

Lecture Notes on Game Theory

Set 3 – Mixed Strategy Equilibria

Bernhard von Stengel

Department of Mathematics, London School of Economics
Houghton St, London WC2A 2AE, United Kingdom
email: stengel@maths.lse.ac.uk

November 16, 2004

3.1 Expected-utility payoffs

A game in strategic form does not always have a Nash equilibrium in which each player deterministically chooses one of his strategies. However, players may instead randomly select from among these *pure* strategies with certain probabilities. Randomizing one's own choice in this way is called a *mixed* strategy. A profile of mixed strategies is called a *mixed equilibrium* if no player can gain on average by unilateral deviation. Nash showed in 1951 that any finite strategic-form game has a mixed equilibrium (J. F. Nash (1951), Non-cooperative games. *Annals of Mathematics* **54**, pp. 286–295). We will show how Nash proved this theorem in Section 3.7 below.

Average (that is, *expected*) payoffs must be considered because the outcome of the game may be random. This requires that each payoff in the game represents an “expected utility”, in the sense that the payoffs can be weighted with probabilities in order to represent the player's preference for a random outcome.




Figure 3.1 One-player decision problem to decide between *comply* and *cheat*, demonstrating expected-utility payoffs. With these numbers, the player is indifferent.

As an example, Figure 3.1 shows a game with a single player who can decide to *comply* with a regulation, to buy a parking permit, or to *cheat* otherwise. The payoff when she

chooses to comply is 0. Cheating involves a 10 percent chance of getting caught and having to pay a penalty, stated as the negative payoff -90 , and otherwise a 90 percent chance of gaining a payoff of 10. With these numbers, cheating leads to a random outcome with an expected payoff of $0.9 \times 10 + 0.1 \times (-90)$, which is zero, so that the player is exactly indifferent between her two available moves.

If the payoffs are monetary amounts, each payoff unit standing for a dollar, say, one would not necessarily assume such a *risk neutrality* on the part of the player. In practice, decision-makers are typically *risk averse*, meaning they prefer the safe payoff of 0 to the gamble with an expectation of 0.

In a game-theoretic model with random outcomes, as in the game above, the payoff is not necessarily to be interpreted as money. Rather, the player's attitude towards risk is incorporated into the payoff figure as well. To take our example, the player faces a punishment or reward when cheating, depending on whether she is caught or not. Suppose that the player's decision only depends on the probability of being caught, which is 0.1 in Figure 3.1, so that she would cheat if that probability was zero. Moreover, set the reward for cheating arbitrarily to 10 units, as in the figure above, and suppose that being caught has clearly defined consequences for the player, like regret and losing money and time. Then there must be a certain probability of getting caught where the player in the above game is indifferent, say 4 percent. This determines the utility $-u$, say, for "getting caught" by the equation

$$0 = 0.96 \times 10 + 0.04 \times (-u)$$

which states equal expected utility for the choices *comply* and *cheat*. This equation is equivalent to $u = 9.6/0.04 = 240$. That is, in the above game, the negative utility -90 would have to be replaced by -240 to reflect the player's attitude towards the risk of getting caught. With that payoff, she will now prefer to comply if the probability of getting caught stays at 0.1.

The point of this consideration is to show that payoffs exist, and can be constructed, that represent player's preference for a risky outcome, as measured by the resulting *expected payoff*. These payoffs do not have to represent money. The existence of such expected-utility payoffs depends on a certain consistency of the player when facing choices with random outcomes. This can be formalized, but the respective theory, known as the *von Neumann–Morgenstern* axioms for expected utility, is omitted here for brevity.

In practice, the risk attitude of a player may not be known. A game-theoretic analysis should be carried out for different choices of the payoff parameters in order to test how much they influence the results. Often, these parameters represent the "political" features of a game-theoretic model, those most sensitive to subjective judgement, compared to the more "technical" part of a solution. In particular, there are more involved variants of the inspection game discussed in the next section. In those more complicated models, the technical part often concerns the optimal usage of limited inspection resources, like maximizing the probability of catching a player who wants to cheat. This, in turn, may imply a "political decision" when to declare that the inspectee has actually cheated. Such models and practical issues are discussed in the book by R. Avenhaus and M. Canty, *Compliance Quantified*, Cambridge University Press, 1996.

