6.972: Game Theory February 8, 2005

Lecture 3: The Nash Equilibrium and Mixed Strategies

Lecturer: Asu Ozdaglar

1 The Nash Equilibrium

Recall the Nash Equilibrium defined in the previous lecture

Definition 1 (Nash Equiliubrium) *A pure strategy Nash Equilibrium of a strategic game* $\langle \mathcal{I}, (S_i), (u_i)_{i\in\mathcal{I}} \rangle$ *is an action profile* $s^* \in S$ *s.t.* $\forall i \in \mathcal{I}$ *the following condition holds*

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \forall s_i \in S_i.$$

Definition 2 (Nash Equiliubrium - Restated) *Let* $\langle \mathcal{I}, (S_i), (u_i)_{i\in\mathcal{I}} \rangle$ *be a strategic game. For any* $s_{-i} \in S_{-i}$*, define the best-response function* $B_i(s_{-i})$*,*

$$B_i(s_{-i}) = \{s_i \in S_i \mid u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \quad \forall s_i' \in S_i\}$$

We say that an action profile s^* *is a Nash Equilibrium iff*

$$s_i^* \in B_i(s_{-i}^*) \quad \forall i.$$

For two players, the set of Nash Equilibria is given by the intersection of the best response curves. Recall, for example, the Cournot model. Below we give two other famous examples of games with pure strategy Nash Equilibria.

1.1 Battle of the Sexes

	Ballet	Soccer
Ballet	2, 1	0, 0
Soccer	0, 0	1, 2

Game 1: Battle of the Sexes

The above game has two Nash Equilibria, namely (Ballet, Ballet) and (Soccer, Soccer).

1.2 Second Price Auction (with Complete Information)

The second price auction game is specified as follows:

- An object to be assigned to a player in $\{1, .., n\}$.
- Each player has her own valuation of the object. Player i's valuation of the object is denoted v_i. We further assume that $v_1 > v_2 > ... > 0$.
- The assignment process is described as follows:

- The players simultaneously submit bids, $b_1, .., b_n$.
- The object is given to the player with the highest bid (or to a random player among the ones bidding the highest value).
- The winner pays the *second* highest bid.

- The Utility function for each of the players is as follows:
  - The winner receives her valuation of the object minus the price she pays;
  - Everyone else receives 0.

Bidding the valuation is a NE We start by showing that the strategy to bid $(b_1, .., b_n) = (v_1, .., v_n)$ is a Nash Equilibrium. First note that if indeed everyone plays according to that strategy, then player 1 receives the object and pays a price v_2. This means that her payoff will be $v_1 - v_2 > 0$, and all other payoffs will be 0. Now, player 1 has no incentive to deviate, since her utility can only decrease. Likewise, for all other players $v_i \neq v_1$, it is the case that in order for v_i to change her payoff from 0 she needs to bid more than v_1, in which case her payoff will be $v_i - v_1 < 0$. Therefore no player has an incentive to deviate from the strategy, assuming that everyone else plays according to it. We conclude that this strategy is a Nash Equilibrium.

Are There More Nash Equilibria? We show that the strategy $(v_1, 0, 0, ..., 0)$ is also a Nash Equilibrium. As before, player 1 will receive the object, and will have a payoff of $v_1 - 0 = v_1$. Using the same argument as before we conclude that none of the players have an incentive to deviate, and the strategy is thus a Nash Equilibrium. We leave as an exercise to show that the strategy $(v_2, v_1, 0, 0, ..., 0)$ is also a Nash Equilibrium.

A Weakly Dominating Nash Equilibrium We end the section on second price auctions by showing that the aforementioned strategy of bidding one's valuation (i.e., $b_i = v_i$), in fact, weakly dominates all other strategies. Consider the following picture proof where B^* represents the maximum of all bids excluding player i's bid, $B^* = \max_{j \neq i} b_j$, and $v*$ is player i's valuation. The vertical axis is utility. The first graph shows the payoff for bidding one's valuation. In the second graph, which represents the case when a player bids lower than their valuation, notice that whenever $b_i \leq B^* \leq v^*$, player i receives utility 0 because she loses the auction to whoever bid B^*. If she would have bid her valuation, he would have positive utility in this region (as depicted in the first graph). Similar analysis is made for the case when a player bids more than their valuation.

