6.254 : Game Theory with Engineering Applications
Lecture 4: Strategic Form Games - Solution Concepts

Asu Ozdaglar
MIT

February 11, 2010

Outline

- Review
- Correlated Equilibrium
- Existence of a Mixed Strategy Equilibrium in Finite Games

- Reading:
  - Fudenberg and Tirole, Chapters 1 and 2.

Rationalizability

- A different solution concept in which a player's belief about the other players' actions is not assumed to be correct (as in a Nash equilibrium), but rather, simply constrained by rationality.

  (1) Players maximize with respect to some (uncorrelated) beliefs about opponent's behavior (i.e., they are rational).
  (2) Beliefs have to be consistent with other players being rational, and being aware of each other's rationality, and so on (but they need not be correct).

- Leads to an infinite regress: "I am playing strategy σ_1 because I think player 2 is using σ_2, which is a reasonable belief because I would play it if I were player 2 and I thought player 1 was using σ_1', which is a reasonable thing to expect for player 2 because σ_1' is a best response to σ_2', ...

Never-Best Response and Strictly Dominated Strategies

Definition

A pure strategy s_i *is* **strictly dominated** *if there exists a mixed strategy* $\sigma_i \in \Sigma_i$ *such that*

$$u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \qquad \text{for all } s_{-i} \in S_{-i}.$$

Definition

A pure strategy s_i *is a* **never-best response** *if for all beliefs* σ_{-i} *there exists* $\sigma_i \in \Sigma_i$ *such that*

$$u_i(\sigma_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i}).$$

- A strictly dominated strategy is a never-best response.
- Does the converse hold?
- Last time, we studied a 3-player example that illustrates a never-best response strategy which is not strictly dominated.

Rationalizable Strategies

Iteratively eliminating never-best response strategies yields rationalizable strategies.

- Start with $\tilde{S}_i^0 = S_i$.
- For each player $i \in \mathcal{I}$ and for each $n \geq 1$,

$$\tilde{S}_i^n = \{s_i \in \tilde{S}_i^{n-1} \mid \exists \quad \sigma_{-i} \in \prod_{j \neq i} \tilde{\Sigma}_j^{n-1} \text{ such that } u_i(s_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i}) \quad \text{for all } s_i' \in \tilde{S}_i^{n-1}\}.$$

- Independently mix over $\tilde{S}_i^n$ to get $\tilde{\Sigma}_i^n$.
- Let $R_i^\infty = \cap_{n=1}^\infty \tilde{S}_i^n$. We refer to the set R_i^∞ as the set of rationalizable strategies of player i.

Rationalizable Strategies

- Since the set of strictly dominated strategies is a strict subset of the set of never-best response strategies, set of rationalizable strategies represents a further refinement of the strategies that survive iterated strict dominance.
- Let NE_i denote the set of pure strategies of player i used with positive probability in any mixed Nash equilibrium.
- Then, we have

$$NE_i \subseteq R_i^\infty \subseteq D_i^\infty,$$

where R_i^∞ is the set of rationalizable strategies of player i, and D_i^∞ is the set of strategies of player i that survive iterated strict dominance.

Remarks:

- For a two-player game, a never-best response strategy is always strictly dominated.
- If beliefs are allowed to be correlated, then a never-best response strategy is always strictly dominated. (Proof relies on the separating hyperplane theorem; check your book or the notes on Stellar)

Correlated Strategies

- In a Nash equilibrium, players choose strategies (or randomize over strategies) independently.
- For games with multiple Nash equilibria, one may want to allow for randomizations between Nash equilibria by some form of communication prior to the play of the game.

Example Consider the Battle of the Sexes game:

	Ballet	Football
Ballet	$1, 4$	$0, 0$
Football	$0, 0$	$4, 1$

Suppose that the players flip a coin and go to the Ballet if the coin is Heads, and to the Football game if the coin is tails, i.e., they randomize between two pure strategy Nash equilibria, resulting in a payoff of $(5/2, 5/2)$ that is not a Nash equilibrium payoff.