3.2 Example: Compliance inspections

Suppose a consumer purchases a license for a software package, agreeing to certain restrictions on its use. The consumer has an incentive to violate these rules. The vendor would like to verify that the consumer is abiding by the agreement, but doing so requires inspections which are costly. If the vendor does inspect and catches the consumer cheating, the vendor can demand a large penalty payment for the noncompliance.

		II	
		<i>comply</i>	<i>→ cheat</i>
I	<i>Don't inspect</i>	0	10
	<i>Inspect</i>	0	-90
		-1	-6

←

Figure 3.2. Inspection game between a software vendor (player I) and consumer (player II).

Figure 3.2 shows possible payoffs for such an inspection game. The standard outcome, defining the reference payoff zero to both vendor (player I) and consumer (player II), is that the vendor chooses *Don't inspect* and the consumer chooses to *comply*. Without inspection, the consumer prefers to *cheat* since that gives her payoff 10, with resulting negative payoff -10 to the vendor. The vendor may also decide to *Inspect*. If the consumer complies, inspection leaves her payoff 0 unchanged, while the vendor incurs a cost resulting in a negative payoff -1 . If the consumer cheats, however, inspection will result in a heavy penalty (payoff -90 for player II) and still create a certain amount of hassle for player I (payoff -6).

In all cases, player I would strongly prefer if player II complied, but this is outside of player I's control. However, the vendor prefers to inspect if the consumer cheats (since -6 is better than -10), indicated by the downward arrow on the right in Figure 3.2. If the vendor always preferred *Don't inspect*, then this would be a dominating strategy and be part of a (unique) equilibrium where the consumer cheats.

The circular arrow structure in Figure 3.2 shows that this game has no equilibrium in pure strategies. If any of the players settles on a deterministic choice (like *Don't inspect* by player I), the best response of the other player would be unique (here *cheat* by player II), to which the original choice would *not* be a best response (player I prefers *Inspect* when the other player chooses *cheat*, against which player II in turn prefers to *comply*). The strategies in a Nash equilibrium must be best responses to each other, so in this game this fails to hold for any pure strategy profile.

What should the players do in the game of Figure 3.2? One possibility is that they prepare for the worst, that is, choose a *max-min* strategy. A max-min strategy maximizes the player's worst payoff against all possible choices of the opponent. The max-min strategy for player I is to *Inspect* (where the vendor guarantees himself payoff -6), and for player II it is to *comply* (which guarantees her payoff 0). However, this is not a Nash equilibrium and hence not a stable recommendation to the two players, since player I could switch his strategy and improve his payoff.

A *mixed strategy* of player I in this game is to *Inspect* only with a certain probability. In the context of inspections, randomizing is also a practical approach that reduces costs. Even if an inspection is not certain, a sufficiently high chance of being caught should deter from cheating, at least to some extent.

The following considerations show how to find the probability of inspection that will lead to an equilibrium. If the probability of inspection is very low, for example one percent, then player II receives (irrespective of that probability) payoff 0 for *comply*, and payoff $0.99 \times 10 + 0.01 \times (-90) = 9$, which is bigger than zero, for *cheat*. Hence, player II will still cheat, just as in the absence of inspection.

If the probability of inspection is much higher, for example 0.2, then the expected payoff for *cheat* is $0.8 \times 10 + 0.2 \times (-90) = -10$, which is less than zero, so that player II prefers to *comply*. If the inspection probability is either too low or too high, then player II has a unique best response. As shown above, such a pure strategy cannot be part of an equilibrium.

Hence, the only case where player II herself could possibly randomize between her strategies is if both strategies give her the same payoff, that is, if she is *indifferent*. As stated and proved formally in Theorem 3.1 below, it is never optimal for a player to assign a positive probability to a pure strategy that is inferior, given what the other players are doing. It is not hard to see that player II is indifferent if and only if player I inspects with probability 0.1, since then the expected payoff for *cheat* is $0.9 \times 10 + 0.1 \times (-90) = 0$, which is then the same as the payoff for *comply*.

