Noncooperative Solutions

MIT 14.126-Game Theory

Sergei Izmalkov
Muhamet Yildiz

(Thanks to Paul Milgrom)

Noncooperative Solutions

Definitions

Normal form games

- A normal form game is a triple (**N**,S,u)
  - **N**={1,…,N} is the (non-empty) set of players
  - $S=S_1\times\ldots\times S_N$ is the set of strategy profiles
    - A mixed strategy for n is a probability distribution π_n on S_n, that is, $\pi_n\in\Delta(S_n)$.
  - $u:S\rightarrow R^N$ is the payoff vector.
    - Define $u(\pi)=\sum\ldots\sum u(s_1,\ldots,s_N)\pi_1(s_1)\ldots\pi_N(s_N)$.
- The payoff of player n is $u_n(\pi)$.
- Strategies may be correlated: $u(\pi)= \sum_s u(s)\pi(s)$

Mixed Strategies?

- Interpretations of mixed equilibria
  - Bluffing?
    - An informal consensus now sees "bluffing" as pooling equilibria in incomplete information games.
  - Uncertainty in other minds
  - Population proportions
  - Harsanyi: Behavioral approximation to pure strategy equilibrium in games with random payoff perturbations

Reasons for Correlation

- People's play can be correlated objectively for reasons outside the game:
  - They can observe a common variable, such as the weather.
  - They can have a common culture which inclines them to common inclinations, unknown to the outside observer, which appear as correlations in behavior.
- Play can also be subjectively correlated.
  - The observer may be learning some unknown aspect of human behavior.
  - The observer may know that each tribe has one chief without knowing who it is. Then, "chief" behavior by one player makes others less likely to exhibit "chief" behavior.

Noncooperative Solutions

Dominance & Rationalizability

Bayesian Rationality

- A rational (Bayesian) player
  - has beliefs about the likely play of others
  - optimizes according to those beliefs
- Theorem. A strategy for player n is a *best reply* to *some* probability distribution of the play of others if and only if it is *not strictly dominated* by any pure or mixed strategy of player n.

7

Proof

- If s_n is strictly dominated by s_n', then s_n' is a better reply to any probability distribution over the strategies of others.
- Suppose s_n is not strictly dominated. Let K be the number of profiles in S_{-n}. Then the following two convex subsets of R^K are disjoint
  - $A=\text{Convex Hull}\{(u_n(s_n',s_{-n});s_{-n}\in S_{-n})|s_n'\in S_n\}$
  - $B=\{z\in R^K|z>u_n(s_n,\cdot)\}$
- By the separating hyperplane theorem, there exists a non-zero vector p such that $p\cdot y\geq p\cdot z$ for all $z\in A$, $y\in B$. By inspection, p is non-negative and can be normalized to the required probability vector.
- Since $y^*=u_n(s_n,\cdot)$ is on the boundary of B, $p.y^*\geq p\cdot z$ for all $z\in A$. **QED**

8

Dominant strategy equilibrium

- Assume for each n there is a strategy s_n^* that strictly dominates all other strategies of n.

Theorem: If all players are rational, then s^* will be played.

- Assume for each n there is a strategy s_n^{**} that weakly dominates all other strategies of n.

Theorem: If all players are rational and cautious, then s^{**} will be played.

Rationalizability Defined

- Definition. A strategy s_j for player j is *rationalizable* if there is a collection of sets $\{Z_n\}$ such that $s_j \in Z_j$ and for all players n,
  - $Z_n \subseteq S_n$
  - For all $s_n \in Z_n$, s_n is a best reply to some belief p_n whose support is a subset of Z_{-n}.
- Note that
  - all the elements of each such Z_n are rationalizable.
  - The union of all such Z_n is the set of rationalizable strategies for player n.

