Chapter 9: Nash Equilibrium

Elements of Decision: Lecture Notes of Intermediate Microeconomics

Charles Z. Zheng
Tepper School of Business, Carnegie Mellon University

Last update: April 21, 2020

In this final chapter we introduce games where every player's decision needs to best respond to the other player's decision. In a nutshell, this is the essence of the interactive reasoning in economics, formalized by the concept Nash equilibrium.

1 Battle of the Sexes and Nash equilibrium

The crucial observation of the following game (called Battle of the Sexes) is that what is best to a player depends on which strategy the other player is to choose. For example, the best strategy for player 1 (row player) is U if player 2 (column player) chooses L, and is D if R. Likewise for player 2, L is the best if player 1 chooses U, and is R if player 1 chooses D.

	L	R
U	4, 2	0, 0
D	0, 0	2, 4

Thus, solving a game means no longer locating the outcome that results from a player's rational decision, but rather characterizing the set of stable points that rational players would decide not to deviate from. This is the concept of Nash equilibrium. Formally, a pair (σ_1, σ_2) of strategies, pure or mixed, is said to be a *Nash equilibrium* iff σ_1 maximizes player 1's expected payoff among all mixed strategies[1] for player 1 as long as player 2 sticks to σ_2, and σ_2 maximizes player 2's expected payoff among all of player 2's mixed strategies as long as player 1 sticks to σ_1. In other words, neither player can gain from a *unilateral* deviation from the pair (σ_1, σ_2).

Put differently, you can think of a player's decision as *responding* to a strategy that he expects the other player is to choose, and the response is to choose a mixed strategy to maximize his own expected payoff provided that the other player does choose what the former has expected. Any such a maximum is called a *best response*. A Nash equilibrium is any strategy pair in which each strategy is a best response to the other strategy.

2 How to calculate Nash equilibria

Let us illustrate with the above game. First, calculate the *pure-strategy Nash equilibria*, those consisting of only pure strategies, through tracing through the chain of best responses:

a. According to the above payoff matrix, $U \to L \to U$. That is, if player 1 is expected to play U then player 2 would respond with L $(2 > 0)$, and if player 2 is to play L then player 1 would choose U $(4 > 0)$. Note that the chain $U \to L \to U$ forms a loop $U \to L \to U \to L \to U \to \cdots U$ that involves no other strategy than the pair (U, L). That means (U, L) is a Nash equilibrium.

[1] Recall from Chapter 7 that pure strategies are just special cases of mixed strategies.

b. Let us switch to the other pure strategy, D, that has not been included by the above chain. Again from the payoff matrix we have $D \to R \to D$, which again forms a loop involving no other strategy than the pair (D, R). Thus (D, R) is another Nash equilibrium.

Since all pure strategies have been covered by some chain of best responses, we have found all pure-strategy Nash equilibria. There are two of them: (U, L) and (D, R).

Second, we find out all the Nash equilibria with *totally mixed strategies*, i.e., at least one player employs a mixed strategy such that any pure strategy of his is to be played with a strictly positive probability. Suppose this player is player 1. Then he must be indifferent between the two pure strategies U and D, otherwise he would have spent zero probability on the inferior pure strategy. It follows, according to the payoff matrix, that player 2 must also employ a totally mixed strategy, otherwise player 1 would have responded with the pure strategy that best responds to whichever pure strategy that player 2 is employing (e.g., if player 2 chooses the pure L then player 1 would have chosen U and U only, rather than mixing between U and D). Thus, in the Battle of the Sexes game, any Nash equilibrium with totally mixed strategies is in the form of

$$(\alpha U + (1-\alpha)D, \beta L + (1-\beta)R)$$

for some numbers α and β such that $0 < \alpha < 1$ and $0 < \beta < 1$. We just need to find out what α and β are each equal to.

