Nash Equilibrium:
Definition

Adam Brandenburger*

Version 01/04/07

1 Introduction

In the games in "Dominance and Iterated Dominance," the iterated-dominance procedure yields a unique choice for each player. In many games, the procedure narrows down the players' choices much less. Consider, for example, the game depicted in Figure 1 below.

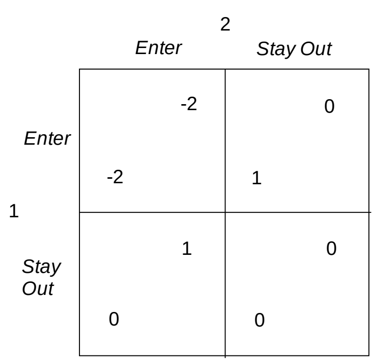

Figure 1

This is a Battle of the Sexes, and can be thought of as a game between two firms, labeled 1 and 2, each of which must decide whether or not to enter a certain business. If one firm enters and the other stays out, the entrant makes money. However, if both firms enter, then both lose money. There is only room for one firm, so to speak.

Neither the choice *Stay Out* nor the choice *Enter* is dominated in this game, so application of the iterated-dominance procedure allows either choice to be

---

*With the assistance of Amanda Friedenberg and Konrad Grabiszewski. nash-01-04-07

made. This seems like an accurate reflection of the uncertainty inherent in this game. There is no reason to rule out any of the four possible outcomes.

Let's look a little more closely at two of these outcomes–that in which both players choose *Stay Out*, and that in which both players choose *Enter*. There is, in a certain sense, an 'inconsistency' in each of these situations. Take, for example, the case in which both players choose *Enter*. Choosing *Enter* is optimal for player 1 only if 1 assigns probability of at least $\frac{2}{3}$ to player 2's choosing *Stay Out*. In fact, player 2 chooses *Enter*. Player 1's choice is based on an incorrect assessment of player 2's choice. Of course, this 'inconsistency' may be nothing more than an accurate reflection of the uncertainty facing player 1. Still, game theory includes a concept, namely that of Nash equilibrium, that rules out such situations.

2 Pure-Strategy Equilibrium

We consider an n-player strategic-form game $\Gamma = \left\langle S^1, ..., S^n, \pi^1, ..., \pi^n \right\rangle$, where S^i is player i's finite set of pure strategies and π^i is player i's payoff function.

Definition 1 *A* ***(pure-strategy) Nash equilibrium*** *specifies a strategy $s^i \in S^i$ for each player i (where $i = 1, \ldots, n$) such that*

$$\pi^i(s^1, \ldots, s^{i-1}, s^i, s^{i+1}, \ldots, s^n) \geq \pi^i(s^1, \ldots, s^{i-1}, r^i, s^{i+1}, \ldots, s^n)$$

for every strategy $r^i \in S^i$.

In words, a Nash equilibrium specifies a strategy for each player, in such a way that each player's strategy yields the player at least as high a payoff as any other strategy of the player, given the strategies of the other players.

Let's calculate the Nash equilibria of the game depicted in Figure 1. There are two: In one equilibrium, player 1 chooses *Enter* and player 2 chooses *Stay Out*. In the other equilibrium, player 1 chooses *Stay Out* and player 2 chooses *Enter*. Note, in particular, that neither the case where both players choose *Stay Out* nor the case where both players choose *Enter* is an equilibrium.

The following easy theorem gives one way to think about the equilibrium concept. (For a formal statement, see the Preliminary Observation in "Epistemic Conditions for Nash Equilibrium," by Robert Aumann and Adam Brandenburger, *Econometrica*, 1995, 63, 1161-1180.)

Theorem 1 *(Informal statement) Suppose that each player is rational and assigns probability one to the strategies chosen by the other players. Then, the strategies constitute a Nash equilibrium.*

While obvious, this theorem should be contrasted with the assertion that if the players in a game are rational, then they must play a Nash equilibrium. This assertion, which one often hears or reads, is false. At least, it is false given the way we are using the concept of rationality in this course. For us, the assumption that the players are rational implies that they choose any strategies that remain after Step 1 of the iterated-dominance procedure. That is, the players can choose any undominated strategies. (Refer back once more to the note "Dominance and Iterated Dominance.") In general, undominated strategies are not equilibrium strategies. The game in Figure 1 is a case in point.