Traffic Intersection Game

Consider a game where two cars arrive at an intersection simultaneously. Row player (player 1) has the option to play U or D, and the column player (player 2) has the option to play L or R with payoffs as follows.

	L	R
U	5, 1	0, 0
D	4, 4	1, 5

- There are two pure strategy Nash equilibria: (U, L) and (D, R).
- To find the mixed strategy Nash equilibria, assume player 1 plays U with probability p and player 2 plays L with probability q. Using the mixed equilibrium characterization, we have

$$5q = 4q + (1 - q) \Rightarrow q = \frac{1}{2}$$
$$5p = 4p + (1 - p) \Rightarrow p = \frac{1}{2}$$

- This implies that there is a unique mixed strategy equilibrium with expected payoff (5/2,5/2).

Traffic Intersection Game

- Assume that there is a publicly observable random variable (such as a fair coin) such that with probability 1/2 (Head), player 1 plays U and player 2 plays L, and with probability 1/2 (Tail), player 1 plays D and player 2 plays R.
- The expected payoff for this play of the game increases to (3,3).
- We show that no player has an incentive to deviate from the "recommendation" of the coin.
- If player 1 sees a Head, he believes that player 2 will play L, and therefore playing U is his best response (similar argument when he sees a Tail).
- Similarly, if player 2 sees a Head, he believes that player 1 will play U, and therefore playing L is his best response (similar argument when he sees a Tail).
- When the recommendation of the coin is part of a Nash equilibrium, no player has an incentive to deviate

Traffic Intersection Game

- With a publicly observable random variable, we can get any payoff vector in the convex hull of Nash equilibrium payoffs.
  - Note that the convex hull of a finite number of vectors $x_1, \ldots, x_k$ is given by

$$\text{conv}(\{x_1, \ldots, x_k\}) = \{x \mid x = \sum_{i=1}^{k} \lambda_i x_i,\ \lambda_i \geq 0,\ \sum_{i=1}^{k} \lambda_i = 1\}$$

- The coin flip is one way of communication prior to the play.
- A more general form of communication is to find a trusted mediator who can perform general randomizations.
- Consider next a more elaborate signalling scheme.
- Suppose the players find a mediator who chooses $x \in \{1, 2, 3\}$ with equal probability $1/3$. She then sends the following messages:
  - If $x = 1$, player 1 plays U, player 2 plays L.
  - If $x = 2$, player 1 plays D, player 2 plays L.
  - If $x = 3$, player 1 plays D, player 2 plays R.

Traffic Intersection Game

- We show that no player has an incentive to deviate from the "recommendation" of the mediator:
    - If player 1 gets the recommendation U, he believes player 2 will play L, so his best response is to play U.
    - If player 1 gets the recommendation D, he believes player 2 will play L, R with equal probability, so playing D is a best response.
    - If player 2 gets the recommendation L, he believes player 1 will play U, D with equal probability, so playing L is a best response.
    - If player 2 gets the recommendation R, he believes player 1 will play D, so his best response is to play R.
- Thus the players will follow the mediator's recommendations.
- With the mediator, the expected payoffs are $(10/3, 10/3)$, strictly higher than what the players could get by randomizing between Nash equilibria.

Correlated Equilibrium

- The preceding examples lead us to the notions of correlated strategies and "correlated equilibrium".
- Let $\Delta(S)$ denote the set of probability measures over the set S.
- Let R be a random variable taking values in $S = \Pi_{i=1}^{n} S_i$ distributed according to π.
  - An instantiation of R is a pure strategy profile and the i^{th} component of the instantiation will be called the recommendation to player i.
  - Given such a recommendation, player i can use conditional probability to form a posteriori beliefs about the recommendations given to the other players.

