

# EQUILIBRIUM IN NON-COOPERATIVE GAMES

JOEL SHAPIRO

ABSTRACT. These notes provide a brief introduction to the theory of non-cooperative games, with emphasis on the notion of Nash Equilibrium, and the proof of its existence by means of the Brouwer Fixed Point Theorem.

In this section we'll formally define the notion of a "non-cooperative game," in the setting of which we'll discuss John F. Nash's fundamental notion of "equilibrium." Then, following Nash, we'll use the Brouwer Fixed-Point Theorem to prove the famous result that won him the 1994 Nobel Prize in Economics: *In a certain natural sense such an equilibrium exists for every "finite" such game.*

The following example illustrates much of what is to follow.

**4.1. Mathematicians go out to dinner.** Twenty mathematicians go out to dinner, agreeing to split the bill equally. Upon arriving at the restaurant they discover that there are only two choices: the \$10 dinner and the \$20 one. Nobody wants to spend \$20. Assuming that each person acts only on the basis of "rational self-interest," what will they do?

Each player reasons: "Everyone else is going to choose the \$10 dinner, so if I choose the \$20 one, I get it for just \$10.50." But everyone else has the same thought, so it seems most likely that—even though nobody wants to—everyone will choose the \$20 dinner. Furthermore, if each player reasons this far ahead, then there's no sense in anyone opting unilaterally for the \$10 dinner, which would then cost that person \$19.50.

This is an example of Nash equilibrium. Later we'll formalize the idea, but for now let's consider a few more examples—all of which involve just two competitors.

**4.2. Four Examples.** Here are four simple two-person games. For each one, the players know all the strategies and all the payoffs, each player's goal being: "Maximize my payoff." What each competitor doesn't know is—at each play of the game—the strategy that the other will use.

**4.2.1. Matching Pennies.** Each player puts down a penny. If both coins show the same face (heads-heads or tails-tails), Player I wins the other player's coin. If the faces differ (heads-tails or tails-heads), then Player II wins Player I's coin. For this game, and each example to follow, it's convenient to represent the situation by a "payoff matrix" whose entries represent the payoffs to each competitor for each possible

---

*Date:* October 16, 2013.

play of the game. For Matching Pennies the payoff matrix for Player I is:

$$\begin{array}{c|cc} \text{I / II} & H & T \\ \hline H & 1 & -1 \\ \hline T & -1 & 1 \end{array}$$

Here the “ $(i, j)$ –entry” of the matrix represents Player I’s payoff upon playing face  $i$  ( $H$  or  $T$ ) to Player II’s face  $j$ . For this game the payoff matrix for Player II is the negative of the one for Player I, i.e., Matching Pennies is a “zero-sum” game—each player’s gain is the other’s loss. In a two-person game it’s often convenient to display the payoffs for both players in a “bi-matrix.” For Matching Pennies this is:

$$\begin{array}{c|cc} \text{I / II} & H & T \\ \hline H & (1, -1) & (-1, 1) \\ \hline T & (-1, 1) & (1, -1) \end{array}$$

where now the entry at row  $i$  and column  $j$  displays, for the situation where Player I plays  $i$  and Player II plays  $j$ , Player I’s payoff in the first coordinate, and Player II’s in the second.

4.2.2. *Rock–Paper–Scissors*. Here each competitor has three strategies: “rock”, “paper”, or “scissors”, which they play simultaneously with the hand signals we all know. The rules are that “paper covers rock,” “scissors cuts paper,” and “rock breaks scissors.” Let’s suppose that at each play of the game, the winner receives one penny from the loser, with neither winning anything if both opt for the same strategy. The situation is captured by the  $3 \times 3$  bi-matrix shown below where, as before, each row represents payoffs to both players for the corresponding strategy of Player I, and each column represents payoffs for the corresponding strategy of Player II.

$$\begin{array}{c|ccc} \text{I \ II} & r & p & s \\ \hline r & (0, 0) & (-1, 1) & (1, -1) \\ \hline p & (1, -1) & (0, 0) & (-1, 1) \\ \hline s & (-1, 1) & (1, -1) & (0, 0) \end{array}$$

4.2.3. *Prisoners’ Dilemma*. This game imagines both players to be prisoners, held in separate interrogation cells. The police are sure the prisoners have committed a serious crime which, if this can be proven, will land each suspect in prison for five years. However there’s only enough evidence to convict the pair of a less serious infraction—one that carries a prison term of only one year. The prisoners are offered this deal: If one of them defects by implicating the other in the more serious crime while the other prisoner does not, then the defector will go free. If each prisoner defects, then each will receive, for their cooperation, a sentence of only three years. The payoff matrix below summarizes the situation, where “L” denotes “stay loyal” and “D” denotes “defect.”

$$\begin{array}{c|cc}
 I \backslash II & L & D \\
 \hline
 L & (-1, -1) & (-5, 0) \\
 D & (0, -5) & (-3, -3)
 \end{array}$$

4.2.4. *Battle of the Sexes.* In this game Players I and II are a couple who wish to spend an evening out. Player I wants to go to the ball game ( $B$ ) while Player II prefers the symphony ( $S$ ). They work on opposite sides of town, and plan to decide that afternoon, via some form of electronic communication, which event to attend. But because of a massive solar flare they can't do this, so each must guess what the other's going to do. Their preferences might be indicated by the payoff matrix below, where the diagonal terms indicate that for each player the favored event is twice as desirable as the alternative, and the cross-diagonal terms show that no choice is so desirable as to warrant going alone.

