Game Theory and Algorithms*
Lecture 2: Nash Equilibria and Examples

February 24, 2011

Summary: We introduce the *Nash Equilibrium*: an outcome (action profile) which is stable in the sense that no player has incentive to deviate. (The concept is most famous since a *mixed* equilibrium exists in every finite strategic game, but in this lecture we consider only *pure* equilibria.) We discuss the relation between iterated elimination and Nash equilibria. Then we give some examples of games where Nash equilibria explain the most plausible outcomes: duopoly (price competition by 2 firms), elections, and auctions. To easily explain the last two, we define *weakly dominant strategies*. Finally, we mention *potential games*, which always have Nash equilibria.

1 Best Responses and Nash Equilibria

In the last lecture we saw that for some games, *iterated elimination (of strictly dominated strategies)* leaves just a single outcome, which is arguably the only plausible outcome (if the players seek to maximize their own utility, and are greedy). The Nash equilibrium is another *solution concept* for strategic games.

Definition 1. Let a_{-i} be a partial action profile for all players but i. Action a_i is a *best response to* a_{-i} if for all actions a'_i of player i,

$$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}).$$

Note that there is always at least one best response (if A_i is finite), and possibly more.

Definition 2. The action profile a is a *Nash equilibrium* if for each player i, a_i is a best response to a_{-i}.

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

Equivalently, an action profile is a Nash equilibrium under the following conditions: if each player i believes everyone else will play according to the profile, then player i can maximize their own utility by also playing according to the profile. (It's possible that some players could have additional actions that would also maximize their utility.) Sometimes we will write NE for "Nash equilibrium."

Nash introduced this concept in 1951, but it generalizes solution concepts considered earlier for particular games by Cournot (1838, oligopoly), Bertrand (1883, oligopoly), von Neumann (1926, zero-sum two-player games), and Hotelling (1929, voting).

For example, let us determine the Nash equilibria of the Prisoners' dilemma game.

p1 \p2	Q	F
Q	2,2	0,3
F	3,0	1,1

For both players, let us compute all of their best responses to their opponent's actions. The following diagram illustrates this idea, where we put a star over $u_i(a)$ when a_i is a best response to a_{-i}. For example, when player 1 chooses Q, the best response by player 2 is F, so in the (Q, F) outcome we change 0, 3 to $0, 3^*$.

p1 \p2	Q	F
Q	2,2	$0, 3^*$
F	$3^*, 0$	$1^*, 1^*$

The definition of a Nash equilibrium is an outcome where each player is using a best response to the other players. So, we conclude that this game has exactly one Nash equilibrium, (F, F).

In any game given explicitly by a table (even for more than 2 players), this starring method can be used to find all the NE. More generally, even if there are infinitely many actions, to find the NE it suffices to determine the intersections of the n sets $\{(a_i, a_{-i}) \mid a_i \text{ is a best response to } a_{-i}\}$, for $i = 1, \ldots, n$.

1.1 Properties

Q: Is it always true that a game has exactly one Nash equilibrium?
A: No: there can be zero equilibria, or more than one equilibrium. The taxpayers' game is an example with no NE. Another example is the game "matching pennies:" two players each choose *heads* or *tails*; player 1 wins a dollar from player 2 if they make the same choices, otherwise player 1 loses a dollar to player 1.

p1 \p2	H	T
H	$1^*, -1$	$-1, 1^*$
T	$-1, 1^*$	$1^*, -1$

p1 \p2	B	S
B	$2^*, 1^*$	0,0
S	0,0	$1^*, 2^*$

The second example is "Bach or Stravinsky" which has two NE. The story is that two friends want to see a concert, and there are two concerts, one by Bach and one by Stravinsky. Player

1 prefers Bach, player 2 prefers Stravinsky, but if they cannot agree on which concert to visit, they both stay home, which is the least preferred outcome. *Coordination games* typically refer to games like this where the players need to have external coordination of their actions in order to get a reasonably good solution.

Nash equilibria have the following relation to the "iterated elimination" algorithm we saw earlier.

Exercise. Show that when we perform iterated elimination of strictly dominated strategies, all Nash equilibria survive (i.e., for every Nash equilibrium a and each player i, the algorithm never deletes action a_i).