With this mixed strategy of player I (*Don't inspect* with probability 0.9 and *Inspect* with probability 0.1), player II is indifferent between her strategies. Hence, she can *mix* them (that is, play them randomly) without losing payoff. The only case where, in turn, the original mixed strategy of player I is a best response is if player I is indifferent. According to the payoffs in Figure 6, this requires player II to choose *comply* with probability 0.8 and *cheat* with probability 0.2. The expected payoffs to player I are then for *Don't inspect* $0.8 \times 0 + 0.2 \times (-10) = -2$, and for *Inspect* $0.8 \times (-1) + 0.2 \times (-6) = -2$, so that player I is indeed indifferent, and his mixed strategy is a best response to the mixed strategy of player II.

This defines the only Nash equilibrium of the game. It uses mixed strategies and is therefore called a *mixed equilibrium*. The resulting expected payoffs are -2 for player I and 0 for player II.

The preceding analysis shows that the game in Figure 3.2 has a mixed equilibrium, where the players choose their pure strategies according to certain probabilities. These probabilities have several noteworthy features.

First, the equilibrium probability of 0.1 for *Inspect* makes player II indifferent between *comply* and *cheat*. As explained in Section 3.1 above, this requires payoffs to be *expected utilities*.

Secondly, mixing seems paradoxical when the player is indifferent in equilibrium. If player II, for example, can equally well *comply* or *cheat*, why should she gamble? In particular, she could *comply* and get payoff zero for certain, which is simpler and safer. The answer is that precisely because there is no incentive to choose one strategy over the other, a player can mix, and only in that case there can be an equilibrium. If player II would *comply* for certain, then the only optimal choice of player I is *Don't inspect*, making the choice of complying not optimal, so this is not an equilibrium.

The least intuitive aspect of mixed equilibrium is that the probabilities depend on the *opponent's payoffs* and not on the player's own payoffs (as long as the qualitative preference structure, represented by the arrows, remains intact). For example, one would expect that raising the penalty -90 in Figure 3.2 for being caught lowers the probability of cheating in equilibrium. In fact, it does not. What does change is the probability of inspection, which is reduced until the consumer is indifferent.

3.3 Bimatrix games

In the following, we discuss mixed equilibria for general games in strategic form. We always assume that each player has only a finite number of given pure strategies. In order to simplify notation, we consider the case of two players. Many definitions and results carry over without difficulty to the case of more than two players.

Recall that a game in strategic form is specified by a finite set of “pure” strategies for each player, and a payoff for each player for each *strategy profile*, which is a tuple of strategies, one for each player. The game is played by each player independently and simultaneously choosing one strategy, whereupon the players receive their respective payoffs.

For two players, a game in strategic form is also called a *bimatrix game* (A, B) . Here, A and B are two payoff matrices. By definition, they have equal dimensions, that is, they are both $m \times n$ matrices, having m rows and n columns. The m rows are the pure strategies i of player I and the n columns are the pure strategies j of player II. For a row i , where $1 \leq i \leq m$, and column j , where $1 \leq j \leq n$, the matrix entry of A is a_{ij} as payoff to player I, and the matrix entry of B is b_{ij} as payoff to player II.

Usually, we depict such a game as a table with m rows and n columns, so that each cell of the table corresponds to a pure strategy pair (i, j) , and we enter both payoffs a_{ij} and b_{ij} in that cell, a_{ij} in the lower-left corner, preferably written in red if we have colours at hand, and b_{ij} in the upper-right corner of the cell, displayed in blue. The “red” numbers are then the entries of the matrix A , the “blue” numbers those of the matrix B . It does not matter if we take two matrices A and B , or a single table where each cell has two entries (the respective components of A and B).

A *mixed strategy* is a randomized strategy of a player. It is defined as a probability distribution on the set of pure strategies of that player. This is played as an “active ran-

domization”: Using a lottery device with the given probabilities, the player picks each pure strategy according to its probability. When a player plays according to a mixed strategy, the other player is not supposed to know the outcome of the lottery. Rather, it is assumed that the opponent knows that the strategy chosen by the player is a random event, and bases his or her decision on the resulting distribution of payoffs. The payoffs are then “weighted with their probabilities” to determine the *expected payoff*, which represents the player’s preference, as explained in Section 3.1.