$$\begin{array}{c|cc}
 I \backslash II & B & S \\
 \hline
 B & (2, 1) & (0, 0) \\
 S & (0, 0) & (1, 2)
 \end{array}$$

4.3. **Nash Equilibrium.** In Prisoners' Dilemma it appears at first glance that both players will be better off if they remain loyal, thus insuring that, although nobody goes free, each gets a light sentence. However if one player believes in the other's unswerving loyalty, then that player will do better by defecting. A quick look at the rows of the payoff matrix shows, for example, that Player I's strategy D *dominates* L in the sense that no matter which strategy Player II chooses, Player I's payoff for D is better than for L. Similarly, looking at the columns of the matrix, we see that Player II's strategy D also dominates L. Thus each player has a dominant strategy, and it appears that—even though the best overall result would be achieved were both to stay loyal—it would be “safest” for each of them to defect.

Such “dominance” is not always present; neither of the other three games described above exhibits it. However in Battle of the Sexes the strategy pairs  $(B, B)$  and  $(S, S)$  share a salient feature of dominant pairs: If a player deviates unilaterally from either pair then that player's payoff decreases. The same is true of in the “Mathematicians go to Dinner” game for the strategy twenty-tuple: “Everyone chooses the \$20 dinner.”

In the middle of the last century John F. Nash made a profound study of this “weaker-than-dominance” notion. Now called “Nash equilibrium,” it has become a cornerstone of modern economic theory. To properly describe it we need to formalize our notion of “non-cooperative game.”

**Definition.** A *non-cooperative game* (henceforth, just a “game”) consists of a set  $\mathcal{P}$  of *players*, where each player  $P$  has:

- (a) a *strategy set*  $\Sigma_P$ , and
- (b) a real-valued *payoff function*  $u_P$  defined on the cartesian product  $\Sigma = \prod_{P \in \mathcal{P}} \Sigma_P$  of strategy sets.

Each element  $\sigma \in \Sigma$  represents a particular play of the game; think of it as a “vector” with coordinates indexed by the players,  $\sigma(P)$  being the strategy chosen by Player  $P$ . Then  $u_P(\sigma)$  is player  $P$ 's payoff

for that particular play of the game. Upon defining  $\mathcal{U} = \{u_P : P \in \mathcal{P}\}$ , we can refer to the game described above as simply the triple  $(\mathcal{P}, \Sigma, \mathcal{U})$ .

In the examples of §4.2 the set  $\mathcal{P}$  of players has just two elements, and the strategy sets for each player are finite (and, at least in these cases, the same for each player). For example, in Rock-Paper-Scissors,  $\mathcal{P} = \{\text{Player I}, \text{Player II}\}$ , the strategy sets for each player are  $\Sigma_I = \Sigma_{II} = \{r, p, s\}$ , and the collection  $\Sigma$  of strategy “vectors” is the nine-element cartesian product  $\Sigma_I \times \Sigma_{II} = \{(x, y) : x, y \in \{r, p, s\}\}$ . The payoff function for Player I is represented by the matrix below;  $u_I(x, y)$  is the number located at the intersection of the  $x$ -row and the  $y$ -column

$$I \backslash II \begin{array}{c} r \quad p \quad s \\ \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix} \end{array}$$

while that of Player II is represented by the negative of this matrix (so, like Matching Pennies, this is a “zero-sum” game). Thus, for example,  $u_I(p, s) = -1 = -u_{II}(p, s)$ .

For the “Mathematicians go to dinner” game our set of players is  $\mathcal{P} = \{P_1, P_2, \dots, P_{20}\}$ , and each has the same strategy set  $\Sigma_j = \{10, 20\}$  ( $1 \leq j \leq 20$ ). Thus the set of strategy vectors  $\Sigma = \prod_{j=1}^{20} \Sigma_j$  is the set of 20-tuples, each coordinate of which is either 10 or 20. The description of the game suggests that Player  $P_j$ 's payoff for strategy vector  $x = (x_1, x_2, \dots, x_{20})$  should be

$$u_j(x) = x_j - \frac{1}{20} \sum_{k=1}^{20} x_k \quad (j = 1, 2, \dots, 20),$$

where on the right-hand side the first term is the value of  $P_j$ 's dinner, and the second is the amount  $P_j$  pays, i.e., the total cost of all twenty dinners, averaged over the participants.

**Definition.** In a non-cooperative game  $(\mathcal{P}, \Sigma, \mathcal{U})$ : if  $P \in \mathcal{P}$  and  $\sigma$  and  $\sigma'$  are strategy vectors that differ only in the  $P$ -coordinate (i.e.,  $\sigma(Q) = \sigma'(Q)$  for all  $Q \in \mathcal{P} \setminus \{P\}$  and  $\sigma'(P) \neq \sigma(P)$ ), then we'll say  $\sigma'$  is a *unilateral change of strategy* for player  $P$ .

With these preliminaries in hand we can formulate the fundamental concept of this section.

**Definition.** For a non-cooperative game  $(\mathcal{P}, \Sigma, \mathcal{U})$ : a *Nash equilibrium* is a strategy vector with the property that: For each player, no unilateral change of strategy results in a strictly better payoff.