To that end, recall from the above that player 1 is indifferent between U and L in this equilibrium. In other words, his expected payoff from playing U is equal to his expected payoff from playing D, given that player 2 sticks to $\beta L + (1-\beta)R$. Thus we obtain an equation

$$\beta \cdot 4 + (1-\beta) \cdot 0 = \beta \cdot 0 + (1-\beta) \cdot 2.$$

Likewise, player 2 is indifferent between L and R as long as player 1 sticks to $\alpha U + (1-\alpha)D$:

$$\alpha \cdot 2 + (1-\alpha) \cdot 0 = \alpha \cdot 0 + (1-\alpha) \cdot 4.$$

The two equations are reduced to

$$\begin{aligned} \beta &= 1/3, \\ \alpha &= 2/3. \end{aligned}$$

We thus obtain the only totally mixed strategy Nash equilibrium

$$\left(\frac{2}{3}U + \frac{1}{3}D, \frac{1}{3}L + \frac{2}{3}R \right).$$

In other words, player 1 chooses U with probability $2/3$, and player 2 chooses L with probability $1/3$.

Note that in the above calculation we determine player 1's equilibrium strategy (pinning down the value of α) through player 2's best response to the strategy (indifference between L and R), and determine player 2's equilibrium strategy through player 1's best response (indifference between U and D). That is an example of interactive reasoning, the essence of game theory.

Note also that when player 2 is expected to play the mixed strategy $\frac{1}{3}L + \frac{2}{3}R$, the pure strategies U and D are each a best response for player 1. (In fact, *any* mixed strategy between U and D is a best response for player 1 to $\frac{1}{3}L + \frac{2}{3}R$; verify that yourself.) But then why isn't $(U, \frac{1}{3}L + \frac{2}{3}R)$ a solution of this game? That is because, should player 1 be expected to play U, player 2 would not have played $\frac{1}{3}L + \frac{2}{3}R$ (check it yourself). Hence $(U, \frac{1}{3}L + \frac{2}{3}R)$ is a self-defeating prediction, whereas $\left(\frac{2}{3}U + \frac{1}{3}D, \frac{1}{3}L + \frac{2}{3}R\right)$ is a self-fulfilling prophesy.

3 Best response and the consumer optimization problem

Recall the idea from the zero-sum game chapter that a player's strategy amounts to a contingent payoff vector. In the Battle of Sexes, a strategy for player 1 corresponds to a bundle (v_L, v_R) from player 1's perspective, with v_L his expected payoff when player 2 chooses L, and v_R that when player 2 chooses R. As in that chapter, we can draw player 1's choice set in the v_L-v_R diagram as the segment DU in Figure 1. There, the dashed lines represent player 1's indifference curves in the case where he expects player 2 to play a mixed strategy $\beta L + (1-\beta)R$.

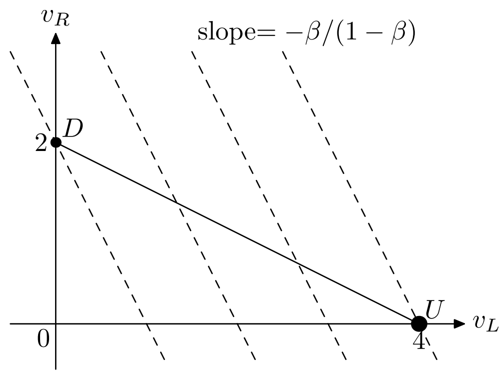

Figure 1: Player 1's decision given player 2's strategy $\beta L + (1-\beta)R$

To understand why the indifference curves are as depicted, recall from the definition of mixed strategy that $\beta L + (1-\beta)R$ means, for any contingent payoff vector (v_L, v_R), that player 1 gets a payoff v_L with probability β, and payoff v_R with probability $1-\beta$. Thus, player 1's expected payoff from any contingent payoff vector (v_L, v_R) is equal to $\beta v_L + (1-\beta)v_R$. In other words, expecting player 2 to play $\beta L + (1-\beta)R$, player 1's utility function is given by

$$W_1(v_L, v_R) = \beta v_L + (1-\beta)v_R$$

for any $(v_L, v_R) \in \mathbb{R}^2$. Recall from an earlier chapter that the above utility function is simply one that represents a perfect substitutes preference. That is why the indifference curves are the dashed lines in Figure 1.

Furthermore, note from the figure that the dashed lines are steeper than DU, whose slope is obviously $-1/2$. Thus we know $\beta/(1-\beta) > 1/2$, which is simplified to $\beta > 1/3$. Namely, Figure 1 depicts a situation where player 1 expects player 2 to play L with a probability larger than 1/3. From the graph, you see that player 1's optimum is the point U there, meaning that he plays the pure strategy U for sure.