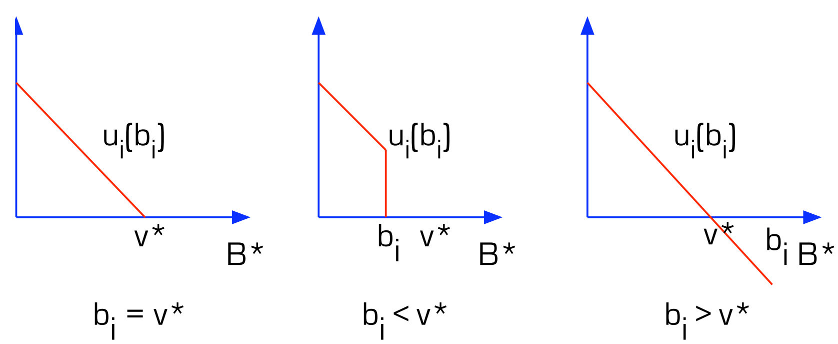

Note that, interestingly, the above argument shows that that there exist Nash Equilibria that are weakly dominated (by other Nash Equilibria).

2 Mixed Strategies

Recall the game of Matching Pennies from the previous lecture.

	Heads	Tails
Heads	$1, -1$	$-1, 1$
Tails	$-1, 1$	$1, -1$

Matching Pennies

It is easy to see that this game does not have a pure Nash Equilibrium (for every pure strategy in this game, both parties have an incentive to deviate). However, if we allow the players to *randomize* over their choice of actions we can find a Nash equilibrium. Assume that player 1 picks "Head" with probability p and "Tail" with probability $1 - p$, and that player 2 picks both "Head" and "Tail" with probability $\frac{1}{2}$. Then, with probability

$$\frac{1}{2}p + \frac{1}{2}(1-p) = \frac{1}{2}$$

player 1 will receive 1. Likewise she will receive -1 with the same probability. This means that if player 2 plays the strategy $(\frac{1}{2}, \frac{1}{2})$ then no matter what strategy player 1 chooses, she will get the same output. Due to the symmetry of the game, we conclude that the randomized strategy $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ is stable. We call such a state, a mixed strategy Nash Equilibrium. In order to formalize this notion we will require some new notation.

2.1 Notation

1. Let Σ_i represent the set of probability distributions over S_i, the set of actions for player i.
2. Let $\sigma_i \in \Sigma_i$ represent a mixed strategy for player i, which is a probability mass function over pure strategies, $s_i \in S_i$.

Assume that a player's mixed strategies are independent randomizations—in other words, all of the players use independent random coins when they sample from their mixed strategies. In our notation, we say $\sigma \in \Sigma = \prod_{i=1}^{I} \Sigma_i$.

Assume that the sets S_i are finite. Let $\sup(\sigma_i)$ denote the support of σ_i, defined as the set $\{s_i \in S_i \mid \sigma_i(s_i) > 0\}$. In other words, since σ_i is a distribution, its support is the set of strategies which are assigned positive probability.

The payoff of a mixed strategy corresponds to the expected value of the pure strategy profiles in its support. More precisely, we can say

$$u_i(\sigma) = \sum_{s \in S} \left(\prod_{j=1}^{I} \sigma_j(s_j) \right) u_i(s).$$

2.2 Definition of Mixed Strategy Nash Equilibrium

Definition 3 *A mixed strategy profile σ^* is a (mixed strategy) Nash Equilibrium if for each player i and for each $\sigma_i \in \Sigma_i$,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*).$$

Note that since $u_i(\sigma_i', \sigma_{-i}^*) = \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, \sigma_{-i}^*)$, it is sufficient to check only *pure* strategy "deviations" when determining whether a given profile is a Nash equilibrium. This leads to the following proposition:

Proposition 1 *A mixed strategy profile σ^* is a (mixed strategy) Nash Equilibrium if for each player i and for each $s_i \in S_i$*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*).$$

2.3 A Useful Lemma in Calculating the Mixed Strategy NE

Lemma 1 *$\sigma^* \in \Sigma$ is a Nash Equilibrium if and only if, for each player $i \in \mathcal{I}$, the following two conditions hold:*

1. *The expected payoff given σ_{-i}^* to every $s_i \in supp(\sigma_i^*)$ is the same.*
2. *The expected payoff given σ_{-i}^* to the actions s_i which are not in the support of σ_i^* must be less than or equal to the expected payoff described in (1).*

Remark: We note that the lemma also extends to the continuous case.