A pure strategy is a special mixed strategy. Namely, consider a pure strategy i of player I. Then the mixed strategy x that selects i with probability one and any other pure strategy with probability zero is effectively the same as the pure strategy i , since x chooses i with certainty. The resulting expected payoff is the same as the pure strategy payoff, since any unplayed strategy has probability zero and hence does not affect the expected payoff, and the pure strategy i is weighted with probability one.

3.4 Matrix notation for expected payoffs

Unless specified otherwise, we assume that in the two-player game under consideration, player I has m strategies and player II has n strategies. The pure strategies of player I, which are the m rows of the bimatrix game, are denoted by $i = 1, \dots, m$, and the pure strategies of player II, which are the n columns of the bimatrix game, are denoted by $j = 1, \dots, n$.

A mixed strategy is determined by the probabilities that it assigns to the player’s pure strategies. For player I, a mixed strategy x can therefore be identified with the m -tuple of probabilities (x_1, x_2, \dots, x_m) that it assigns to the pure strategies $1, 2, \dots, m$ of player I. We can therefore consider x as an element of m -space (written \mathbb{R}^m). We assume that the vector x with m components is a *row vector*, that is, a $1 \times m$ matrix with a single row and m columns. This will allow us to write expected payoffs in a short way.

A mixed strategy y of player II is an n -tuple of probabilities y_j for playing the pure strategies $j = 1, \dots, n$. That is, y is an element of \mathbb{R}^n . We write y as a *column vector*, as $(y_1, y_2, \dots, y_n)^\top$, that is, the row vector (y_1, y_2, \dots, y_n) *transposed*. Transposition in general applies to any matrix. The transpose B^\top of the payoff matrix B , for example, is the $n \times m$ matrix where the entry in row j and column i is b_{ij} , since transposition means exchanging rows and columns. A column vector with n components is therefore considered as an $n \times 1$ matrix; transposition gives a row vector, a $1 \times n$ matrix.

Normally, all vectors are considered as column vectors, so \mathbb{R}^n is equal to $\mathbb{R}^{n \times 1}$, the set of all $n \times 1$ matrices with n rows and one column. We have made an exception in defining a mixed strategy x of player I as a row vector. Whether we mean row or column vectors will be clear from the context.

Suppose that player I uses the mixed strategy x and that player II uses the mixed strategy y . With these conventions, we can now succinctly express the expected payoff to player I as xAy , and the expected payoff to player II as yBx .

In order to see this, recall that the *matrix product* CD of two matrices C and D is defined when the number of columns of C is equal to the number of rows of D . That is, C is a $p \times q$

matrix, and D is a $q \times r$ matrix. The product CD is then a $p \times r$ matrix with entry $\sum_{k=1}^q c_{ik}d_{kj}$ in row i and column j , where c_{ik} and d_{kj} are the respective entries of C and D . Matrix multiplication is associative, that is, for another $r \times s$ matrix E the matrix product CDE is a $p \times s$ matrix, which can be computed either as $(CD)E$ or as $C(DE)$.

For mixed strategies x and y , we read xAy and xBy as matrix products. This works because x , considered as a matrix, is of dimension $1 \times m$, both A and B are of dimension $m \times n$, and y is of dimension n . The result is a 1×1 matrix, that is, a single real number.

It is best to think of xAy being computed as $x(Ay)$, that is, as the product of a row vector x that has m components with a column vector Ay that has m components. (The matrix product of two such vectors is also known as the *scalar product* of these two vectors.) The column vector Ay has m rows. We denote the entry of Ay in row i by $(Ay)_i$ for each row i . It is given by

$$(Ay)_i = \sum_{j=1}^n a_{ij}y_j \quad \text{for } 1 \leq i \leq m. \quad (1)$$

That is, the entries a_{ij} of row i of player I's payoff matrix A are multiplied with the probabilities y_j of their columns, so $(Ay)_i$ is the *expected payoff to player I when playing row i* . One can also think of y_j as a linear coefficient of the j th column of the matrix A . That is, Ay is the linear combination of the column vectors of A , each multiplied with its probability under y . This linear combination Ay is a vector of expected payoffs, with one expected payoff $(Ay)_i$ for each row i .