Indeed, the above theorem suggests that one way to decide whether the equilibrium assumption is a good one in a particular situation is to think about how plausible it is that the players will correctly predict one another's choices.

For example, in the game depicted in Figure 1, there seems to be no good reason to require that this degree of consistency should arise. The game is naturally thought of as being played just once, between two players who may well have no history of prior interaction. It truly seems like a game played in a 'fog,' in which wrong guesses about another player's choices are as much to be expected as right guesses. However, there are other contexts in which equilibrium analysis is appropriate. One might be a situation in which the players meet and play the same game over and over again. Here, there could be some kind of dynamical process, in which the players adapt their behavior to one another until they have no more reason to change what they are doing. At this point, the players' choices might well constitute an equilibrium.

3 Mixed-Strategy Equilibrium

Not every game possesses a Nash equilibrium as defined above. To see this, consider the game depicted in Figure 2 below. It can be checked that none of the four possibilities–(U, L), (D, L), (U, R), or (D, R)–is a Nash equilibrium.

So far, we have assumed that each player makes a definite choice of strategy from his set of strategies, i.e. chooses a so-called *pure* strategy. Let us now allow the players also to choose strategies at random, i.e. to choose *mixed* strategies as defined in "Dominance and Iterated Dominance." For example, player 1 might choose *Up* with probability $\frac{1}{2}$ and *Down* with probability $\frac{1}{2}$, while player 2 might choose *Left* with probability $\frac{1}{4}$ and *Right* with probability $\frac{3}{4}$. In this case, the expected payoff to player 1 is

$$\frac{1}{2} \times \left[(\frac{1}{4} \times 2) + (\frac{3}{4} \times 0) \right] + \frac{1}{2} \times \left[(\frac{1}{4} \times 0) + (\frac{3}{4} \times 1) \right] = \frac{5}{8}.$$

Likewise, the expected payoff to player 2 is

$$\frac{1}{4} \times \left[(\frac{1}{2} \times 0) + (\frac{1}{2} \times 1) \right] + \frac{3}{4} \times \left[(\frac{1}{2} \times 1) + (\frac{1}{2} \times 0) \right] = \frac{1}{2}.$$

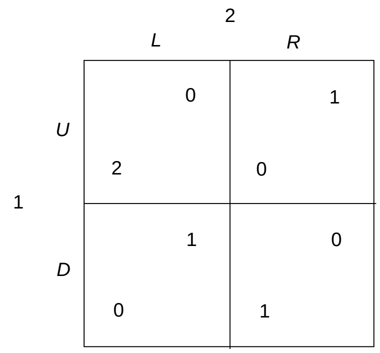

Figure 2

In general, if player 1 chooses the mixed strategy $\sigma^1 \in \Delta(S^1)$, player 2 chooses the mixed strategy $\sigma^2 \in \Delta(S^2)$, ..., player n chooses the mixed strategy $\sigma^n \in \Delta(S^n)$, we will denote the resulting expected payoff to player i by $v^i(\sigma^1, \sigma^2, \ldots, \sigma^n)$. Formally,

$$v^i\left(\sigma^1, \ldots, \sigma^n\right) = \sum_{(s^1,\ldots,s^n)\in S^1\times\cdots\times S^n} \sigma^1(s^1) \times \cdots \times \sigma^n(s^n)\pi^i(s^1, \ldots, s^n),$$

where we follow the notation in "Dominance and Iterated Dominance."

Definition 2 *A (**mixed-strategy) Nash equilibrium** specifies a mixed strategy $\sigma^i \in \Delta(S^i)$ for each player i (where $i = 1, \ldots, n$) such that*

$$v^i(\sigma^1, \ldots, \sigma^{i-1}, \sigma^i, \sigma^{i+1}, \ldots, \sigma^n) \geq v^i(\sigma^1, \ldots, \sigma^{i-1}, \rho^i, \sigma^{i+1}, \ldots, \sigma^n)$$

for every mixed strategy $\rho^i \in \Delta(S^i)$.

In words, a mixed-strategy Nash equilibrium specifies a mixed strategy for each player, in such a way that each player's mixed strategy yields the player at least as high an expected payoff as any other mixed strategy of the player, given the mixed strategies of the other players.