Intuitively, the lemma means that for a player i, every action in the support of a Nash equilibrium is a best response to σ_{-i}^*. This lemma follows from the fact that if the strategies in the support have different payoffs, then it would be better to just take the pure strategy with the highest expected payoff. This would contradict the assumption that σ^* is a Nash equilibrium. Using the same argument, it follows that the pure strategies which are not in the support must have lower (or equal) expected payoffs. More formally,

Proof: Let σ^* be a mixed strategy Nash equilibrium, and let $E_i^* = u_i(\sigma_i^*, \sigma_{-i}^*)$ denote the expected utility for player i. Then, by Proposition 1 above, it holds that $E_i^* \geq u_i(s_i, \sigma_{-i}^*)$ for all $s_i \in S_i$. (which already proves the second part of the lemma). We proceed to show the first part of the lemma, namely that $E_i^* = u_i(s_i, \sigma_{-i}^*)$ for all $s_i \in \text{supp}(\sigma_i^*)$. Assume, for contradiction that this is not the case, i.e., that there exist an action $s_i' \in \text{supp}(\sigma_i^*)$ such that $u_i(s_i', \sigma_{-i}^*) > E_i^*$. This means that

$$E_i^* = \sum_{s_i \in S_i} \sigma_i^*(s_i) u_i(s_i, \sigma_{-i}^*) > E_i^*,$$

since $u_i(s_i, \sigma_{-i}^*) \geq E_i^*$ for all $s_i \in S_i$, except for s_i' where the the inequality is strict. This is a contradiction. $\square$

Application of the Lemma Let us return to the Battle of the Sexes Game.

	Ballet	Soccer
Ballet	$2,1$	$0,0$
Soccer	$0,0$	$1,2$

Battle of the Sexes

Recall that this game has 2 pure Nash Equilibria. We show that it has one *unique* mixed strategy Nash Equilibrium. First, note that by using the Lemma (and inspecting the game tableau) it is easy to see that there are no Nash Equilibria where only one of the players randomizes over its actions. Now, assume instead that player 1 chooses the action "Ballet" with probability p and "Soccer" with probability $1-p$, and that player 2 picks "Ballet" with probability q and "Soccer" with probability $1-q$. Using the Lemma on player 1's actions, we obtain the following equation:

$$2q + 0 * (1-q) = 0 * q + 1 * (1-q)$$

Next, applying the lemma on player 2's actions, we obtain:

$$1 * p + 0 * (1-p) = 0 * p + 2 * (1-p)$$

We conclude that the only possible mixed strategy Nash Equilibrium is when $q = \frac{1}{3}$ and $p = \frac{2}{3}$.

3 Strict Dominance by a Mixed Strategy

Consider the following game.

	Left	Right
Up	$2,0$	$-1,0$
Middle	$0,0$	$0,0$
Down	$-1,0$	$2,0$

Game 4: Strict Dominance by Mixed Strategies

Note that Player 1 has no pure strategies that strictly dominate Middle. However, the expected outcome of the strategy $(\frac{1}{2}, 0, \frac{1}{2})$ is strictly higher than the outcome of playing Middle. We thus extend the notion of strict domination also to mixed strategies, and say that Middle is strictly dominated by the strategy $(\frac{1}{2}, 0, \frac{1}{2})$. More formally,

Definition 4 (Strict Domination by Mixed Strategies) *An action s_i is strictly dominated if there exists a mixed strategy $\sigma_i' \in \Sigma_i$ such that $u_i(\sigma_i', s_{-i}) > u_i(s_i, s_{-i})$, for all $s_{-i} \in S_{-i}$.*

Remark: It is shown in the text book that a Nash Equilibrium can never be strictly dominated. However, as we saw in the Second Price Auction game, there do exist Nash Equilibria that are weakly dominated.

4 Rationalizability

In the Nash equilibrium concept, each player's action is optimal *conditioned* on the *belief* that the other players also play their Nash equilibrium strategy (i.e., the Nash Equilibrium strategy is only optimal if the belief about the other player is correct). Let us now consider a different solution concept in which a player's belief about the other players' actions in not assumed to be correct, but rather, simply constrained by rationality. To begin, let us define a belief.

Definition 5 *A* belief *of player i about the other players' actions is a probability measure over the set S_{-i}, which we denote as $\Delta(S_{-i})$.*

Remark: Our definition allows 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 picking their strategies together instead of individually.

As an example, consider the following game:

	Q	F
Q	4, 2	0, 3
X	1, 1	1, 0
F	3, 0	2, 2

If player 1 believes that player 2 will play Q, then playing Q is rational for player 1 since 4 is better than the other outcomes 1 and 3 respectively. On the other hand, notice that playing X is *never a best response*, regardless of what strategy is chosen by the other player, since playing F always results in better payoffs ($3 > 2, 2 > 1$). In general, we can argue that a strictly dominated strategy will never be a best response, regardless of a player's beliefs about the other players' actions.