Correlated Equilibrium

Definition

A **correlated equilibrium** *of a finite game is a joint probability distribution* $\pi \in \Delta(S)$ *such that if* R *is a random variable distributed according to* π *then*

$$\sum_{s_{-i} \in S_{-i}} \mathrm{Prob}(R = s | R_i = s_i) \left[u_i(s_i, s_{-i}) - u_i(t_i, s_{-i}) \right] \geq 0$$

for all players i*, all* $s_i \in S_i$ *such that* $\mathrm{Prob}(R_i = s_i) > 0$*, and all* $t_i \in S_i$*.*

- A distribution π is defined to be a correlated equilibrium if no player can ever expect to unilaterally gain by deviating from his recommendation, assuming the other players play according to their recommendations.

Characterization of Correlated Equilibrium

We have the following useful characterization for correlated equilibria in finite games.

Proposition

A joint distribution $\pi \in \Delta(S)$ is a correlated equilibrium of a finite game if and only if

$$\sum_{s_{-i} \in S_{-i}} \pi(s) \left[u_i(s_i, s_{-i}) - u_i(t_i, s_{-i}) \right] \geq 0 \quad (1)$$

for all players i and all $s_i, t_i \in S_i$ such that $s_i \neq t_i$.

Characterization of Correlated Equilibrium

Proof.

- Using the definition of conditional probability, we can rewrite the definition of a correlated equilibrium as

$$\sum_{s_{-i}\in S_{-i}} \frac{\pi(s)}{\sum_{t_{-i}\in S_{-i}} \pi(s_i, t_{-i})} \left[u_i(s_i, s_{-i}) - u_i(t_i, s_{-i})\right] \geq 0$$

  for all i, all $s_i \in S_i$ such that $\sum_{t_{-i}\in S_{-i}} \pi(s_i, t_{-i}) > 0$, and all $t_i \in S_i$.
- The denominator does not depend on the variable of summation so it can be factored out of the sum and cancelled, yielding the simpler condition that (1) holds for all i, all $s_i \in S_i$ such that $\sum_{t_{-i}\in S_{-i}} \pi(s_i, t_{-i}) > 0$, and all $t_i \in C_i$.
- But if $\sum_{t_{-i}\in S_{-i}} \pi(s_i, t_{-i}) = 0$ then the left hand side of (1) is zero regardless of i and t_i, so the equation always holds trivially in this case.

Characterization of Correlated Equilibrium

- Another equivalent convenient characterization: a joint distribution $\pi \in \Delta(S)$ is a correlated equilibrium of a finite game if and only if for all i and s_i with $\pi(s_i) > 0$ (i.e., marginal distribution),

$$\sum_{s_{-i} \in S_{-i}} \pi(s_{-i} \mid s_i) \left[u_i(s_i, s_{-i}) - u_i(t_i, s_{-i}) \right] \geq 0,$$

for all $t_i \in S_i$.

Remarks:

- Any mixed Nash equilibrium is a correlated equilibrium.
- The set of correlated equilibria is a convex set.
- An immediate implication of the preceding two statements is that the set of correlated equilibria contains the convex hull of the set of (mixed) Nash equilibria.

Departure Function Characterization

- We can alternatively think of correlated equilibria as joint distributions corresponding to recommendations which will be given to the players as part of an extended game.
- The players are then free to play any function of their recommendation (this is called a departure function) as their strategy in the game.
- If it is a Nash equilibrium of this extended game for each player to play his recommended strategy (i.e., to use the identity departure function), then the distribution is a correlated equilibrium.

Proposition

A joint distribution $\pi \in \Delta(S)$ is a correlated equilibrium of a finite game if and only if

$$\sum_{s \in S} \pi(s) \left[u_i(s_i, s_{-i}) - u_i(\zeta_i(s_i), s_{-i}) \right] \geq 0 \quad (2)$$

for all players i and all functions $\zeta_i : S_i \to S_i$.

Departure Function Characterization

Proof.

