6.972: Game Theory February 14, 2008

Lecture 4: Correlated Rationalizability, Correlated Equilibrium

Lecturer: Asu Ozdaglar

1 Review

- **Strategic form game** $\langle \mathcal{I}, (S_i), (u_i) \rangle$.

- **Pure strategy Nash equilibrium**: We showed that for finite games, there does not always exist a pure strategy Nash equilibrium.

- **A mixed strategy** $\sigma_i \in \Sigma_i$ is defined as a randomization over available actions of a player (i.e., it is a probability measure over S_i). In defining the Nash equilibrium, we assume that the players randomize independently, i.e., a mixed strategy Nash equilibrium $\sigma = \prod_j \sigma_j$ is a mixed strategy profile in which the players mix independently.

- **Rationalizability**: We defined a never-best response and argued that rational players will not play "never-best responses". This motivates us to consider iteratively eliminating never-best responses from pure strategy sets of players. One interesting question is how this elimination procedure compares with the "iterated strict dominance" that we studied earlier, i.e., whether this process allows us to eliminate more strategies from the play of the game based on rationality assumptions. We have shown in the last lecture that while every strictly dominated strategy is a never-best response, the converse need not be true, i.e., a never-best response need not be strictly dominated (see the example in the Lecture 3 notes). Hence, eliminating never-best responses refines iterated strict dominance.

  In the next section, we show however that, if we allow the players to have correlated beliefs, i.e., it need not be a mixed strategy profile for the opponents, then the process of eliminating never-best responses and strictly dominated strategies yields the same set of strategies at the end.

2 Correlated Rationalizability

In this section, we allow a player to believe that the other players' actions are correlated—in other words, that the other players might be in a coalition and thus pick their strategies together instead of individually. To capture this idea, we slightly modify our definition of a belief.

Definition 2.1 *A* belief *of player i about the other players' actions is a probability measure over the set S_{-i}, which we denote as $\Delta(S_{-i})$.*

Note that we allow correlation in our belief. Recall, $\Delta(S)$ denotes a probability distribution over S. One possible type of distribution is the product distribution $S_1 \times S_2 \times ... \times S_I$, which denotes independent mixing between the I players. In general however, the distribution $\Delta(S)$ allows correlation in the strategies of players.

Definition 2.2 *An action $s_i \in S_i$ is a rational action if there exists a belief $\alpha_{-i} \in \Delta(S_{-i})$ such that s_i is a best response to α_{-i}.*

To define rationalizable strategies, we eliminate actions that are not best responses to any belief, or in other words, we eliminate actions that are *never-best responses.* Let us next recall the definitions of "never-best response" strategy and "strictly dominated" strategy.

Definition 2.3

1. *An action s_i is a* never-best response *if for all beliefs α_{-i} there exists $\sigma_i \in \Sigma_i$ such that $u_i(\sigma_i, \alpha_{-i}) > u_i(s_i, \alpha_{-i})$.*
2. *An action s_i is strictly dominated if there exists $\sigma_i \in \Sigma_i$ such that $u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.*

It is straightforward to show that a strictly dominated action is a never-best response. Does the other direction hold? We have shown in the previous lecture that it doesn't hold if beliefs are independent mixings. In this lecture, we will show that this direction holds for correlated beliefs.

3 Strict Dominance & Correlated Rationalizability

We first formally define the process of iterative elimination of strictly dominated strategies.

Algorithm 3.1 (Strict Dominance Iteration) *Let* $S_i^0 = S_i$ *and* $\Sigma_i^0 = \Sigma_i$. *For each player* $i \in \mathcal{I}$ *and for each* $n \geq 1$, *we define* S_i^n *as*

$$\begin{aligned} S_i^n = \{s_i \in S_i^{n-1} \mid \text{ there is no } \quad & \sigma_i \in \Sigma_i^{n-1} \text{ such that} \\ & u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}^{n-1}\}. \end{aligned}$$

Independently mix over S_i^n *to get* Σ_i^n. *Let* $D_i^\infty = \cap_{n=1}^\infty S_i^n$. *We refer to the set* D_i^∞ *as the set of strategies of player* i *that survive iterated strict dominance.*

We next formally define the process of iterative elimination of never-best response strategies. Recall our notation that $\Delta(A)$ denotes the set of probability measures over the set A [implying that the set $\Delta(S_{-i})$ denotes the set of all probability measures over the set S_{-i}, including independent mixings].