More precisely: “ $\sigma^* \in \Sigma$  is a Nash equilibrium” means:

$$\text{If } P \in \mathcal{P} \text{ and } \sigma \in \Sigma, \text{ with } \sigma(Q) = \sigma^*(Q) \text{ for every } Q \in \mathcal{P} \setminus \{P\}, \text{ then } u_P(\sigma) \leq u_P(\sigma^*).$$

Nash equilibrium is perhaps best interpreted in terms of “best response.” To avoid notational complications, let's consider this notion just for two-person games. In this setting Player I and Player II have strategy sets  $X$  and  $Y$  respectively, and  $\Sigma = X \times Y$ , i.e. the set of ordered pairs  $(x, y)$ , where  $x$  ranges through  $X$  and  $y$  through  $Y$ .

**Definition.** To say that Player I's strategy  $x^* \in X$  is a *best response* to Player II's  $y \in Y$  means that

$$u_I(x^*, y) = \max_{x \in X} u_I(x, y)$$

Similar language defines Player II's best response to a given strategy of Player I. With this terminology, for a two-person game a Nash equilibrium is a pair of strategies, each of which is a best response to the other.

In terms of bi-matrices, it's easy to check if such mutual best-response strategy pairs exist, and if so, to identify them: For each column of the matrix, mark off (say by an underline) the entry or entries for which Player I's payoff is largest (the "best response(s)" of Player I to Player II's strategy for that column). Similarly, for each row, mark off (say by an overline) each entry for which Player II's best response is largest. Any entry that has two marks is a Nash equilibrium. It's easy to check this way that  $(D, D)$  is a Nash equilibrium for Prisoners' Dilemma, that Rock-Paper-Scissors and Matching Pennies have no Nash equilibrium, and that Battle of the Sexes has *two* Nash equilibria:  $(B, B)$  and  $(S, S)$ .

It must be emphasized that, as the examples above show:

- Not every game has a Nash equilibrium (e.g., Matching Pennies, Rock-Paper-Scissors).
- If it exists, Nash equilibrium need not be unique (e.g., Battle of the Sexes).
- If it exists, Nash equilibrium need not provide the best outcome for each—or for any—player (e.g., Mathematicians go to Dinner, Prisoners' Dilemma).

**4.4. Mixed Strategies.** Let's return to Matching Pennies. In a sequence of consecutive plays, how can each competitor guard against falling into a pattern discoverable by the other? Each could choose heads or tails *randomly*, for example tossing the coin rather than just putting it down. They could even use biased coins; say Player I chooses a coin with probability  $p \in [0, 1]$  of coming up Heads (and  $1 - p$  of Tails), while Player II opts for one that has probability  $q \in [0, 1]$  of Heads. So now we have the makings of a new game, with the same players, but with ("mixed") strategies that are *probability distributions* over the original ("pure") strategies. Let's represent the mixed strategy "Heads with probability  $p$  and Tails with probability  $1 - p$ " by the vector  $(p, 1 - p) = pH + (1 - p)T$  where now  $H = (1, 0)$  denotes the "pure" strategy of choosing Heads with probability 1, and  $T = (0, 1)$  the pure strategy of choosing Tails with probability 1. The strategy set for each player can now be represented geometrically as the line segment in  $\mathbb{R}^2$  joining  $H = (1, 0)$  to  $T = (0, 1)$ .

We need payoff functions for this new game, and for these we choose *expected payoffs*, calculated in the obvious way: Player I's expected payoff for playing the pure strategy Heads to Player II's mixed strategy  $\sigma_q := (q, 1 - q) := qH + (1 - q)T$  is:

$$u_I(H, \sigma_q) := 1 \cdot q + (-1) \cdot (1 - q) = 2q - 1,$$

while for playing Tails to Player II's  $\sigma_q$  it is

$$u_I(T, \sigma_q) := -1 \cdot q + 1 \cdot (1 - q) = 1 - 2q.$$

Thus Player I's expected payoff for playing strategy  $\sigma_p$  to Player II's  $\sigma_q$  is:

$$(1) \quad u_I(\sigma_p, \sigma_q) = p \cdot u_I(H, \sigma_q) + (1 - p) \cdot u_I(T, \sigma_q) = (2p - 1)(2q - 1).$$

Similarly, Player II's expected payoff for playing strategy  $\sigma_q$  to Player I's  $\sigma_p$  is

$$u_{II}(\sigma_p, \sigma_q) = (1 - 2q)(2p - 1) = -u_I(\sigma_p, \sigma_q),$$

thus extending the "pure-strategy" zero-sum relationship between the two players' payoffs to our new situation. These are the payoff functions for our new game, which we might term the "mixed extension" of the original game.

Recall that the original matching-pennies game had no Nash equilibrium. However for the mixed-strategy extension *there is one*, namely the strategy pair  $(\sigma_{1/2}, \sigma_{1/2})$  wherein each player chooses to flip a fair coin. The (expected) payoffs for this strategy pair are not spectacular—namely zero—but neither player's payoff can be unilaterally improved. For example, if Player II chooses to play  $\sigma_{1/2}$  then, according to (1) above, Player I's expected payoff is  $u_I(\sigma_p, \sigma_{1/2}) = 0$  for any  $p \in [0, 1]$ , and similarly for Player II when Player I chooses to play  $\sigma_{1/2}$ . In other words, any strategy  $\sigma_p$  with  $(0 \leq p \leq 1)$  is a *best response* for Player I to Player II's strategy  $\sigma_{1/2}$ , and vice versa.

*Remark.* This "indifference of best response" is not an accident. It arises from the linearity of mixed-strategy payoffs in each variable, and will be important for our later work. In the next section we'll treat this phenomenon more carefully.