Exercise. Show that if the final result of iterated elimination is a single outcome, then that outcome is a Nash equilibrium of the original game. [We solved this exercise in lecture.]

1.2 Discussion

The NE concept is obviously not able to predict the outcome of a game (since there can be 0 or ≥ 2 of them) and it doesn't typically help us to figure out a "good" strategy for ourself. Nonetheless, here are some ways that Nash equilibria arise naturally.

- The definition is the weakest notion of a *predictable state*: if we make a global prediction that all players will abide by a Nash equilibrium, and all players believe that their opponents will abide by this plan, then each player has no incentive to deviate from this plan.

- Consequently, if a game has any single plausible "obviously" predictable outcome, it should be a Nash equilibrium.

- You can think of a Nash equilibria as a *self-enforcing agreement* or a *stable social convention* similarly: an agreement from which no player has incentive to deviate. This is illustrated by the two equilibria of the following "Drive on which side of the road?" game, modelling two cars driving in opposite directions on the same road.

p1 \p2	L	R
L	0,0	$-1, -1$
R	$-1, -1$	0,0

- Nash equilibria are fixed points of the following process, called *best response dynamics*. First, we start at an arbitrary action profile a. Then, each player i in turn is given a chance to adjust their choice a_i — they want to increase their own utility $u_i(a)$ if possible. We repeat this, giving all players a chance to change. We stop once we go through an entire round of players and nobody wants to change. Then the possible action profiles where this process can terminate are precisely Nash equilibria.
  Q: What happens when we execute best response dynamics in the taxpayer game?

- Similar to best response dynamics, NE can be viewed as the final fixed result of a *learning* or *evolution* process. (There is a more specific notion of an *evolutionarily stable strategy* but it is beyond the scope of our course.)

- Possibly the most important reason is *Nash's theorem*, that every finite game has a *mixed* Nash equilibrium (more on this in a later lecture). Thus, it is a concept that can be applied to any game.

- The NE concept unifies results from many examples studied before Nash. We will see some of these historical examples next.

2 Examples

2.1 Tragedy of the Commons

Consider the game with n companies, each of which owns a factory. It costs 3 dollars to install a pollution-controller in their factory. For each player who does not install a pollution-controller, each player pays 2 dollars. So $A_i = \{\textbf{I}\text{nstall}, \textbf{P}\text{ollute}\}$ and

$$u_i(a) = \begin{cases} -2(\#\text{ of Ps in } a), & a_i = P \\ -2(\#\text{ of Ps in } a) - 3, & a_i = I. \end{cases}$$

What are the Nash Equilibria of this game? Consider the best response functions. Fix a player i and the choices a_{-i} of their opponents, and let k be the number of Ps in a_{-i}. Then $u_i(I, a_{-i})$ is $-2k - 3$, while $u_i(P, a_{-i})$ is $-2k - 2$. So Polluting is always a better response than Installing (in fact P strictly dominates I); there cannot be any Installer in a NE since they have incentive to deviate to P, and we see the all-P outcome is a Nash equilibrium.

The result is a little surprising since the equilibrium outcome gives utility $-2n$ to each player, whereas if they all installed pollution-controllers they would each have utility -3. (Note, for $n = 2$ this is a Prisoners' dilemma.)

2.2 Voter Participation

For a non-trivial example of finding Nash equilibria in a many-player game, we give the following exercise from Martin Obsorne's book.

Exercise. In an election, there are $2k$ players: k of them support the candidate A, and k of them support the other candidate B. Each player can **V**ote or **A**bstain. The candidate with the largest number of voting supporters wins the election (they tie if they have the same number). Each voter gets a payoff of $+2$ if their candidate wins, 0 in a tie, and -2 if their candidate loses. Since the voters are lazy, their utility incurs a charge of -1 if they vote. (So each player's utility function can take on 6 values, the best is $+2$ (winning without voting), and the worst is -2 (voting and losing).) Find the Nash equilibria of this game.

Exercise. Consider the same game with 3 players, where A has 2 supporters and B has 1 supporter. What are the Nash equilibria? (Hint: it is qualitatively different from the previous exercise.)