Iterating to Rationalizability

- First step
  - Every player is (Bayesian) "rational." Therefore, every player plays only strategies that are not strictly dominated.
  - Create a new strategic form by eliminating dominated strategies from the original game.
- Iteration n+1
  - Every player knows that others will play only strategies remaining from iteration n. Rational players choose a best reply, that is, a strategy that is not strictly dominated in the new game.
  - Create strategic form n+1 by eliminating dominated strategies from strategic form n.
- If the original game is finite, eventually no changes are made.

11

Iterated strict dominance & Rationalizability

- Theorem. When the iterative procedure terminates, the remaining strategies for each player n are precisely n's rationalizable strategies.
- Proof sketch.
  - Each stage eliminates only strategies that are not rationalizable.
  - Remaining strategies are necessarily rationalizable.

12

CK of Rationality & Rationalizability

Theorem (Bernheim, Pearce): Assume that the payoffs and the rationality are common knowledge. Then, each player must play a rationalizable strategy. Moreover, given any rationalizable strategy profile s, there exists a hierarchy of beliefs at which s is played.

Nash Eq. & Rationalizability

- Every Nash equilibrium is rationalizable.
- Theorem (Milgrom and Roberts): Assume a supermodular game on a complete lattice. Then, there exist the smallest and the largest Nash equilibria, x and y, respectively. If z is rationalizable, then $x \leqq z \leqq y$.
- Corollary: If a supermodular game has a unique Nash equilibrium, then it has a unique rationalizable strategy.

Rationalizability in Linear Cournot duopoly

- $P = 1 - q_1 - q_2;$
- $c_1 = c_2 = c$

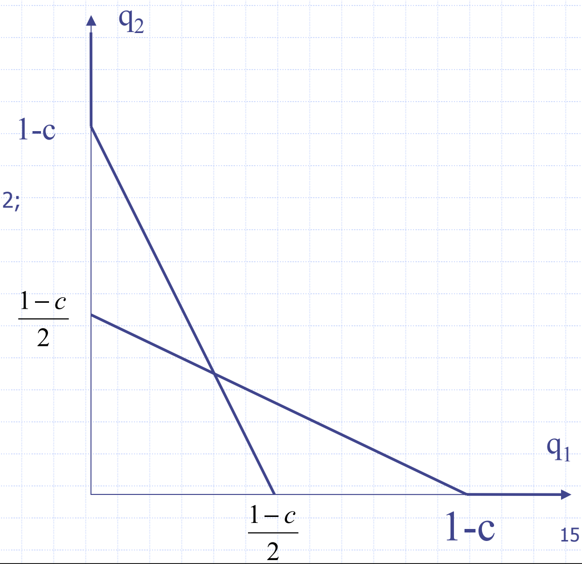

Example (Robustness)

	L	R
T	2,1.001	1001,1
B	1,1	1000,1000

Noncooperative Solutions

Nash Equilibrium & Its Existence Theorems

Nash Equilibrium

- A normal form game is a triple (**N**,S,u)
  - **N**={1,…,N} is the (non-empty) set of players
  - $S=S_1\times\ldots\times S_N$ is the set of strategy profiles
    - A mixed strategy for n is a probability distribution π_n on S_n, that is, $\pi_n\in\Delta(S_n)$.
  - $u:S\rightarrow R^N$ is the payoff vector.
    - Define $u(\pi)=\sum\ldots\sum u(s_1,\ldots,s_N)\pi_1(s_1)\ldots\pi_N(s_N)$.
- Equilibrium: A mixed strategy profile $\pi\in\Pi$ is a Nash equilibrium if

$$(\forall n\in \mathbf{N})(\forall \pi'_n\in\Delta(S_n))u_n(\pi)\geq u_n(\pi'_n,\pi_{-n})$$

Epistemic conditions for NE

Theorem (Aumann & Brandenburger): In a 2-person game, assume that payoff functions, the rationality of players, and their conjectures are all mutually known. Then, the conjectures constitute a Nash equilibrium.

For n>2 players, we need common prior assumption and common knowledge of conjectures.