*For "Matching Pennies" there is no other Nash equilibrium.* To see why, let's examine the strategy pair  $(\sigma_p, \sigma_q)$  where, say,  $q \neq 1/2$ . Suppose for the moment that Player I is not playing a pure strategy, i.e.  $\sigma_p$  with  $0 < p < 1$ . If  $q < 1/2$ , so Player II is more likely to play Tails, then common sense (and the formula for  $u_I$ ) dictates that Player I's best response is to replace the mixed strategy  $\sigma_p$  by the pure strategy  $T = \sigma_0$ , which results in a strictly larger payoff. Similarly, if  $q > 1/2$  and  $p \neq 1$  then Player I does strictly better by playing Heads, i.e.  $\sigma_1$ . Thus  $(\sigma_p, \sigma_q)$  is not a Nash equilibrium.

If  $p = 0$ , so Player I is already playing the pure strategy Tails, then, since we already know that our game has no pure-strategy equilibrium, we may as well assume Player II is *not* playing a pure strategy. Thus  $0 < q < 1$  and Player II can get an improved payoff by switching to the pure strategy  $H = \sigma_1$ —the best response to Player I's strategy of Tails. Thus  $(\sigma_0, \sigma_q)$  is not a Nash equilibrium; similarly, neither is  $(\sigma_1, \sigma_q)$ .

*In summary:* for any  $p \in [0, 1]$  and any  $q \in [0, 1] \setminus \{1/2\}$  the mixed strategy pair  $(\sigma_p, \sigma_q)$  is not a Nash equilibrium. Similarly, the same is true for any  $q$  and any  $p \neq 1/2$ . Thus  $(\sigma_{1/2}, \sigma_{1/2})$  is the unique Nash equilibrium for the mixed-strategy extension of Matching Pennies.  $\square$

**Exercise.** Show that, for the mixed-strategy extension of Prisoners' Dilemma, the strategy pair  $(D, D)$  is still the only Nash equilibrium.

**4.5. The mixed-strategy extension of Rock-Paper-Scissors.** Recall that for the Rock-Paper-Scissors game the payoff matrix for Player I is:

$$I \backslash II \begin{array}{c} r \quad p \quad s \\ \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix} \end{array}.$$

Let's omit the labels and just write this payoff matrix as

$$A = \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}$$

recalling that the corresponding matrix for Player II is just  $-A$ .

Suppose  $y = (y_1, y_2, y_3)$  is a probability vector (non-negative components that sum to 1), and suppose that, in the mixed-strategy extension of our game, Player II uses the strategy  $y = y_1r + y_2p + y_3s$ . Then the pure-strategy payoffs to Player I are found by matrix multiplication:

$$\begin{bmatrix} u_I(r, y) \\ u_I(p, y) \\ u_I(s, y) \end{bmatrix} = Ay^T = A \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} -y_2 + y_3 \\ y_1 - y_3 \\ -y_1 + y_2 \end{bmatrix}$$

Thus the mixed-strategy payoff to Player I upon playing the strategy  $x = x_1r + x_2p + x_3s$  (where  $x = (x_1, x_2, x_3)$  is a probability vector) is

$$u_I(x, y) = xAy^T = [x_1, x_2, x_3] \begin{bmatrix} -y_2 + y_3 \\ y_1 - y_3 \\ -y_1 + y_2 \end{bmatrix},$$

i.e.,

$$(2) \quad \begin{aligned} u_I(x, y) &= (-y_2 + y_3)x_1 + (y_1 - y_3)x_2 + (-y_1 + y_2)x_3 \\ &= (x_2 - x_3)y_1 + (-x_1 + x_3)y_2 + (x_1 - x_2)y_3 \end{aligned}$$

(here, and henceforth, we identify a mixed strategy with its probability vector). We see from these equations that if Player II employs the "uniform" strategy  $\sigma^* = \frac{1}{3}r + \frac{1}{3}p + \frac{1}{3}s$ , then *any* mixed strategy  $\sigma = x_1r + x_2p + x_3s$  for Player I is a best response in that  $u_I(\sigma, \sigma^*) = 0$ . In other words, a unilateral change of strategy cannot improve Player I's payoff. Since  $u_{II} = -u_I$ , the same is true for Player II if Player I uses the strategy  $\sigma^*$ . Thus the strategy pair  $(\sigma^*, \sigma^*)$  is a Nash equilibrium for the mixed-strategy extension of Rock-Paper-Scissors.

*Rock-Paper-Scissors has no other Nash equilibrium!*

Before proving this we need a few preliminaries that will also help in the proof of Nash's existence theorem .

**Definition.** (a) Let  $\Pi_n$  denote the set of *probability vectors* in  $\mathbb{R}^n$ , i.e. the collection of vectors, each of which has non-negative coordinates that sum to 1. In other words,  $\Pi_n$  is the *convex hull* of the standard unit vectors  $\{e_1, e_2, \dots, e_n\}$  for  $\mathbb{R}^n$ . Clearly it's a compact, convex subset of  $\mathbb{R}^n$ .

(b) The *support* of a vector  $x \in \Pi_n$  is the set of vectors  $e_j$  in the representation  $x = \sum_j x_j e_j$  for which  $x_j \neq 0$ .

Crucial to all that follows is a simple “maximum principle” for real-valued functions on  $\Pi_n$  having the form “linear plus constant.” Each such function  $f$  is continuous and “respects convex combinations” in the sense that if  $x$  is a convex combination of a finite set of vectors, then  $f(x)$  is a convex combination—with the same coefficients—of its values at those vectors. The continuity of  $f$  and compactness of  $\Pi_n$  guarantee that  $f$  attains its maximum at some point of  $\Pi_n$ . This point has special properties.