2.3 Duopoly

We now give a classical example of modelling price competition by two companies. Just like *monopoly* means a sector of the economy controlled by a single company, *duopoly* means one controlled by two companies, and *oligopoly* means control by many companies. There are two prominent historical examples; in both of them, consumers make a choice modelled by a supply-demand curve.

Cournot Oligopoly	Bertrand Oligopoly
- companies pick amounts to sell	- companies pick prices
- consumers determine price	- consumers determine amount to buy

One major difference from the games considered so far is that the A_i will not be finite sets; players will be able to make a continuum of choices.

2.3.1 Cournot Duopoly

In this game, each firm $i = 1, 2$ picks a quantity q_i of goods to produce ($A_i = \mathbb{R}_{\geq 0}$). They produce the same type of good, so we let $Q = q_1 + q_2$ denote the total amount of goods produced. Each company incurs a cost $C_i(q_i)$ to produce their goods. The market is willing to buy all of the goods, at a price per unit $P(Q)$ depending on the total amount produced. The utility to each company is its income minus their costs.

For simplicity, we study the model where C_i and P are as simple as possible:

$$C_i(q_i) = c \cdot q_i, \qquad P(Q) = \max\{0, \alpha - Q\}$$

or in essence, linear production cost, and inverse-linear demand. Thus the game has

$$u_i(q) = P(q_1 + q_2) \cdot q_i - C(q_i) = \begin{cases} q_1 \cdot (\alpha - q_1 - q_2 - c), & \text{if } q_1 + q_2 \leq \alpha \\ -q_1 \cdot c, & \text{otherwise.} \end{cases}$$

We assume $\alpha > c > 0$ (if $\alpha \leq c$, no profit is possible).
Q: If there were only one firm, what would the optimal profit be?
A: In this case pick we would pick q to maximize $q(\alpha - q - c)$, which happens at $q = \frac{\alpha - c}{2}$. The utility would be $\frac{(\alpha-c)^2}{4}$.
Q: Back in the two-player game, what are the Nash equilibria, and the equilibrium profits?

The most direct way to answer this question is to find the *best response functions* for each player: $B_1(q_2) = \arg\max_{q_1 \in \mathbf{R}_{\geq 0}} u_1(q_1, q_2)$ and likewise for $B_2(q_1)$. (These may be set-valued functions in general, or have *no value*, but we ignore this when possible for ease of exposition.) Then, (q_1, q_2) is an equilibrium exactly when $q_1 = B_1(q_2)$ and $q_2 = B_2(q_1)$ both hold: we find the intersection of the graphs of the best response functions.

To compute the best response function, we seek to maximize $u_1(q_1, q_2)$ for a fixed q_2.

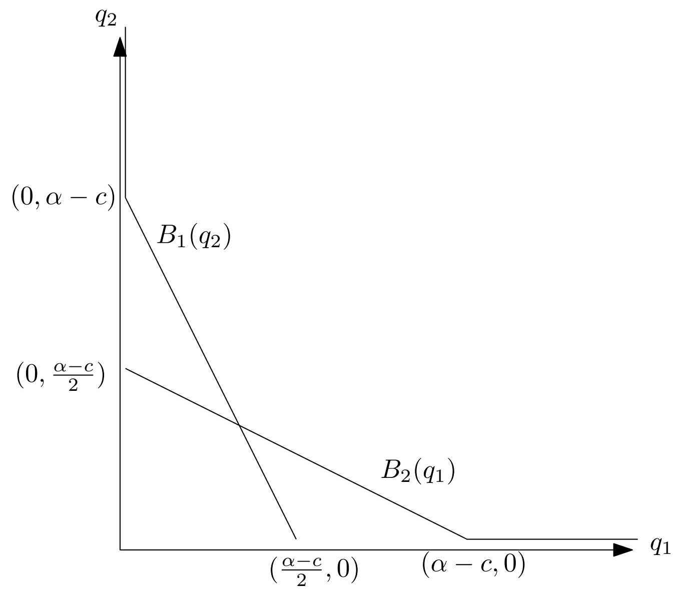

- First, note that if $q_2 \geq \alpha - c$, then $P(Q) \leq c$ no matter how we pick q_1, so a nonzero profit is impossible for firm 1. In fact the unique best response is $q_1 = 0$.

