

- Evolutionary game theory studies mutation (ESS) and selection process (not in this lecture)
- The stable states often coincide with solution concepts from the "rational" framework
- Evolutionary game theory does not explain **how** a population arrives at such a strategy
  $\Rightarrow$ Learning in games and behavioral game theory

The "best" textbook: Weibull, *Evolutionary game theory*, 1995

THANKS EVERYBODY

See you next week!

And keep checking the website for new materials as we progress:

http://www.coss.ethz.ch/education/GT.html