**Lemma 4.1.** *Suppose  $f: \Pi_n \rightarrow \mathbb{R}$  has the form linear + constant, and attains its maximum at  $x^* \in \Pi_n$ . Then  $f \equiv f(x^*)$  on the support of  $x^*$ , and hence on the convex hull of that support.*

*Proof.* Let  $m = f(x^*) = \max\{f(x) : x \in \Pi_n\}$ . Then  $f(e_j) \leq m$  for each  $j$ . Suppose  $f(e_k) = m - \varepsilon$  for some  $k$  and some  $\varepsilon > 0$ . Let  $J = \{j : e_j \in \text{support of } x^*\}$ , so that  $x^* = \sum_{j \in J} \lambda_j e_j$ , where the  $\lambda_j$ 's are non-negative and sum to one. Thus

$$m = f(x^*) = \sum_{j \in J} \lambda_j f(e_j) \leq \sum_{j \in J \setminus \{k\}} \lambda_j m + \lambda_k (m - \varepsilon) = m - \lambda_k \varepsilon < m,$$

which is a contradiction. Thus  $f(e_j) = m$  for every  $e_j$  in the support of  $x^*$ . Since  $f$  respects convex combinations,  $f \equiv m$  on the entire convex hull of the support of  $x^*$ .  $\square$

*The Principal of Indifference.* In our work so far on the mixed-strategy extensions of both Matching Pennies and Rock-Paper-scissors we found that if, in a Nash equilibrium pair, both strategies contain in their supports all the pure strategies, then for each player: *unilateral deviation from the equilibrium does not change the payoff.* Lemma 4.1 shows that this phenomenon is no accident:

**Corollary 4.2.** *In the mixed-strategy extension of a two-person game: Each player's payoff function is constant on the convex hull of the support of that player's best response to a strategy of the other player.*

*Proof.* Suppose  $x^*$  is Player I's best response to Player II's strategy  $y$ . Then  $u_I(\cdot, y)$  is linear on the strategy set of Player I, and attains its maximum at  $x^*$ . By Lemma 4.1 it's therefore constant on the convex hull of the support of  $x^*$ . Similarly,  $u_{II}(x, \cdot)$  is constant on the the convex hull of the support of any best response by Player II to Player I's strategy  $x$ .  $\square$

In other words: The collection of Player I's best responses to a given strategy of Player II is the convex hull of Player I's pure strategy best responses. Let's not forget, however, that the best-response set we're talking about here could be just a single point (example: Prisoners' Dilemma).

*Uniqueness of Nash equilibrium for Rock-Paper-Scissors.* We've already seen that if each player mixes the r, p, and s strategies uniformly (i.e., with probability 1/3 each), then the resulting strategy pair is a Nash equilibrium for the game. The proof that it's the only one breaks down into two major cases. For the first of these, let  $\Pi_3^\circ$  denote the “interior” of  $\Pi_3$ , i.e. the set of vectors in  $\Pi_3$ , no coordinate of which is zero.

CASE I. *There is no new Nash equilibrium in  $(\Pi_3^\circ \times \Pi_3) \cup (\Pi_3 \times \Pi_3^\circ)$ .*

Fix  $y \in \Pi_3$ . Suppose  $x^* \in \Pi_3^\circ$  is a best response for Player I to Player II's strategy  $y$ , i.e. suppose

$$u_I(x^*, y) = \max\{u_I(x, y) : x \in \Pi_3\}$$

Thus  $u_I(x^*, y)$  is the maximum value of the linear function  $u_I(\cdot, y)$  over  $\Pi_3$ , so by the previous Lemma, that linear function also attains this value on the convex hull of its support which ( $x^* \in \Pi_3^\circ$ ) is all of  $\Pi_3$ . In other words,  $u_I(\cdot, y)$  is identically constant on  $\Pi_3$ , so by the second line of expression (2) we see that the coefficients of  $x^*$  must all be the same, i.e., equal to  $1/3$ .

Similarly, if  $y^* \in \Pi_3^\circ$  is a best response for Player II to Player I's strategy  $x \in \Pi_3$ , then  $y^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ . Thus:

*For the mixed-strategy extension of Rock-Paper-Scissors, the only Nash equilibrium in  $(\Pi_3^\circ \times \Pi_3) \cup (\Pi_3 \times \Pi_3^\circ)$  is the one found previously—in which each player employs the uniform strategy  $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ .*

For the second case, let  $\partial\Pi_3$  denote the “boundary” of  $\Pi_3$ , i.e. the collection of vectors in  $\Pi_3$  for which some coordinate is zero. Since  $\Pi_3$  is the union of its “interior” and its “boundary” it remains to establish:

CASE II: *No point of  $\partial\Pi_3 \times \partial\Pi_3$  can be a Nash equilibrium.*