What does player 2's decision problem look like if she expects player 1 to play U for sure? It is depicted by Figure 2. There, the segment RL stands for her choice set, and the vertical dashed lines her indifference curves provided that player 1 is playing U with probability one. Given this pure strategy, player 2's expected payoff from any contingent payoff vector (v_U, v_D) is equal to $1 \cdot v_U + 0 \cdot v_D = v_U$, namely, her preference on the contingent payoff vectors corresponds to the utility function

$$W_2(v_U, v_D) = v_U$$

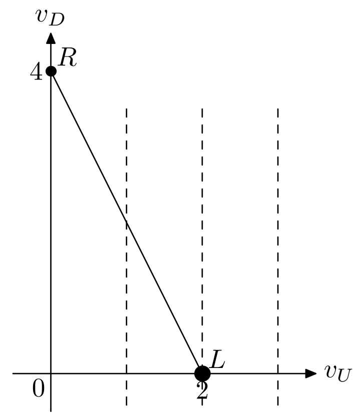

Figure 2: Player 2's decision given player 2's strategy U

for all $(v_U, v_D) \in \mathbb{R}^2$. Hence her indifference curves, the graphs of the equations v_U = constant, are just the vertical lines. It is obvious from the graph that her optimum is L, namely, to play the pure strategy L for sure.

Now that player 2 is playing L for sure, the β in Figure 1 should have been one, and the indifference curves there should not have been slanting but rather be vertical. In other words, the indifference curves depicted in Figure 1 is not a stable situation, as it would predict that player 2 will choose L for sure thereby altering the slopes of the indifference curves. Put differently, at no equilibrium would player 2 play strategy L with a probability $\beta \in (1/3, 1)$.

The purpose of this section is to show you how a Nash equilibrium is related to the decision problems studied in earlier chapters. Each player's equilibrium strategy is like a consumer's optimum, with "consumption bundles" there corresponding to contingent payoff vectors here, except for two new elements: first, his budget set is the convex hull generated by his pure strategies (cf. the zero-sum game chapter); second, his indifference curves (or rather indifference lines) are determined by the strategy that he expects his opponent to play, and for his expectation to be correct, the strategy that he expects his opponent to play should be the opponent's optimum as well.

4 Cournot competition

Let us consider a duopoly game of economic relevance, Cournot competition. Two firms supply the same kind of outputs to a market, where the inverse demand function (cf. Chapter 1) is

$$P(q) := \begin{cases} a - bq & \text{if } 0 \le q \le a/b \\ 0 & \text{if } q \ge a/b \end{cases} \tag{1}$$

for some parameters $a > 0$ and $b > 0$. The two firms choose their output quantities simultaneously. Given firm 1's quantity q_1 and firm 2's quantity q_2, the market price is $P(q_1 + q_2)$ and at that price each firm i sells its output q_i and incurs the cost cq_i, where c is a parameter such that $0 \le c < a$.

Note that this is not a Stackelberg game considered in the previous chapter, and we cannot solve it through backward induction. We shall solve for a pure-strategy Nash equilibrium. To that end, denote (q_1^*, q_2^*) for a Nash equilibrium. By definition of the equilibrium, firm 1 in choosing q_1^*

must be maximizing its profit among all output quantities q_1 provided that firm 2 sticks to q_2^*. In other words, q_1^* solves the decision problem

$$\max_{q_1 \in \mathbb{R}_+} (a - b(q_1 + q_2^*) - c)\, q_1 \tag{2}$$

This is similar to the monopoly problem we have seen in Chapter 1, with q_2^* here treated as a parameter to the decision maker (firm 1). Following the method there, we solve for q_1^* through the first-order condition for any $q_1^* > 0$. We shall verify at the end that $q_1^* > 0$ is a restriction that renders no loss of generality and that, given the objective defined above, the second-order condition is always satisfied by any solution of the first-order condition.