Algorithm 3.2 (Correlated Rationalizability Iteration) *Start with* $\tilde{S}_i^0 = S_i$. *Then, for each player* $i \in \mathcal{I}$ *and for each* $n \geq 1$,

$$\begin{aligned} \tilde{S}_i^n = \{s_i \in \tilde{S}_i^{n-1} \mid \text{ there exists } \quad & \alpha_{-i} \in \Delta(\tilde{S}_{-i}^{n-1}) \text{ such that} \\ & u_i(s_i, \alpha_{-i}) \geq u_i(s_i', \alpha_{-i}) \quad \text{for all } s_i' \in \tilde{S}_i^{n-1}\}. \end{aligned}$$

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.

At an intuitive level, rationalizability asks the question "what players might do", whereas strict dominance asks the question "what players won't do, and what they won't do conditional on other players not doing certain things, etc.". Note that both iterated strict dominance and rationalizability never eliminate any strategy played with positive probability in a Nash equilibrium. Indeed both these concepts could be quite weak. Most games, including many games with a unique Nash equilibrium are not dominance solvable.

In the next proposition, we show the equivalence of iterated strict dominance and correlated rationalizability.

Proposition 3.3 *Consider a strategic form game $\langle \mathcal{I}, (S_i), (u_i) \rangle$ with finite number of players and finite strategy spaces S_i. Then $S_i^n = \tilde{S}_i^n$ for all n.*

Proof. We first note that s_i strictly dominated implies that s_i is a never-best response. Hence $\tilde{S}_i^n \subseteq S_i^n$ for all n. We next show that if a pure strategy is a never-best response, then it is strictly dominated. We show this by contraposition.

Assume that for some player i the pure strategy $\bar{s}_i \in S_i$ is not strictly dominated. This implies that, for all $\sigma_i \in \Sigma_i^n$, there exists some $s_{-i} \in S_{-i}^n$ such that

$$u_i(\sigma_i, s_{-i}) \le u_i(\bar{s}_i, s_{-i}). \tag{1}$$

Consider the set S_{-i}. Letting $|S_{-i}| = d$, we can write $S_{-i} = \{(s_{-i})_1, \ldots, (s_{-i})_d\}$. We define the set

$$\bar{U}_i = \Big\{ x \in \mathbb{R}^d \mid \text{there exists some } \sigma_i \in \Sigma_i \text{ such that } x_j \le u_i(\sigma_i, (s_{-i})_j) \text{ for all } j \in \{1, \ldots, d\} \Big\}.$$

For each mixed strategy σ_i of player i, we have a point in $\bar{U}_i$, which specifies player i's payoff for every pure strategy of its opponents.

It is straightforward to verify that the set U_i is nonempty and convex. Moreover, it follows by Eq. 1, that the vector $\bar{x} = [u_i(\bar{s}_i, (s_{-i})_j)]_{j \in \{1,\ldots,d\}}$ is not an interior point of $\bar{U}_i$. By the Supporting Hyperplane Theorem, there exists some $\alpha_{-i} \in \mathbb{R}^d \neq 0$ such that

$$\alpha_{-i}^T \, \bar{x} \ge \alpha_{-i}^T \, x \qquad \text{for all } x \in \bar{U}_i.$$

Since $\alpha_{-i} \neq 0$, by appropriate normalization, we can assume that $\sum_{j=1}^d (\alpha_{-i})_j = 1$. Moreover, in view of the structure of the set $\bar{U}_i$, it can be seen that $\alpha_{-i} \ge 0$. Hence α_{-i} is a (correlated) belief. From the preceding relation, we have

$$\alpha_{-i}^T \Big[u_i(\bar{s}_i, s_j) \Big]_{j \in \{1,\ldots,d\}} \ge \alpha_{-i}^T \Big[u_i(\sigma_i, s_j) \Big]_{j \in \{1,\ldots,d\}} \qquad \text{for all } \sigma_i \in \Sigma_i,$$

which implies that

$$u_i(\bar{s}_i, \alpha_{-i}) \geq u_i(\sigma_i, \alpha_{-i}) \qquad \text{for all } \sigma_i \in \Sigma_i.$$

Hence $\bar{s}_i$ is a best response to the belief α_{-i}. This shows that $S_i^n \subseteq \tilde{S}_i^n$ for all n and completes the proof. ■

Final Lesson:

- 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.
- The latter two sets are equal when beliefs are allowed to be correlated.

4 Correlated Equilibrium

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

Example [Battle of the Sexes (BoS)]: Suppose in the BoS game, 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 $(3/2, 3/2)$ that is not a Nash equilibrium payoff.