Furthermore, xAy is the expected payoff to player I when the players use x and y , since

$$x(Ay) = \sum_{i=1}^m x_i(Ay)_i = \sum_{i=1}^m x_i \sum_{j=1}^n a_{ij}y_j = \sum_{i=1}^m \sum_{j=1}^n (x_i y_j) a_{ij}. \quad (2)$$

Because the players choose their pure strategies i and j independently, the probability that they choose the pure strategy pair (i, j) is the product $x_i y_j$ of these probabilities, which is the coefficient of the payoff a_{ij} in (2).

Analogously, xBy is the expected payoff to player II when the players use the mixed strategies x and y . Here, it is best to read this as $(xB)y$. The vector xB , as the product of a $1 \times m$ with an $m \times n$ matrix, is a $1 \times n$ matrix, that is, a row vector. Each column of that row corresponds to a strategy j of player II, for $1 \leq j \leq n$. We denote the respective column entry by $(xB)_j$. It is given by $\sum_{i=1}^m x_i b_{ij}$, which is the scalar product of x with the j th column of B . That is, $(xB)_j$ is the expected payoff to player II when player I plays x and player II plays the pure strategy j . If these numbers are multiplied with the column probabilities y_j and added up, then the result is the expected payoff to player II, which in analogy to (2) is given by

$$(xB)y = \sum_{j=1}^n (xB)_j y_j = \sum_{j=1}^n \left(\sum_{i=1}^m x_i b_{ij} \right) y_j = \sum_{j=1}^n \sum_{i=1}^m (x_i y_j) b_{ij}, \quad (3)$$

which is the expected payoff to player II.

3.5 Convex combinations and mixed strategy sets

It is useful to regard mixed strategy vectors as geometric objects. A mixed strategy x of player I assigns probabilities x_i to the pure strategies i . The pure strategies, in turn, are special mixed strategies, namely the unit vectors in \mathbb{R}^m , for example $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ if $m = 3$. The mixed strategy (x_1, x_2, x_3) is then a linear combination of the pure strategies, namely $x_1 \cdot (1, 0, 0) + x_2 \cdot (0, 1, 0) + x_3 \cdot (0, 0, 1)$, where the linear coefficients are just the probabilities. Such a linear combination is called a *convex* combination since the coefficients sum to one and are nonnegative.




Figure 3.3 The line through the points x and y is given by the points $x + p(y - x)$ where $p \in \mathbb{R}$. Examples are point a for $p = 0.6$, point b for $p = 1.5$, and point c when $p = -0.4$. The line segment connecting x and y results when p is restricted to $0 \leq p \leq 1$.

Figure 3.3 shows two points x and y , here in the plane, but the picture may also be regarded as a suitable view of the situation in a higher-dimensional space. The line that goes through the points x and y is obtained by adding to the point x , regarded as a vector, any multiple of the difference $y - x$. The resulting vector $x + p \cdot (y - x)$, for $p \in \mathbb{R}$, gives x when $p = 0$, and y when $p = 1$. Figure 3.3 gives some examples a , b , c of other points. When $0 \leq p \leq 1$, like for point a , the resulting points give the *line segment* joining x and y . If $p > 1$, then one obtains points on the line through x and y on the other side of y relative to x , like the point b in Figure 3.3. For $p < 0$, the corresponding point, like c in Figure 3.3, is on that line but on the other side of x relative to y .

The expression $x + p(y - x)$ can be rewritten as $(1 - p)x + py$, where the given points x and y appear only once. This expression (with $1 - p$ as the coefficient of the first vector and p of the second) shows how the line segment joining x to y corresponds to the real interval $[0, 1]$ for the possible values of p , with the endpoints 0 and 1 of the interval corresponding to the endpoints x and y , respectively, of the line segment.