Taking the derivative of the objective with respect to q_1, plugging into the derivative $q_1 = q_1^*$ and setting the derivative obtained thereof to zero, we obtain the first-order condition

$$a - bq_1^* - bq_2^* - c - bq_1^* = 0,$$

i.e.,

$$q_1^* = \frac{a-c}{2b} - \frac{q_2^*}{2}. \tag{3}$$

By the same token, thinking from the perspective of firm 2 (switching the roles between subscripts 1 and 2), we have

$$q_2^* = \frac{a-c}{2b} - \frac{q_1^*}{2}. \tag{4}$$

Plug this equation to the previous one to obtain

$$q_1^* = \frac{a-c}{2b} - \frac{1}{2}\left(\frac{a-c}{2b} - \frac{q_1^*}{2}\right) = \frac{a-c}{4b} + \frac{q_1^*}{4}.$$

Solve this equation for q_1^* to get

$$q_1^* = \frac{a-c}{3b}.$$

Plugging this into the previous equation for q_2^*, we have

$$q_2^* = \frac{a-c}{2b} - \frac{1}{2}\frac{a-c}{3b} = \frac{a-c}{2b}\left(1 - \frac{1}{3}\right) = \frac{a-c}{3b}.$$

Thus, in the Nash equilibrium, each firm chooses the output quantity $(a-c)/(3b)$.

While we did not check the second-order condition for q_1^* and q_2^* in the above calculation, the condition is automatically satisfied because the objective is concave in the sense that its derivative with respective to the choice variable is decreasing when you enlarge the choice variable: You can check that the second-derivative of the objective with respect to q_1 is equal to $-2b$.

Finally, we explain why $q_1^* > 0$ and $q_2^* > 0$ are true in any Nash equilibrium. First, for each firm say firm 1, $q_1^* < (a-c)/b$. Otherwise, $q_1^* \geq (a-c)/b$; by Eq. (1), the market price is less than or equal to

$$a - b\left(\frac{a-c}{b} + 0\right) = a - (a-c) = c,$$

and strictly so if $q_2 > 0$. Consequently, the per-unit profit is nonpositive to the other firm, firm 2, and negative if $q_2 > 0$. Hence firm 2's best response is $q_2^* = 0$. But that in turns implies that firm 1's best response cannot be $q_1^* \geq (a-c)/b$: doing so gives it nonpositive profit, while any $q_1 \in (0, a/b)$ yields a positive profit, with $q_2^* = 0$. Second, it follows from the first observation that $q_2^* > 0$: By the first observation that $q_1^* < (a-c)/b$, Eq. (1) implies that the market price is larger than c if $q_2 = 0$ and for any positive q_2 that is sufficiently close to zero. Thus, firm 2 can always have a positive profit with some $q_2 > 0$, while $q_2 = 0$ yields only zero profit. Hence $q_2^* > 0$. Repeating the above two-step reasoning with the roles between firms 1 and 2 switched, we also get $q_1^* > 0$. Thus, there is no loss of generality to assume that $q_1^* > 0$ and $q_2^* > 0$ in any Nash equilibrium (q_1^*, q_2^*) of our Cournot game.

5 Intersection of the best responses

In the previous section, we derive Eq. (3) for any possible value of q_2^*. (Although we put a star on its shoulder to signify that it is an equilibrium quantity that we are after, nothing about q_2^* being an equilibrium quantity is used in the derivation of (3).) Thus, by the same derivation of (3), we obtain firm 1's best response $q_1^*(q_2)$ to any quantity q_2 that it expects firm 2 to choose:

$$q_1^*(q_2) = \frac{a-c}{2b} - \frac{q_2}{2}.$$

Likewise, the same derivation of (4) gives us firm 2's best response $q_2^*(q_1)$ to any quantity q_1 that it expects firm 1 to choose:

$$q_2^*(q_1) = \frac{a-c}{2b} - \frac{q_1}{2}.$$

The two best responses are graphed as the two solid lines in Figure 3. There, the intersection of the two best responses, point E, is exactly the Nash equilibrium. One way to understand why the intersection is the equilibrium is to recall the fact that we solved the equilibrium (q_1^*, q_2^*) in the previous section by plugging (4) into (3), which is exactly what one does in finding the intersection between two graphs.