Suppose  $(x^*, y^*) \in \partial\Pi_3 \times \partial\Pi_3$  is a Nash equilibrium. Since  $\partial\Pi_3$  is a union of three line segments in  $\mathbb{R}^3$ , there are nine cases to consider. It's enough to do one of them—the others being similar. Suppose, e.g. that  $x^* = \xi^*r + (1 - \xi^*)p = (\xi^*, 1 - \xi^*, 0)$ , and  $y^* = (\eta^*, 0, 1 - \eta^*)$ , where  $0 \leq \xi^*, \eta^* \leq 1$ . In other words, at our supposed equilibrium, Player I does not play Scissors and Player II does not play Paper. Since Player I's strategy  $x^*$  is a best response to Player II's  $y^*$ , the linear function  $u_I(\cdot, y^*)$  has its maximum on  $\Pi_3$  at  $x^*$ , and so is constant on the convex hull of the support of  $x^*$ , namely the line segment between  $r = (1, 0, 1)$  and  $p = (0, 1, 0)$ . In other words the linear function (calculated from the first line of (2))

$$\xi \rightarrow u_I((\xi, 1 - \xi, 0), y^*) = (2 - 3\eta^*)\xi + (2\eta^* - 1) \quad 0 \leq \xi \leq 1$$

is constant, hence its “ $\xi$ -coefficient”  $2 - 3\eta^*$  is zero, i.e.  $\eta^* = 2/3$ , hence  $y^* = (2/3, 0, 1/3)$ .

A similar calculation (best achieved by using the second line of (2)) using the fact that  $y^*$  is Player II's best response to  $x^*$  shows that  $x^* = (2/3, 1/3, 0)$ . To see that the strategy pair  $(x^*, y^*)$  is not a Nash equilibrium for (the mixed-strategy extension of) Rock-Paper-Scissors, let's add these strategies to the original (pure-strategy) game, obtaining the following payoff matrix for Player I:

$$\begin{array}{c|cccc} I \backslash II & r & p & s & y^* \\ \hline r & 0 & -1 & 1 & \frac{1}{3} \\ p & 1 & 0 & -1 & \frac{1}{3} \\ s & -1 & 1 & 0 & -\frac{2}{3} \\ x^* & \frac{1}{3} & -\frac{2}{3} & \frac{1}{3} & \frac{1}{3} \end{array}$$

Upon looking down the last column we see that  $x^* = \frac{2}{3}r + \frac{1}{3}p$  is Player I's best response to Player II's strategy  $y^* = \frac{2}{3}r + \frac{1}{3}s$ . However  $y^*$  is *not* Player II's best response to Player I's  $x^*$ , since  $U_{II}(x^*, y^*) = -u_I(x^*, y^*) = -1/3$ , whereas  $u_{II}(x^*, p) = 2/3$ . Thus the strategy pair  $(x^*, y^*)$  is not a Nash equilibrium.  $\square$

**Exercise.** For the mixed-strategy extension of “Battle of the Sexes:”

- (a) Show that the mixed-strategy pair  $(\frac{2}{3}B + \frac{1}{3}S, \frac{1}{3}B + \frac{2}{3}S)$  is a Nash equilibrium, and that this, along with the previously found Nash equilibria  $((B, B)$  and  $(S, S))$  are the only ones.
- (b) How does the payoff for this new equilibrium pair compare to the payoffs for the previously found ones?
- (c) Under the new equilibrium pair (as opposed to the old ones) is it more or less probable that the players will meet up?

**4.6. Nash's Theorem.** We're finally ready to state and prove Nash's famous theorem which establishes the existence of Nash equilibria for mixed-strategy extensions of finite games. As usual, we'll confine ourselves to the special case of two players.

**Theorem.** *For any non-cooperative finite game, the mixed-strategy extension has a Nash equilibrium.*

*Proof.* For notational simplicity we'll do the proof in detail for 2-player games, and then point out how the argument applies as well to the general situation.

(a) *Getting started:* Suppose Player I has  $m$  pure strategies and Player II has  $n$  of them. Then, as we just observed for "Rock-Paper-Scissors," the payoff functions for each player can be represented in terms of  $m \times n$  matrices  $A$  and  $B$ :

$$(3) \quad \begin{aligned} u_I(x, y) &= xAy^T \\ u_{II}(x, y) &= xBy^T \end{aligned}$$

where the probability vectors  $x \in \Pi_m$  and  $y \in \Pi_n$  are regarded as row matrices, and the superscript "T" means "transpose." Both payoff functions are, in each variable separately, linear and continuous.

*Best response:* Since  $\Pi_n$  is a compact subset of  $\mathbb{R}^n$ , for each fixed  $y \in \Pi_n$  the payoff function  $u_I(\cdot, y)$  attains its maximum thereon, say at  $x^* \in \Pi_m$ . Continuing with the terminology introduced in the previous section, we say "the strategy  $x^*$  is a *best response* for Player I to Player II's strategy choice  $y$ ." By Lemma 4.1, every vector in the support of  $x^*$ , and indeed in the convex hull of this support, is also a best response to  $y$ . In particular, for each  $y \in \Pi_n$ , Player I has a *pure strategy* (i.e. a standard unit vector in  $\mathbb{R}^m$ ) that is a best response to Player II's strategy  $y$ .

Similar remarks apply to the idea of Player II's best response to Player I's strategy  $x$ .