In general, a convex combination of points z_1, z_2, \dots, z_k in some space is given as any linear combination $p_1 \cdot z_1 + p_2 \cdot z_2 + \dots + p_k \cdot z_k$ where the linear coefficients p_1, \dots, p_k are nonnegative and sum to one. The previously discussed case corresponds to $z_1 = x$, $z_2 = y$, $p_1 = 1 - p$, and $p_2 = p \in [0, 1]$.

A set of points is called *convex* if it contains with any points z_1, z_2, \dots, z_k also their convex combinations. Equivalently, one can show that a set is convex if it contains with any two points also the line segment joining these two points; one can then obtain combinations of k points for $k > 2$ by iterating convex combinations of only two points.

The coefficients in a convex combination can also be regarded as probabilities, and conversely, a probability distribution on a finite set can be seen as a convex combination of the unit vectors.

In a two-player game with m pure strategies for player I and n pure strategies for player II, we denote the sets of mixed strategies of the two players by X and Y , respectively:

$$\begin{aligned} X &= \{(x_1, \dots, x_m) \mid x_i \geq 0 \text{ for } 1 \leq i \leq m, \sum_{i=1}^m x_i = 1\}, \\ Y &= \{(y_1, \dots, y_n)^\top \mid y_j \geq 0 \text{ for } 1 \leq j \leq n, \sum_{j=1}^n y_j = 1\}. \end{aligned} \tag{4}$$

For consistency with Section 3.4, we assume that X contains row vectors and Y column vectors, but this is not an important concern.




Figure 3.4 Examples of player I's mixed strategy set X when $m = 2$ (left) and $m = 3$ (right), as the set of convex combinations of the unit vectors.

Examples of X are shown in Figure 3.4. When $m = 2$, then X is just the line segment joining $(1, 0)$ to $(0, 1)$. If $m = 3$, then X is a triangle, given as the set of convex combinations of the unit vectors, which are the vertices of the triangle.

It is easily verified that in general, X and Y are convex sets.

3.6 The best response condition

A mixed strategy equilibrium is a profile of mixed strategies such that no player can improve his expected payoff by unilaterally changing his own strategy. In a two-player game, an

equilibrium is a pair (x, y) of mixed strategies such that x is a best response to y and vice versa. That is, player I cannot get a better expected payoff than xAy by choosing any other strategy than x , and player II cannot improve her expected payoff xBy by changing y .

It seems not easy to decide if x is a best response to y among all possible mixed strategies, that is, if x maximizes xAy for all x in X , since X is an infinite set. However, the following theorem, known as the *best response condition*, shows how to recognize this. This theorem is not difficult but important to understand. We discuss it afterwards.

Theorem 3.1 (Best response condition.) *Let x and y be mixed strategies of player I and II, respectively. Then x is a best response to y if and only if for all pure strategies i of player I,*

$$x_i > 0 \implies (Ay)_i = \max\{(Ay)_k \mid 1 \leq k \leq m\}. \quad (5)$$

Proof. Recall that $(Ay)_i$ is the i th component of Ay , which is the expected payoff to player I when playing row i , according to (1). Let $u = \max\{(Ay)_k \mid 1 \leq k \leq m\}$, which is the maximum of these expected payoffs for the *pure* strategies of player I. Then

$$\begin{aligned} xAy &= \sum_{i=1}^m x_i (Ay)_i = \sum_{i=1}^m x_i (u - (u - (Ay)_i)) = \sum_{i=1}^m x_i u - \sum_{i=1}^m x_i (u - (Ay)_i) \\ &= u - \sum_{i=1}^m x_i (u - (Ay)_i). \end{aligned} \quad (6)$$

Since for any pure strategy i , both x_i and the difference of the maximum payoff u and the payoff $(Ay)_i$ for row i is nonnegative, the sum $\sum_{i=1}^m x_i (u - (Ay)_i)$ is also nonnegative, so that $xAy \leq u$. The expected payoff xAy achieves the maximum u if and only if that sum is zero, that is, if $x_i > 0$ implies $(Ay)_i = u$, as claimed. \square

Consider the phrase “ x is a best response to y ” in the preceding theorem. This means that among all mixed strategies in X of player I, x gives maximum expected payoff to player I. However, the pure best responses to y in (5) only deal with the pure strategies of player I. Each such pure strategy corresponds to a row i of the payoff matrix. In that row, the payoffs a_{ij} are multiplied with the column probabilities y_j , and the sum over all columns gives the expected payoff $(Ay)_i$ for the pure strategy i according to (1). This pure strategy is a best response if and only if no other row gives a higher payoff.