Does the converse hold? That is, if a strategy is *never a best response* for any belief, then is it strictly dominated? The answer depends on our definition of "belief."

Remark: In this lecture, the formal notion of *rationalizability* is not offered. However, we shall intuitively describe it as iterated removal of never-best responses.

4.1 Never-Best Response Does Not Imply Strict Dominance

Let us assume that the probability measure over S_{-i} must consist of an *independent mix*—in other words, that beliefs do not allow other players' actions to be correlated.

In this case, the following example illustrates a never-best response which is not strictly dominated. Consider the following three-player game in which all of the player's payoffs are the same. Player 1 chooses C or D, player 2 chooses A or B and player 3 chooses M_i for $i = 1, 2, 3, 4$.

	A	B
C	8	0
D	0	0

M_1

	A	B
C	4	0
D	0	4

M_2

	A	B
C	0	0
D	0	8

M_3

	A	B
C	3	3
D	3	3

M_4

First, observe that playing M_2 is never a best response to any mixed strategy of players 1 and 2. To show this, let p represent the probability with which player 1 chooses A and let q represent the probability that player 2 chooses C. (Remark: we explicitly assume here that p and q are independent.) The payoff for player 3 when she plays M_2 is

$$u_3(M_2, p, q) = 4pq + 4(1-p)(1-q) = 8pq + 4 - 4p - 4q$$

Suppose, by contradiction, that this is a best response for some choice of p, q. This implies the following inequalities:

$$\begin{aligned} 8pq + 4 - 4p - 4q &\geq u_3(M_1, p, q) = 8pq \\ &\geq u_3(M_3, p, q) = 8(1-p)(1-q) = 8 + 8pq - 8p - 8q \\ &\geq u_3(M_4, p, q) = 3 \end{aligned}$$

By reducing the top two equations, we have the following inequalities:

$$\begin{aligned} p + q &\leq 1 \\ p + q &\geq 1 \end{aligned}$$

Thus $p + q = 1$, and substituting into the third inequality, we have $pq \geq 3/8$. Substituting again, we have $p^2 - p + \frac{3}{8} \leq 0$ which has no positive roots since the left side factors into $(p - \frac{1}{2})^2 + (\frac{3}{8} - \frac{1}{4})$.

On the other hand, by inspection, we can see that M_2 is not strictly dominated, since in some cases, it has a higher payoff than M_4.

4.2 Never-best Response Implies Strict Dominance

In the more general case, when we assume that a belief can also include the situation in which other players' strategies are correlated, we can show that strict dominance and never-best responses are equivalent notions.

The proof will be given in the next lecture. We now recall some preliminaries.

Some Preliminaries

Definition 6 (Convex Set) *A convex set is a set C such that $\forall x, y \in C$, the point $\lambda x + (1 - \lambda y) \in C$ for all $\lambda \in [0, 1]$.*

Intuitively, if you take two points in a convex set and draw a line between them, then all of the points on the line must also be contained in the set.

Definition 7 (Compact Set) *A compact set is a set, which is closed and bounded.*

(Remark: The formal definition of compact set is different, but equivalent to the one stated above. A closed set is one in which every sequence of points in the set converges to a point which is also in the set, and a bounded set is one which is contained in a ball of some finite radius.)

Definition 8 (Interior Point) *A point $x \in A$ is an* interior point *of A if there exists a neighborhood of x which is contained in A.*

A *boundary point* of a set C is a point $x \in C$ which is not in the interior of C.

A *supporting hyperplane* to a set C at a boundary point $x_0 \in C$ is the set

$$\{x \mid a^T x = a^T x_0\},$$

where a is a nonzero vector such that $a^T y \leq a^T x_0$ for all $y \in C$.

With a picture, a supporting hyperplane is a plane which is "tangential" to the set C at the boundary point x_0. A useful theorem which often comes up in analysis is that every convex set has a supporting hyperplane.

Theorem 1 (Supporting Hyperplane Theorem) *Let C be a non-empty convex set in R^n. For every boundary point $\bar{x}$ of C, there exists an $a' \in R$ where $a' \neq 0$ such that*

$$a'\bar{x} \leq a'x \quad \forall x \in C$$

The hyperplane $H = \{x \mid a'x = a'\bar{x}\}$ is called the supporting hyperplane.

For a proof, see the posted handout on the Supporting Hyperplane Theorem. The following is a pictorial representation of the theorem.