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 clever randomizations, as illustrated in the next example.

Example [Traffic Intersection Game]: Consider a game where two cars arrive at an intersection simultaneously. Row player has the option to play U or D, and the Column player 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. Then we have

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

There is a unique mixed strategy equilibrium with expected payoff (5/2,5/2).

Case 1: Assume that there is a publicly observable random variable (such as a fair coin) such that

with probability 1/2 (H) → Player 1 plays U and Player 2 plays L

with probability 1/2 (T) → Player 1 plays D and Player 2 plays R.

Then the expected payoff increases to (3,3).

Claim: It is an "equilibrium" to follow the coin.

Proof: We need to check the incentives of each player:

- If Row sees an H, he believes that Column will play L → he plays U (similar when he sees a T).
- If Column sees an H, he believes that Row will play D → he plays R (similar when he sees a T).

Case 1: 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 $x = 1$ → tell Row to play U, Column to play L.
- If $x = 2$ → tell Row to play D, Column to play L.

- If $x = 3 \rightarrow$ tell Row to play D, Column to play R.

Claim: It is an "equilibrium" to follow the mediator's advice.

Proof: We need to check the incentives of each player.

- If Row hears U, he believes Column will play $L \rightarrow$ play U.
- If Row hears D, he believes Column will play L, R with probability 1/2,1/2 $\rightarrow$ play D.
- If Column hears L, he believes Row will play U, D with probability 1/2,1/2 $\rightarrow$ play L.
- If Column hears R, he believes Row will play $D \rightarrow$ play R.

Thus the players will follow the mediator's suggestion. With the mediator, the expected payoffs are $(10/3, 10/3)$, strictly higher than what the players could get by randomizing between Nash equilibria.

The above example suggests the players may get higher expected payoff if, for example, the game can be transformed to a game with a signaling device. Then, the strategies in the game would be functions from signals to pure strategies. The formal analysis of such a mechanism leads to the notion of the *correlated equilibrium.* Please read Chapter 2 of Fudenberg-Tirole for this analysis.

For our purposes, what really matters in correlated equilibrium is the induced (potentially correlated) probability distribution over the strategy profiles. Therefore, with some abuse of terminology, we will directly define the correlated equilibrium as a probability distribution over the strategy profiles.

Definition 4.1 *A correlated equilibrium is a probability distribution function $p(\cdot)$ over the pure strategies $S = \prod_{i=1}^{n} S_i$ such that for all $i \in \mathcal{I}$ and every $s_i, t_i \in S_i$,*

$$\sum_{s_{-i} \in S_{-i}} p(s_i, s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i}} p(s_i, s_{-i}) u_i(t_i, s_{-i}).$$

In other words,

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

for all $i \in \mathcal{I}$ and $s_i, t_i \in S_i$.

Note that a correlated equilibrium can also be equivalently defined in terms of conditional distributions: A probability distribution $p(\cdot)$ is a correlated equilibrium if, for every player i and ever $s_i \in S_i$ with $p(s_i) > 0$, we have

$$\sum_{s_{-i} \in S_{-i}} p(s_{-i}|s_i) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} p(s_{-i}|s_i) u_i(t_i, s_{-i}), \qquad \text{for all } t_i \in S_i.$$

Intuition: A "universal device" chooses a strategy $s \in S$ at random according to distribution $p(\cdot)$. Each player receives the "recommendation" $s_i \in S_i$. Given his recommendation s_i, player i can form the conditional probability distribution of the opponents' recommendations. The distribution $p(\cdot)$ is a correlated equilibrium if no player can profit by unilaterally deviating from his recommendation, assuming all other players play according to their recommendation.

A correlated equilibrium results if each player playing the recommended strategy does not have any incentive to deviate.

4.1 Properties of a Correlated Equilibrium

Property 1: Any mixed Nash equilibrium is a correlated equilibrium.

Property 2: With a publicly observable random variable, the players can get any payoff vector in the "convex hull" of the Nash equilibrium payoffs. The convex hull of a set X, denoted by $\text{conv}(X)$, is the set of all convex combinations of the elements of set X. In particular, the convex hull of finitely many points $\{x_1, \ldots, x_k\}$ is given by

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

Property 3: Correlated equilibria always exist in finite games. We will show next week that a Nash equilibrium exists in a finite game, which shows the existence of a correlated equilibrium indirectly. We will also exploit the special structure of the correlated equilibrium to show its existence without adhering to fixed point arguments.

Property 4: The sets of correlated equilibrium distributions and payoffs are convex.