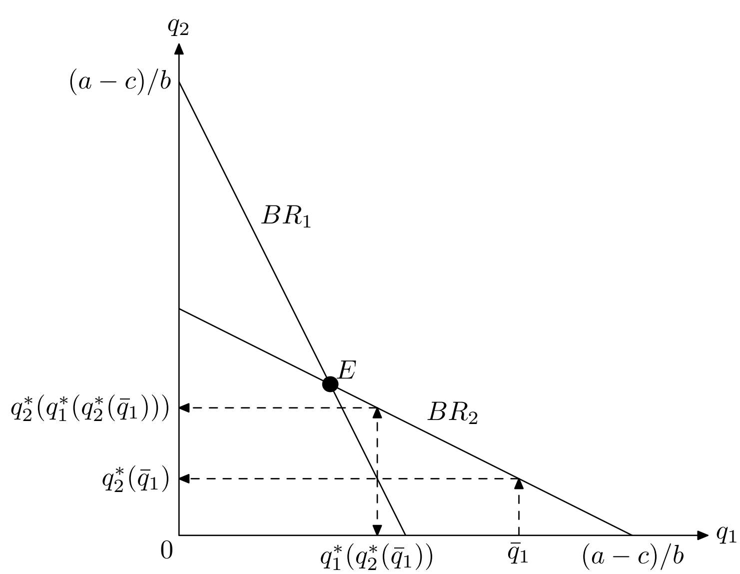

Figure 3: E: Nash equilibrium; BR_i: Firm i's best response to the other firm

There is another, more insightful way to understand why the intersection point E in Figure 3 is the equilibrium: Pick any quantity $\bar{q}_1$ that is different from the horizontal coordinate of the point E. Draw the vertical line whose horizontal coordinate is $\bar{q}_1$ and let the line intersect with the line BR_2 in the figure. The vertical coordinate of that intersection, labeled $q_2^*(\bar{q}_1)$ in the figure, is firm 2's best response to $\bar{q}_1$ (why?). Then draw the horizontal line whose vertical coordinate is $q_2^*(\bar{q}_1)$ and find its intersection with the graph of BR_1 in the figure. The horizontal coordinate of this intersection point, labeled $q_1^*(q_2^*(\bar{q}_1))$ there, is firm 1's best response to firm 2's best response to firm 1's previously expected quantity q_1 (again, why?). Note that $q_1^*(q_2^*(\bar{q}_1)) \neq \bar{q}_1$. Namely, a prediction that firm 1 would choose $\bar{q}_1$ is self-defeating: Should firm 1 be expected to do that, firm 2 would react in such a manner that firm 1 would choose a different quantity. This illustrates that if a point (q_1, q_2) is not a Nash equilibrium, the players' best responses would move away from that point.

By contrast, start with the intersection E between the two best response graphs and apply the same procedure as above, following the best response to the horizontal coordinate of E and then the best response to that best response, so on and so forth. You see that this procedure leaves the coordinates of E completely unchanged. That E is unmoved by the procedure is simply due to the fact that E is an intersection between the best responses. And being an intersection of the best responses is precisely the definition of Nash equilibrium: each player's strategy best responds to the other player's strategy. Thus, we have come across an insight that will be important when you move onto advance economics: *Nash equilibrium means a fixed point of best responses*.

6 Exercises

1. Solving all the Nash equilibria (pure strategy and totally mixed strategy) of the following "Chicken Game":

	Dove	Hawk
Dove	3, 3	1, 4
Hawk	4, 1	0, 0

2. Consider the following game:

	L	m	R
U	1, 2	−2, 1	0, 0
M	−2, 1	1, 2	0, 0
D	0, 0	0, 0	1, 1

   a. Find all the pure-strategy Nash equilibria

   b. Is there any Nash equilibrium where player 1 totally mixes U and M and does not use D, and player 2 totally mixes L and m and does not use R? If Yes, calculate the equilibrium explicitly (i.e., what is the probability for each pure strategy?). If No, explain why.