Let's look for a mixed-strategy Nash equilibrium of the game depicted in Figure 2. In fact, let's consider the pair of mixed strategies we suggested above. The mixed strategy of player 2 is optimal for player 2, given the mixed strategy employed by player 1. To see this, note that if player 1 chooses *Up* with probability $\frac{1}{2}$ and *Down* with probability $\frac{1}{2}$, then either *Left* or *Right* yields player 2 an expected payoff of $\frac{1}{2}$. A little thought then shows that player 2 will still have an expected payoff of $\frac{1}{2}$ if she chooses between Left and Right at random–whether with the given probabilities of $\frac{1}{4}$ and $\frac{3}{4}$, or with any other probabilities. Now turn to player 1. If player 2 chooses *Left* with probability $\frac{1}{4}$ and *Right* with probability $\frac{3}{4}$, then *Up* yields player 1 an expected payoff of $\frac{1}{2}$, while *Down* yields player 1 an expected payoff of $\frac{3}{4}$. It follows that player 1 could increase his expected payoff from $\frac{5}{8}$ to $\frac{3}{4}$, by choosing *Down* with probability 1 rather than choosing *Up* with probability $\frac{1}{2}$ and *Down* with probability $\frac{1}{2}$. So, these mixed strategies do not constitute a mixed-strategy equilibrium.

Next, let player 1 continue to choose *Up* with probability $\frac{1}{2}$ and *Down* with probability $\frac{1}{2}$, but suppose now that player 2 chooses *Left* with probability $\frac{1}{3}$ and *Right* with probability $\frac{2}{3}$. We know that player 2's mixed strategy will still be optimal for her. And now, either *Up* or *Down* yields player 1 an expected payoff of $\frac{2}{3}$, so that it is also optimal for player 1 to choose between *Up* and *Down* with the given probabilities. Evidently, we have found a mixed-strategy equilibrium of the game. This is a particular instance of the following famous result ("Non-Cooperative Games," by John Nash, *Annals of Mathematics*, 1951, 54, 286-295).

Theorem 2 *(Nash) Every n-person game in which each player has finitely many strategies possesses at least one mixed-strategy Nash equilibrium.*

We give a proof of this result in the note "Nash Equilibrium: Existence."

As far as calculating mixed-strategy Nash equilibria is concerned, the following proposition is helpful. It generalizes the technique we used above in analyzing the game in Figure 2. (The proof is left as an exercise.)

Proposition 1 *A collection $(\sigma^1, \ldots, \sigma^n)$ of mixed strategies is a Nash equilibrium if and only if for each player $i = 1, \ldots, n$, it is true that $\sigma^i(s^i) > 0$ implies*

$$v^i(\sigma^1, \ldots, \sigma^{i-1}, s^i, \sigma^{i+1}, \ldots, \sigma^n) \geq v^i(\sigma^1, \ldots, \sigma^{i-1}, r^i, \sigma^{i+1}, \ldots, \sigma^n)$$

for every $r^i \in S^i$.

In words, the proposition says that the equilibrium condition is that for each player i, the mixed strategy σ^i can give positive weight only to pure strategies s^i that maximize player i's expected payoff, given the mixed strategies used by the other players.

4 Relationship to Iterated Dominance

Example 1 *In the game depicted in Figure 3 below,*[1] *all strategies are iteratively undominated. But only the pair of strategies (M, C) constitutes a (pure or mixed) Nash equilibrium. (Verify this assertion.)*

[1] Similar to an example in "Rationalizable Strategic Behavior," by B. Douglas Bernheim, *Econometrica*, 52, 1984, 1007-1028.

			2	
		L	*C*	*R*
	T	7, 0	0, 5	0, 7
1	*M*	5, 0	2, 2	5, 0
	B	0, 7	0, 5	7, 0

Figure 3

This example is an illustration of the following general relationship between the iteratively undominated strategies and the Nash-equilibrium strategies. (We leave the proof of this result as an exercise.)

Theorem 3 *Fix an n-player strategic-form game $\Gamma = \left\langle S^1, ..., S^n, \pi^1, ..., \pi^n \right\rangle$ and an associated Nash equilibrium $(\sigma^1, \ldots, \sigma^n)$. For each player i, if $\sigma^i(s^i) > 0$ (i.e. the strategy s^i is played with positive probability), then the strategy s^i is iteratively undominated.*

The converse is certainly false, as Example 1 shows. That is, it is entirely possible that an iteratively undominated strategy does not 'appear' in any Nash equilibrium.