*Measuring improvement:* Since Player I's pure strategies contain a best response to Player II's strategy  $y \in \Pi_n$ , we can measure how much a strategy  $x \in \Pi_m$  differs from that best response by considering the functions:

$$\delta_i(x, y) := \max\{u_I(e_i, y) - u_I(x, y), 0\} \quad (x \in \Pi_m, y \in \Pi_n)$$

for  $i = 1, 2, \dots, m$ . Let  $\delta(x, y)$  be the vector in  $\mathbb{R}^m$  whose  $i$ -th component is  $\delta_i(x, y)$ . Then clearly:

$$x \text{ is a best response to } y \text{ if and only if } \delta(x, y) = 0.$$

*New strategies from old:* For  $(x, y) \in \Pi_m \times \Pi_n$  consider the map  $T_I: \Pi_m \times \Pi_n \rightarrow \mathbb{R}^m$  defined by

$$(4) \quad T_I(x, y) = \frac{x + \delta(x, y)}{1 + \sum_k \delta_k(x, y)}$$

The coordinates of  $T_I(x, y)$  are all non-negative, and they sum to one, i.e.,  $T_I$  maps  $\Pi_m \times \Pi_n$  into  $\Pi_m$ , and clearly does so continuously.

CLAIM:  $x \in \Pi_m$  is a best response to  $y \in \Pi_n$  if and only if  $T_I(x, y) = x$ .

*Proof of CLAIM.* We've observed that  $x$  is a best response to  $y$  if and only if  $\delta(x, y) = 0$ , and from the definition of  $T_I$  we see that if  $\delta(x, y) = 0$  then  $T_I(x, y) = x$ . So only the converse is in question.

We'll prove this converse in the contrapositive direction: Suppose  $x \in \Pi_m$  is *not* a best response to  $y \in \Pi_n$ . We wish to prove that  $T(x, y) \neq x$ .

We've observed that Player I has a pure strategy  $e_i$  (the  $i$ -th unit vector for  $\mathbb{R}^m$ ) that *is* a best response, so for some index  $i$  between 1 and  $m$  we know that  $\delta_i(x, y) > 0$ , hence  $\sum_k \delta_k(x, y) > 0$ .

On the other hand, since  $x = \sum_k x_k e_k$  is not a best response to  $y$ , the linearity of  $u_I$  in the first variable yields:

$$u_I(e_i, y) > u_I(x, y) = \sum_j x_j u_I(e_j, y).$$

Since the right-hand side is a weighted average of the numbers  $u_I(e_j, y)$ , at least one of these must be strictly less than the left-hand side. Let  $k$  denote the index of that number. Then  $\delta_k(x, y) = 0$ , hence the  $k$ -th coordinate of  $T_I(x, y)$  is

$$\frac{x_k}{1 + \sum_j \delta_j(x, y)} < x_k,$$

the strict inequality guaranteed by the fact that, on the left side, both the numerator and the sum in the denominator are positive. In particular:  $T_I(x, y) \neq x$ . This proves the CLAIM.

*The dénouement:* So far we've created a continuous map  $T_I: \Pi_m \times \Pi_n \rightarrow \Pi_m$  with the property that  $x \in \Pi_m$  is a best response for Player I to Player II's strategy  $y \in \Pi_n$  if and only if  $T_I(x, y) = x$ . In an entirely similar way there's a continuous map  $T_{II}: \Pi_m \times \Pi_n \rightarrow \Pi_n$  such that  $y \in \Pi_n$  is Player II's best response to Player I's strategy  $x \in \Pi_m$  if and only if  $T_{II}(x, y) = y$ . Finally, define  $T: \Pi_m \times \Pi_n \rightarrow \Pi_m \times \Pi_n$  by

$$T(x, y) = (T_I(x, y), T_{II}(x, y)) \quad (x, y) \in \Pi_m \times \Pi_n.$$

Then  $T$  is a continuous map of the compact, convex subset  $\Pi_m \times \Pi_n$  of  $\mathbb{R}^{m+n}$  into itself, so by the *Brouwer Fixed-Point Theorem* it has a fixed point  $(x^*, y^*)$ . Thus  $x^*$  and  $y^*$  are best responses to each other, i.e. the strategy pair  $(x^*, y^*)$  is a Nash equilibrium.

This completes the proof for 2-player games. To see that the argument works in general, note that throughout the argument above, Player II's the strategy  $y$  is "inert." It could as well be the strategy vector for the other  $N - 1$  players in an  $N$ -player game. That observation having been made, it's a routine matter to extend the proof given above to the more general situation. I leave the details to the reader.  $\square$

**4.7. The Minimax Theorem.** Suppose  $X$  and  $Y$  are any sets and  $u: X \times Y \rightarrow \mathbb{R}$  is any real-valued function. Then for any  $x_0 \in X$  and  $y_0 \in Y$  we have

$$\min_{y \in Y} u(x_0, y) \leq u(x_0, y_0) \leq \max_{x \in X} u(x, y_0)$$

whereupon

$$(5) \quad \max_{x \in X} \min_{y \in Y} u(x, y) \leq \min_{y \in Y} \max_{x \in X} u(x, y).$$

If we interpret  $X$  and  $Y$  as strategies for a two-player game, and  $u$  as the payoff function for the  $X$ -strategy player, then  $\max_{x \in X} u(x, y_0)$  is that player's best payoff when the other's strategy is  $y_0$ ; i.e. the right-hand side of (5) is the  $X$ -player's "worst-best" payoff. Similarly,  $\min_{y \in Y} u(x, y)$  is the  $X$ -player's worst payoff for playing strategy  $x$ , hence the left-hand side of (5) is that player's "best-worst" payoff. Thus inequality (5) says: "worst-best is better than best-worst."

For the special case of a mixed-strategy extension of a finite, two-person non-cooperative game, Nash's theorem shows that there is actually *equality* in (5). This is the famous *Minimax Theorem* which was, until the advent of Nash's Theorem, regarded to be the fundamental theorem of game theory.