By substituting $t_i = \zeta_i(s_i)$ into (1) and summing over all $s_i \in S_i$ we obtain (2) for any i and any $\zeta_i : S_i \to S_i$. For the converse, define ζ_i for any $s_i, t_i \in S_i$ by

$$\zeta_i(r_i) = \begin{cases} t_i & r_i = s_i \\ r_i & \text{else.} \end{cases}$$

Then all the terms in (2) except the s_i terms cancel yielding (1). ∎

Nash's Theorem

Theorem

(Nash) *Every finite game has a mixed strategy Nash equilibrium.*

- Implication: matching pennies game necessarily has a mixed strategy equilibrium.
- Why is this important?
  - Without knowing the existence of an equilibrium, it is difficult (perhaps meaningless) to try to understand its properties.
  - Armed with this theorem, we also know that every finite game has an equilibrium, and thus we can simply try to locate the equilibria.

Definitions

- A set in a Euclidean space is compact if and only if it is bounded and closed.
- A set S is **convex** if for any $x, y \in S$ and any $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda) y \in S$.

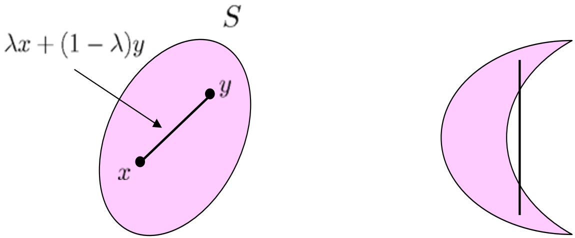

convex set

not a convex set

Weierstrass's Theorem

Theorem

(Weierstrass) *Let A be a nonempty compact subset of a finite dimensional Euclidean space and let $f : A \to \mathbb{R}$ be a continuous function. Then there exists an optimal solution to the optimization problem*

$$\begin{aligned} \text{minimize} \quad & f(x) \\ \text{subject to} \quad & x \in A. \end{aligned}$$

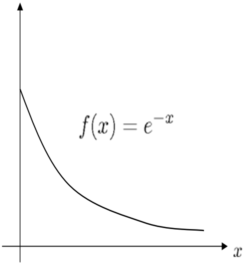

$\min_{x \geq 0} \; e^{-x} = 0$

There exists no optimal x that attains it

Kakutani's Fixed Point Theorem

Theorem

(Kakutani) *Let A be a non-empty subset of a finite dimensional Euclidean space. Let $f : A \rightrightarrows A$ be a correspondence, with $x \in A \mapsto f(x) \subseteq A$, satisfying the following conditions:*

- *A is a compact and convex set.*
- *$f(x)$ is non-empty for all $x \in A$.*
- *$f(x)$ is a convex-valued correspondence: for all $x \in A$, $f(x)$ is a convex set.*
- *$f(x)$ has a closed graph: that is, if $\{x^n, y^n\} \to \{x, y\}$ with $y^n \in f(x^n)$, then $y \in f(x)$.*

Then, f has a fixed point, that is, there exists some $x \in A$, such that $x \in f(x)$.

Kakutani's Fixed Point Theorem—Graphical Illustration

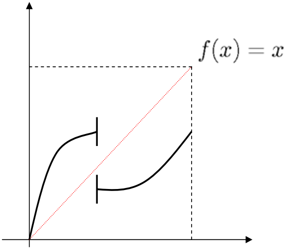

$f(x)$ is not convex-valued

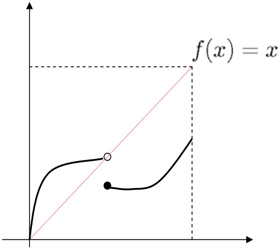

$f(x)$ does not have a closed graph

Proof of Nash's Theorem

- Recall that σ^* is a (mixed strategy) Nash Equilibrium if for each player i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \qquad \text{for all } \sigma_i \in \Sigma_i.$$

- Define the best response correspondence for player i $B_i : \Sigma_{-i} \rightrightarrows \Sigma_i$ as

$$B_i(\sigma_{-i}) = \left\{\sigma_i' \in \Sigma_i \mid u_i(\sigma_i', \sigma_{-i}) \geq u_i(\sigma_i, \sigma_{-i}) \text{ for all } \sigma_i \in \Sigma_i\right\}.$$

- Define the set of best response correspondences as

$$B(\sigma) = [B_i(\sigma_{-i})]_{i \in \mathcal{I}}.$$

- Clearly

$$B : \Sigma \rightrightarrows \Sigma.$$

Proof (continued)

- The idea is to apply Kakutani's theorem to the best response correspondence $B : \Sigma \rightrightarrows \Sigma$. We show that $B(\sigma)$ satisfies the conditions of Kakutani's theorem.