The first point of the theorem is that the condition whether a pure strategy is a best response or not is very easy to check, as one only has to compute the m expected payoffs $(Ay)_i$ for $i = 1, \dots, m$. For example, if player I has three pure strategies ($m = 3$), and the expected payoffs in (1) are $(Ay)_1 = 4$, $(Ay)_2 = 4$, and $(Ay)_3 = 3$, then only the first two strategies are pure best responses. If these expected payoffs are 3, 5, and 3, then only the second strategy is a best response. Clearly, at least one pure best response exists, since the numbers $(Ay)_k$ in (5) have their maximum u for at least one k . The theorem states that only pure best responses i may have positive probability x_i if x is to be a best response to y .

A second consequence of Theorem 3.1, used also in its proof, is that a mixed strategy can never give a higher payoff than the best pure strategy. This is intuitive since “mixing”

amounts to *averaging*, which is an average weighted with the probabilities, in the way that the overall expected payoff xAy in (6) is obtained from those in (1) by multiplying (weighting) each row i with weight x_i and summing over all rows i , as shown in (2). Consequently, any pure best response i to y is also a mixed best response, so the maximum of xAy for $x \in X$ is the same as when x is restricted to the unit vectors in \mathbb{R}^m that represent the pure strategies of player I.

3.7 Existence of mixed equilibria

In this section, we give the original proof of John Nash from 1951 that shows that any game with a finite number of players, and finitely many strategies per player, has a mixed equilibrium. This proof uses the following theorem about continuous functions.

Theorem 3.2 (Brouwer's Fixed Point Theorem) *Let S be a subset of some space \mathbb{R}^N that is convex and compact,¹ and let f be a continuous function from S to S . Then f has at least one fixed point, that is, a point s in S so that $f(s) = s$.*

Theorem 3.3 (Nash [1951].) *Every finite game has at least one equilibrium in mixed strategies.*

Proof. We will give the proof for two players, to simplify notation. It extends in the same manner to any finite number of players. The set S that is used in the present context is the product of the sets of mixed strategies of the players. Let X and Y be the sets of mixed strategies of player I and player II as in (4), and let $S = X \times Y$.

Then the function $f: S \rightarrow S$ that we are going to construct maps a pair of mixed strategies (x, y) to another pair $f(x, y) = (\bar{x}, \bar{y})$. Intuitively, a mixed strategy probability x_i (of player I, similarly y_j of player II) is changed to \bar{x}_i , such that it will decrease if the pure strategy i does worse than the average of all pure strategies. In equilibrium, all pure strategies of a player that have positive probability do equally well, so no sub-optimal pure strategy can have a probability that is reduced further. This means that the mixed strategies do not change, so this is indeed equivalent to the fixed point property $(x, y) = (\bar{x}, \bar{y}) = f(x, y)$.

In order to define f as described, consider the following functions $\chi: X \times Y \rightarrow \mathbb{R}^m$ and $\psi: X \times Y \rightarrow \mathbb{R}^n$ (we do not worry whether these vectors are row or column vectors; it suffices that \mathbb{R}^m contains m -tuples of real numbers, and similarly \mathbb{R}^n contains n -tuples). For each pure strategy i of player I, let $\chi_i(x, y)$ be the i th component of $\chi(x, y)$, and for each pure strategy j of player II, let $\psi_j(x, y)$ be the j th component of $\psi(x, y)$. The functions χ and ψ are defined by

$$\chi_i(x, y) = \max\{0, (Ay)_i - xAy\}, \quad \psi_j(x, y) = \max\{0, (xB)_j - xBy\},$$

for $1 \leq i \leq m$ and $1 \leq j \leq n$. Recall that $(Ay)_i$ is the expected payoff to player I against y when he uses the pure strategy i , and that $(xB)_j$ is the expected payoff to player II against x

¹In this context, a set is compact if it is closed (containing any points near the set) and bounded.