**Theorem.** *For any  $m \times n$  real matrix  $A$  there exist vectors  $x^* \in \Pi_m$  and  $y^* \in \Pi_n$  such that:*

$$(6) \quad \max_{x \in \Pi_m} \min_{y \in \Pi_n} xAy^T = x^*Ay^{*T} = \min_{y \in \Pi_n} \max_{x \in \Pi_m} xAy^T$$

*Proof.* Consider the mixed-strategy extension of the two-person game where Player I's payoff matrix is  $A$  and Player II's is  $-A$ . Nash's theorem asserts that there exists a mixed-strategy pair  $x^* \in \Pi_m$ ,  $y^* \in \Pi_n$  such that

$$xAy^{*T} \leq x^*Ay^{*T} \quad \forall x \in \Pi_m \quad \text{and} \quad x^*(-A)y^T \leq x^*(-A)y^{*T} \quad \forall y \in \Pi_n,$$

i.e.,

$$xAy^{*T} \leq x^*Ay^{*T} \leq x^*Ay^T \quad \forall (x, y) \in \Pi_m \times \Pi_n.$$

Thus for each  $(x, y) \in \Pi_m \times \Pi_n$ :

$$\max_{x \in \Pi_m} xAy^{*T} \leq x^*Ay^{*T} \leq \min_{y \in \Pi_n} x^*Ay^T,$$

from which it follows that

$$\min_{y \in \Pi_n} \max_{x \in \Pi_m} xAy^T \leq x^*Ay^{*T} \leq \max_{x \in \Pi_m} \min_{y \in \Pi_n} xAy^T.$$

It follows from (5) that, in the above chain of inequalities, the right-hand member is  $\leq$  the left-hand one, which establishes (6).  $\square$

The Minimax Theorem asserts that for every finite, two-person, zero-sum game there is a number  $V$ , and mixed strategies  $x^*$  for Player I and  $y^*$  for Player II, such that Player I's payoff for strategy pair  $(x^*, y^*)$  is  $V$ , Player II's payoff is  $-V$ , and neither player's payoff can improve by a unilateral change of strategy. Here, of course,  $V = x^*Ay^{*T}$ , where  $A$  is the payoff matrix for Player I, so each such game has a definite *value*.

**4.8. Notes.** The "Mathematicians go to dinner" game comes from Erica Klarreich's article [2]. It's a very simple version of a situation that game theorists call "The Tragedy of the Commons," more serious instances of which concern resource depletion and environmental degradation (see [1] for the original paper on this). In [2] Klarreich gives a nice nontechnical description of game theory and Nash's influence on it.

Before Nash the fundamental theorem of game theory was von Neumann's Minimax Theorem [10] which, according to Kuhn and Tucker [3]:

... was the source of a broad spectrum of technical results, ranging from his extensions of the Brouwer fixed point theorem, developed for its proof, to new and unexpected methods for combinatorial problems.

The extension mentioned above of the Brouwer Fixed-Point Theorem is actually a precursor to Kakutani's set-valued extension of Brouwer's theorem. Kakutani's theorem was used by Nash to give an essentially one-page proof of his existence theorem [8].

For his work on equilibrium in non-cooperative games, Nash won the 1994 Nobel Prize in Economics. In the view of the noted economist Roger Myerson [4, esp. §1 and §6]:

... Nash's theory of noncooperative games should now be recognized as one of the outstanding intellectual advances of the twentieth century. The formulation of Nash equilibrium has had a fundamental and pervasive impact in economics and the social sciences which is comparable to that of the discovery of the DNA double helix in the biological sciences.

Nash's statement and proof (the one we used here) of his famous theorem occurred first in his Ph.D. thesis [7]. The first publication of this material was the short proof mentioned above, published in the Proceedings of the National Academy of Sciences [8].

At the height of his career—in his early 30's—Nash succumbed to schizophrenia, and disappeared from scientific life for about 30 years. Miraculously, he recovered in time to receive his 1994 Nobel Prize. About a month after he won the Nobel Prize, Nash's recovery was announced to the world by Sylvia Nasar in a *New York Times* article [5]. Nasar went on to write a full length biography of Nash [6], which was later adapted—*very* freely—into a major motion picture.

#### REFERENCES

- [1] Garrett Hardin, *The tragedy of the commons*, Science 162 (1968) 1243–1248.
- [2] Erica Klarreich, *The mathematics of strategy*, PNAS Classics–Game Theory. Available online at <http://www.pnas.org/site/classics/classics5.xhtml>
- [3] H. W. Kuhn and A. W. Tucker, John von Neumann's work in the theory of games and mathematical economics, Bull. Amer. Math. Soc. 64 (1958) 100–122
- [4] Roger B. Myerson, *Nash equilibrium and the history of economic theory*, J. Econ. Lit. 37 (1999) 1067–1082.
- [5] Sylvia Nasar, *The last years of a Nobel laureate*, The New York Times, November 13, 1994.
- [6] Sylvia Nasar, *A Beautiful Mind*, Simon & Schuster 1998.
- [7] John F. Nash, *Non-cooperative games*, Doctoral thesis, Princeton University, 1950
- [8] John F. Nash, *Equilibrium points in N-person games*, Proc. Nat. Acad. Sci. 36 (1950), 48–49.
- [9] John F. Nash, *Non-cooperative games*, Annals of Mathematics, Second Series 54 (1951) 286–295.
- [10] John von Neumann, A model of general economic equilibrium, The Review of Economic Studies 13 (1945) 1–9.