- Σ is compact, convex, and non-empty.
  - By definition
    $$\Sigma = \prod_{i \in \mathcal{I}} \Sigma_i$$
    where each $\Sigma_i = \{x \mid \sum_j x_j = 1\}$ is a *simplex* of dimension $|S_i| - 1$, thus each Σ_i is closed and bounded, and thus compact. Their product set is also compact.

- $B(\sigma)$ is non-empty.
  - By definition,
    $$B_i(\sigma_{-i}) = \arg\max_{x \in \Sigma_i} u_i(x, \sigma_{-i})$$
    where Σ_i is non-empty and compact, and u_i is linear in x. Hence, u_i is continuous, and by Weirstrass's theorem $B(\sigma)$ is non-empty.

Proof (continued)

3. $B(\sigma)$ is a convex-valued correspondence.
   - Equivalently, $B(\sigma) \subset \Sigma$ is convex if and only if $B_i(\sigma_{-i})$ is convex for all i. Let $\sigma_i', \sigma_i'' \in B_i(\sigma_{-i})$.
   - Then, for all $\lambda \in [0,1] \in B_i(\sigma_{-i})$, we have

$$u_i(\sigma_i', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i,$$

$$u_i(\sigma_i'', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i.$$

   - The preceding relations imply that for all $\lambda \in [0,1]$, we have

$$\lambda u_i(\sigma_i', \sigma_{-i}) + (1-\lambda) u_i(\sigma_i'', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i.$$

   By the linearity of u_i,

$$u_i(\lambda \sigma_i' + (1-\lambda)\sigma_i'', \sigma_{-i}) \geq u_i(\tau_i, \sigma_{-i}) \qquad \text{for all } \tau_i \in \Sigma_i.$$

   Therefore, $\lambda \sigma_i' + (1-\lambda)\sigma_i'' \in B_i(\sigma_{-i})$, showing that $B(\sigma)$ is convex-valued.

Proof (continued)

4. $B(\sigma)$ has a closed graph.
   - Suppose to obtain a contradiction, that $B(\sigma)$ does not have a closed graph.
   - Then, there exists a sequence $(\sigma^n, \hat{\sigma}^n) \to (\sigma, \hat{\sigma})$ with $\hat{\sigma}^n \in B(\sigma^n)$, but $\hat{\sigma} \notin B(\sigma)$, i.e., there exists some i such that $\hat{\sigma}_i \notin B_i(\sigma_{-i})$.
   - This implies that there exists some $\sigma'_i \in \Sigma_i$ and some $\epsilon > 0$ such that

$$u_i(\sigma'_i, \sigma_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i}) + 3\epsilon.$$

   - By the continuity of u_i and the fact that $\sigma^n_{-i} \to \sigma_{-i}$, we have for sufficiently large n,

$$u_i(\sigma'_i, \sigma^n_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) - \epsilon.$$

Proof (continued)

- [step 4 continued] Combining the preceding two relations, we obtain

$$u_i(\sigma'_i, \sigma^n_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i}) + 2\epsilon \geq u_i(\hat{\sigma}^n_i, \sigma^n_{-i}) + \epsilon,$$

  where the second relation follows from the continuity of u_i. This contradicts the assumption that $\hat{\sigma}^n_i \in B_i(\sigma^n_{-i})$, and completes the proof.

- The existence of the fixed point then follows from Kakutani's theorem.
- If $\sigma^* \in B(\sigma^*)$, then by definition σ^* is a mixed strategy equilibrium.

6.254 Game Theory with Engineering Applications
Spring 2010