- Otherwise, when $q_2 < \alpha - c$, a positive profit is obtained for any positive $q_1 < \alpha - c$. In this regime the utility is given by $q_1 \cdot (\alpha - q_1 - q_2 - c)$, a quadratic in q_1 maximized at $q_1 = (\alpha - c - q_2)/2$.

Thus, the best response function can be written as $B_1(q_2) = \max\{0, (\alpha - c - q_2)/2\}$.

To compute the intersection of the best response curves, it is helpful to draw a diagram with q_1 on one axis and q_2 on the other (above); it shows that there is exactly one intersection of $q_1 = B_1(q_2)$ and $q_2 = B_2(q_1)$. It occurs at the intersection of the lines $q_1 = (\alpha - c - q_2)/2$ and $q_2 = (\alpha - c - q_1)/2$. Solving, this gives the unique Nash equilibrium

$$q_1 = q_2 = (\alpha - c)/3.$$

We then compute that the equilibrium profit for each firm is $(\alpha - c)^2/9$.

Q: How can *both* of the firms do better?

A: Recall that in the single-player variant, a profit of $\frac{(\alpha-c)^2}{4}$ was possible, which is greater than the total profit $2(\alpha - c)^2/9$ in this equilibrium. Moreover from the consumers' perspective, and for the linear production functions C_i, two companies "look" the same as one company producing their aggregate amount. They can do better by both producing half of the single-player optimal quantity, i.e. in $q_1 = q_2 = (\alpha - c)/4$ they both increase their profits to $(\alpha - c)^2/8$.

You can make some observations about this result: it is similar to the Prisoners' dilemma; and it suggests that monopolies do better than duopolies.

2.3.2 Bertrand Duopoly

This alternative model of oligopoly was given as a response to Cournot's model. As mentioned earlier, this time the action set for each company is its choice of price p_i (so its price is selected from $A_i = \mathbb{R}_{\geq 0}$). We will assume that the consumers choose to buy *only from the cheapest* company, and that they purchase a quantity stipulated by a *demand function* (in contrast with the inverse-demand function in Cournot's model). Being simple again, we'll pick the demand function

$$D(p) = \max\{0, \alpha - p\}$$

where p is the minimum price. As mentioned, the consumers will buy $D(p)$ units of the good, from the company with the lowest price; in the event of a k-way tie, they will split their demand equally amongst the companies. Again, take the cost of producing q units to be the linear function $c \cdot q$ for a fixed $c < \alpha$. The utility is again the income minus the cost, which we reiterate in a simplified form for clarity,

$$u_1(p_1, p_2) = \begin{cases} 0, & \text{if } p_1 > p_2; \\ (p_1 - c) \cdot \max\{0, \alpha - p_1\}, & \text{if } p_1 < p_2; \\ (p_1 - c) \cdot \max\{0, \alpha - p_1\}/2, & \text{if } p_1 = p_2. \end{cases}$$

(u_2 is symmetric). Here is the main, surprising, difference of this model from Cournot's model:

Theorem 3. *The only Nash equilibrium is* (c, c).

This is surprising when you calculate the profits in equilibrium: both firms are selling exactly at the same cost as the manufacturing cost, and have a profit of zero!

The straightforward way to prove this theorem is to use the same general strategy we used earlier: compute the best response functions for both players, and then find all intersections of their graphs. It is, however, more complicated in this case since the utility functions are not continuous. Here is an outline of a simplified version of the proof, which is easier to see after you finish the proof once the "hard way."

Exercise.

1. Show that (c, c) is a Nash equilibrium.
2. Show that (p_1, p_2) is not a Nash equilibrium if $p_1 < c$ or $p_2 < c$.
3. Show that (p_1, p_2) is not a Nash equilibrium if $p_1 = c$ and $p_2 > c$ (or vice-versa).
4. Show that (p_1, p_2) is not a Nash equilibrium if $p_1 > c$ and $p_2 > c$.