Nash's Existence Equilibrium

- Theorem (Nash). Suppose the number of players and the strategy sets are finite. Then, there exists a Nash equilibrium.
- Proof. Let Δ be the set of mixed strategy profiles and consider the map $f:\Delta\rightarrow\Delta$ given as follows. $f_n(\pi)$ is the probability distribution on S_n that assigns to any strategy s_n the probability:

$$f_n(\pi)(s_n)=\frac{\max(0,\pi_n(s_n)+u_n(s_n,\pi_{-n})-u_n(\pi))}{\sum_{s'_n\in S_n}\max(0,\pi_n(s'_n)+u_n(s'_n,\pi_{-n})-u_n(\pi))}$$

Proof, continued

- Since this function is continuous and Δ is convex and compact, f has a fixed point.
- By inspection, the fixed point has the properties that
  - Every one of n's strategies that is played with zero probability has an expected profit no higher than $u_n(\pi)$.
  - Every one of n's pure strategies that is played with positive probability has the same expected profit.
- Therefore, the fixed point is a Nash equilibrium. **QED**

Concave Payoff Functions

- Theorem. Let $S_1,\ldots,S_N$ be convex, compact subsets of a Euclidean space. Suppose that for all n, $u_n:S\rightarrow R$ is continuous and that for all s_{-n}, $U(s_n)=u_n(s_n,s_{-n})$ is concave. Then, there exists a Nash equilibrium strategy profile $s\in S$.
- Notes:
  - On its face, this is a "pure strategy" Nash equilibrium existence theorem.
  - Given any finite game, the corresponding game in "mixed strategies" is linear in the strategies and hence satisfies the stated hypotheses.

Proof

- Consider the "best reply" correspondence f:S→S, given by:

$$f_n(s) = \arg\max u_n(\ , s_{-n})$$
$$= \{s''_n \in S_n \mid (\forall s'_n \in S_n) u_n(s''_n, s_{-n}) \geq u_n(s'_n, s_{-n})\}$$

- Observe (next slide) that f has a "closed graph" and is convex valued.
- The Kakutani fixed point theorem applies (as do several others). By construction, a fixed point is a Nash equilibrium. **QED**

Details

- Kakutani's Fixed Point Theorem.
  - Let Δ be a convex, compact subset of R^N and let f:Δ→Δ be a multifunction ("correspondence") such that for all x∈Δ, f(x) is convex and such that the graph of f is closed. Then, there exists x∈Δ such that x∈f(x).
- Proving closed graph
  - Let $\{s^k\}$ be a sequence of strategy profiles converging to s^*; let $r^k \in f(s^k)$; and let r^* be an accumulation point of $\{r^k\}$. We limit attention to a convergent subsequence of $\{r^k\}$. By continuity of u,

$$(\forall s'_n \in S_n) u_n(r^*_n, s^*_{-n}) = \lim_{k\to\infty} u_n(r^k_n, s^k_{-n})$$
$$\geq \lim_{k\to\infty} u_n(s'_n, s^k_{-n}) = u_n(s'_n, s^*_{-n})$$

Other Existence Proofs

- The various "general" existence proofs of Nash equilibria are based on fixed point theorems.
  - Some are topological theorems
  - Later in the term, we will encounter a lattice-based fixed point theorem.

Noncooperative Solutions

Correlated Equilibrium

Motivation

- People's play can be correlated objectively for reasons outside the game:
  - They can observe a common variable, such as the weather.
  - They can have a common culture which inclines them to common inclinations, unknown to the outside observer, which appear as correlations in behavior.
- Play can also be subjectively correlated.
  - The observer may be learning some unknown aspect of human behavior.
  - The observer may know that each tribe has one chief without knowing who it is. Then, "chief" behavior by one player makes others less likely to exhibit "chief" behavior.