3. Consider a coordination game described below:

Each player, independently, writes on a piece of paper his name and a number from $\{0, 1, 2\}$ as his *bid*; then the papers are collected and announced. A player's payoff is equal to twice the minimum of all the bids minus his bid. For example, if three players bid 1, 1, and 2 respectively, then the first player's payoff is equal to $2 \times \min\{1, 1, 2\} - 1 = 1$, and the third player's payoff equal to $2 \times 1 - 2 = 0$; but if the second player bids 0, then the first player's payoff is equal to -1 and the third player's payoff equal to -2, while the second's equal to 0.

Suppose that there are only two players.

a. Write down the payoff matrix of this two-player game. (Label a player's pure strategies by 0, 1 and 2, corresponding to the bids.)

b. Find all the pure-strategy Nash equilibrium.

c. Is there any Nash equilibrium where each player totally mixes all three pure strategies? If Yes, calculate the equilibrium explicitly (i.e., what is the probability for each pure strategy?). If No, explain why.

d. Is there any Nash equilibrium where each player totally mixes between 0 and 2 and does not use 1 at all? If Yes, calculate it explicitly. If No, explain why.

4. Suppose $a = 10$, $b = 2$ and $c = 0$ in the Counot games. Work out the Nash equilibrium through mimicking the steps in the chapter. (The general formula at the end of the chapter of course applies, but it is more instructional to work out the solution yourself instead of simply plugging the numbers into the general formula.)

5. Consider a Cournot game with the same general parameters as in the chapter except that there are n firms, with n a parameter that can be any integer larger than or equal to two. Thus, if firm i chooses output quantity q_i $(i = 1, \dots, n)$, then the total output quantity is $q_1 + \cdots + q_n$ and the market price for each firm is equal to

$$P(q_1 + \cdots + q_n) = a - b(q_1 + \cdots + q_n).$$

a. Pick any firm i. Write down the first-order condition for firm i's output quantity q_i^* to best respond to those chosen by other firms, i.e., $(q_1^*, \dots, q_{i-1}^*, q_{i+1}^*, \dots, q_n^*)$, which in game theory is denoted by q_{-i}^*.

b. Restrict attention to *symmetric Nash equilibrium*, symmetric in the sense that $q_1^* = q_2^* = \cdots = q_n^* = q^*$ for some common q^*. Plug this equation into the first-order condition and solve for q^*. Then:

   i. What is the market price equal to in this equilibrium?

   ii. When $n \to \infty$:

      A. what is the limit of the market price equal to?

      B. what is the limit of each firm's profit equal to?

      C. what is the limit of the total profit—across all the n firms—equal to?

6. To vaccinate or not to vaccinate one's children against the flu? This debate, with the moralistic and conspiracy-theory rhetorics on both sides peeled off, can be reduced to the following game-theoretic situation: Say there are only two families in a community. If both families vaccinate their children, both families gain from the protection of the vaccine, but both also suffer a little bit due to the (slight) risk of vaccination. If neither families vaccinate their children, both lose a lot due to the likelihood of infection. If one family vaccinates its children while the other family does not, both families gain the protection of the vaccine, with the vaccinated one protected by the vaccine and the one that opts out enjoys having a neighbor who is not infected; meanwhile, the family that opts out of vaccination also gains from being free of the vaccination risk, while the other family loses that little bit due to vaccination risk.

   a. Find a game presented in this chapter, in the lecture or in the other exercises, that captures the above-described game-theoretic situation. Why is it impossible for every family to opt for vaccination?

   b. Consider the totally mixed strategy Nash equilibrium of the game located in the previous step. Suppose that everyone in a population plays that equilibrium. What fraction among the population opts out of vaccination?

7. Suppose that the monetary payoffs in the Battle of the Sexes (BOTS) are:

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

   Now add the following twist: Before the BOTS is played, player A (the row player) has the option of giving 2 dollars to player B (the column player) for free. In deciding whether to give B the 2 dollars, player A simultaneously chooses his action in the BOTS. Player B chooses her action in the BOTS after she has observed player A's dollar-giving decision and before she sees A's choice in the BOTS.

   a. Write down the normal form of this enlarged game (cf. Chapter 8).

   b. Solve the normal form by iterated elimination of weakly dominated strategies.

   c. Contrast the solution obtained from the previous step to the set of Nash equilibria of the original BOTS that has no such money-giving twist. What is the moral of this fable?