when she uses the pure strategy j . Moreover, xAy and xBy are the overall expected payoffs to player I and player II, respectively. So the difference $(Ay)_i - xAy$ is positive if the pure strategy i gives more than the average xAy against y , zero if it gives the same payoff, and negative if it gives less. The term $\chi_i(x, y)$ is this difference, except that it is replaced by zero if the difference is negative. The term $\psi_j(x, y)$ is defined analogously. Thus, $\chi(x, y)$ is a nonnegative vector in \mathbb{R}^m , and $\psi(x, y)$ is a nonnegative vector in \mathbb{R}^n . The functions χ and ψ are continuous.

The pair of vectors (x, y) is now changed by replacing x by $x + \chi(x, y)$ in order to get \bar{x} , and y by $y + \psi(x, y)$ to get \bar{y} . Both sums are nonnegative. The only problem is that in general, these new vectors are no longer probabilities since their components do not sum to one. For that purpose, they are “re-normalized” by the following functions $r : \mathbb{R}^m \rightarrow \mathbb{R}^m$ and $s : \mathbb{R}^n \rightarrow \mathbb{R}^n$, defined by their components r_i and s_j , that is, $r(x) = (r_1(x), \dots, r_m(x))$, and $s(y) = (s_1(y), \dots, s_n(y))$:

$$r_i(x_1, \dots, x_m) = \frac{x_i}{\sum_{k=1}^m x_k}, \quad s_j(y_1, \dots, y_n) = \frac{y_j}{\sum_{k=1}^n y_k}.$$

Clearly, if $x_i \geq 0$ for $1 \leq i \leq m$ and $\sum_{k=1}^m x_k > 0$, then $r(x)$ is defined and is a probability distribution, that is, an element of the mixed strategy set X . Analogously, $s(y) \in Y$.

The function $f: X \rightarrow Y$ is now defined by

$$f(x, y) = (r(x + \chi(x, y)), s(y + \psi(x, y))).$$

What is a fixed point (x, y) of that function, so that $f(x, y) = (x, y)$? Consider the smallest pure strategy payoff $(Ay)_i$ against y , that is, $(Ay)_i = \min_k (Ay)_k$. Then $(Ay)_i \leq xAy$, which is proved analogously to (6), so the component $\chi_i(x, y)$ of $\chi(x, y)$ is zero. This means that the respective term $x_i + \chi_i(x, y)$ is equal to x_i . Conversely, consider some other pure strategy l of player I that gets the maximum payoff $(Ay)_l = \max_l (Ay)_k$. If that payoff is better than the average xAy , then clearly $\chi_l(x, y) > 0$, so that $x_l + \chi_l(x, y) > x_l$. Since $\chi_k(x, y) \geq 0$ for all k , this implies $\sum_{k=1}^m (x_k + \chi_k(x, y)) > 1$, which is the denominator in the re-normalization with r in $r(x + \chi(x, y))$. This re-normalization will now *decrease* the value of x_i for the pure strategy i with $(Ay)_i \leq xAy$, so the relative weight x_i of the pure strategy i decreases, or is unchanged if $x_i = 0$. But $(Ay)_l > xAy$ can only occur if there is some sub-optimal strategy i (with $(Ay)_i \leq xAy < (Ay)_l$) that has positive probability x_i , which can be used instead of the strategy l that gives minimum expected payoff $(Ay)_i$. In that case, $r(x + \chi(x, y))$ is *not equal* to x , so that $f(x, y) \neq (x, y)$.

Analogously, if $\psi(x, y)$ has some component that is positive, then the respective pure strategy of player II has a better payoff than xBy , so $y \neq s(y + \psi(x, y))$ and (x, y) is not a fixed point of f . In that case, y is also not a best response to x .

Hence, the function f has a fixed point (x, y) if and only if both $\chi(x, y)$ and $\psi(x, y)$ are zero in all components. But that means that xAy is the maximum possible payoff $\max_i (Ay)_i$ against y , and xBy is the maximum possible payoff $\max_j (xB)_j$ against x , that is, x and y are mutual best responses. The fixed points (x, y) are therefore exactly the Nash equilibria of the game. \square