27

Definition

- Given a finite strategic form game (N,S,u), a *correlated strategy profile* consists of these elements:
  - A finite probability space (Ω,π)
  - For each player n, an information partition P_n of Ω.
  - For each player n, a strategy $\sigma_n:\Omega \rightarrow S_n$ measurable with respect to P_n.
- The correlated strategy profile is a *correlated equilibrium* if for each n and each strategy τ_n measurable with respect to P_n,

$$\sum_{\omega\in\Omega} u_n(\sigma(\omega))\pi(\omega) \geq \sum_{\omega\in\Omega} u_n(\tau_n(\omega),\sigma_{-n}(\omega))\pi(\omega)$$

28

Characterization Result

- Theorem. Every probability distribution π of strategy profiles in a correlated equilibrium can be achieved by setting the probability space to (S,π), the partition P_n be sets of the form $\{s\in S|s_n=a\}$, and the strategy profile σ so that $\sigma_n(s)=s_n$.
  - We henceforth abbreviate by saying "π is a correlated equilibrium" to mean that $((S,\pi),P,\sigma)$ as defined in the theorem is a correlated equilibrium.

Example

	L	R
T	2,1	0,0
B	0,0	1,2

	L	R
T	p_1	p_2
B	p_3	p_4

	L	R
T	$\frac{p_1}{p_1+p_2}$	$\frac{p_2}{p_1+p_2}$
B	$\frac{p_3}{p_3+p_4}$	$\frac{p_4}{p_3+p_4}$

	L	R
T	$\frac{p_1}{p_1+p_3}$	$\frac{p_2}{p_2+p_4}$
B	$\frac{p_3}{p_1+p_3}$	$\frac{p_4}{p_2+p_4}$

- Conditions:

$$\frac{2p_1}{p_1+p_2}+\frac{0p_2}{p_1+p_2}\geq\frac{0p_1}{p_1+p_2}+\frac{1p_2}{p_1+p_2}$$

  - $2p_1 \geq p_2$
  - $p_4 \geq 2p_3$
  - $2p_4 \geq p_2$
  - $p_1 \geq 2p_3$

Additional Results

- Theorem. If π is a probability distribution of strategy profiles at a Nash equilibrium, then π is a correlated equilibrium.
- Theorem. The set of correlated equilibria π is a closed, convex set.
- Proof. π is a correlated equilibrium if and only if it satisfies the following set of linear inequalities, for all players n functions $f_n:S_n \to S_n$.

$$\sum_{s \in S} \pi(s) u(s) \geq \sum_{s \in S} \pi(s) u\left(f_n(s_n), s_{-n}\right)$$

- Theorem. Any strategy played with positive probability at a correlated equilibrium is rationalizable.

31

Noncooperative Solutions

Extensive Forms and Sequential Rationality

32

Extensive Forms

- An extensive form consists of:
  - The set of players, **N**.
  - The set of histories (sequences) H, which includes the empty sequence and has the property that if $(a_1,\ldots,a_K) \in H$ and L<K, then $(a_1,\ldots,a_L) \in H$.
    - A history $(a_1,\ldots,a_L)$ is *complete* (or *terminal*) if there is no K>L such that $(a_1,\ldots,a_K) \in H$. All infinite histories $(a_1,\ldots,a_l,\ldots)$ are complete.
    - For any non-terminal history write $A(h) = \{a | (h,a) \in H\}$
  - A payoff function mapping complete histories into payoffs for each player.
  - A function P that assigns to each non-terminal history $h \in H$ a player n or chance c (the player who "moves" at h); if P(h)=c, assign also a probability distribution on A(h).
  - A partition (I^k) of non-terminal h with $P(h) \in N$ such that, if h,h' $\in I^k$, then A(h) = A(h') and P(h) = P(h').

33

Definitions

- A game game is said to *have perfect recall* iff no player forgets what he knew and what he has done. We will always assume perfect recall.
- A *perfect information game* is a game in which all information sets are singleton.

34

Sequential Rationality

- Given any information set I_n, where player n moves, player n has
  - a probability distribution $\mu(.|I_n)$ on I_n
  - and a probability distribution on the others' play in the "continuation game" (which may not be a "subgame"),

  representing his beliefs conditional on the even that I_n is reached.
- A player n is *sequentially rational* iff, at any information set I_n he moves, he maximizes his expected payoff conditional on that I_n is reached and according to his beliefs at I_n.

35

Conditional Dominance

- A strategy is conditionally dominated at I_n iff its restriction to the "continuation game" at I_n is strictly dominated at that game for every probability distribution on I_n.
- A sequentially rational player never plays a conditionally dominated strategy.

36

Iterated conditional dominance

- First step
  - Every player is sequentially rational. Therefore, every player plays only strategies that are not conditionally dominated.
  - Create a new game by eliminating conditionally dominated strategies from the original game.
- Iteration n+1
  - Every player knows that others will play only strategies remaining from iteration n. Sequentially rational players choose a best reply at any information set, that is, a strategy that is not conditionally dominated in the new game.
  - Create strategic form n+1 by eliminating conditionally dominated strategies from strategic form n.
- If the original game is finite, eventually no changes are made.

"Theorem"

If the game and the players' sequential rationality are common knowledge, then they will play a strategy profile that survives iterated conditional dominance.

Centipede game

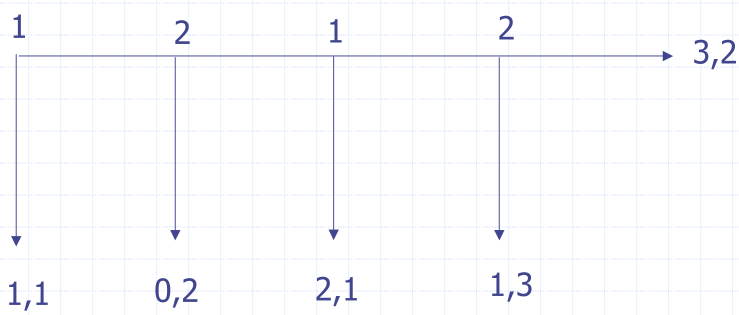

Backwards Induction

- Iterated conditional dominance in a perfect-information game with finite histories is called backwards induction.

Sequential Equilibrium

- An *assessment* is a pair $(s,\mu(.|.))$ of a strategy profile s and a function $\mu(.|.)$ that gives a conditional probability distribution $\mu(.|I_n)$ at every information set I_n.
- An assessment $(s,\mu(.|.))$ is *sequentially rational* iff each s_n is a best response to s_{-n} at each information set I_n of n according to $\mu(.|I_n)$.
- An assessment $(s,\mu(.|.))$ is *consistent* iff there is a sequence $((s^k,\mu^k(.|.)))_k$ of assessments s.t.
  - $(s^k,\mu^k(.|.)) \rightarrow (s,\mu(.|.))$ in Euclidean metric;
  - each s^k is completely mixed, and
  - $\mu^k(.|.)$ is derived from s^k using Bayes' rule.

Sequential Equilibrium

Definition (Kreps, Wilson): A sequential equilibrium is an assessment that is both sequentially rational and consistent.

Example

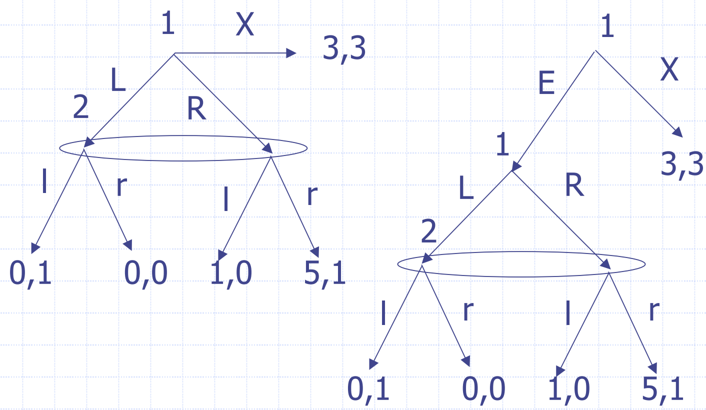

43