

BEEM105 – Game Theory and Applications	Dieter Balkenborg
Lecture	Departments of Economics
2 × 2-games	University of Exeter

1 Mixed strategies in 2 × 2-games

1\2	L	R
T	a_{11}	a_{12}
B	a_{21}	a_{22}

We now consider *mixed strategies*. Let $0 \leq p \leq 1$ be the probability with which the row player chooses her bottom strategy and let $0 \leq q \leq 1$ be the probability with which the column plays his right strategy. A *mixed strategy pair* (or a *mixed strategy combination*) is then described by a pair of numbers (p, q) . In general I find it often convenient to write a mixed strategy as a formal sum like $(1 - p)T + pB$. So $\frac{2}{3}T + \frac{1}{3}B$ would be the mixed strategy where the row player uses his bottom strategy with probability $p = \frac{1}{3}$. The mixed strategy where the row player takes his bottom strategy with certainty would be $0T + 1B = B$, it would be identified with his pure strategy B .

Given a mixed strategy pair (p, q) the expected payoff of the row player is

$$u_1(p, q) = (1 - p)((1 - q)a_{11} + qa_{12}) + p((1 - q)a_{21} + qa_{22})$$

whereas the expected payoff of the column player is

$$u_2(p, q) = (1 - q)((1 - p)b_{11} + pb_{12}) + q((1 - p)b_{21} + pb_{22})$$

The payoff functions are *bilinear*, i.e. for given p the payoff $u_i(p, q)$ ($i = 1, 2$) is a linear function in q and for given q it is a linear function in p . In our discussion of zero sum games we will look at the functions $u_1(p, 0)$ and $u_1(p, 1)$ as they vary in p , i.e. we will graph the payoff of the row player against his own strategy choice p . In our discussion of Nash equilibria for arbitrary 2×2 -games we will consider the functions $u_1(0, q)$ and $u_1(1, q)$ as functions in q and graph the payoffs of the player against strategy choices of his opponent.

Instead of $u_1(p, 0)$ one could also write $u_1(p, T)$ etc.

2 Maximin Strategies in zero sum game

We consider the zero-sum game

	L	R
T	+3	-3
B	-3	+1

As is standard for zero-sum games, only the payoff to the row players is indicated. The column player's utility is $u_2(p, q) = -u_1(p, q)$.


The row player's expected utility is

$$u_1(p, q) = 3(1 - p)(1 - q) - 3((1 - p)q + p(1 - q)) + pq$$

In the following we *keep the strategy of the column player fixed*. We the look at how the *row player's payoff depends on his own choice of probability p* with which he plays his bottom strategy. For instance, for $q = 0, 0.5, 1$ we obtain

$$\begin{aligned} u_1(p, 0) &= 3(1 - p) - 3p \\ u_1(p, 0.5) &= \frac{3}{2}(1 - p) - \frac{3}{2} + \frac{1}{2}p \\ u_1(p, 1) &= -3(1 - p) + p \end{aligned}$$

The resulting functions in p only are drawn in the following diagram on the left.



We see from the graph on the left that for any p the minimal payoff player 1 might get is $\min \{u_1(p, 0), u_1(p, 1)\}$ and the most he could get is $\max \{u_1(p, 0), u_1(p, 1)\}$. Any other expected payoff $u_1(p, q)$ is in between. When he chooses p such that

$$u_1(p, 0) = u_1(p, 1) \Leftrightarrow -3(1-p) + p = 3(1-p) - 3p \Leftrightarrow p = 0.6$$

then he will get the expected payoff -0.6 , regardless of how his opponent chooses. As indicated in the graph on the right, the minimal payoff the row player can get when he uses his bottom strategy with probability p is

$$\min_{0 \leq q \leq 1} u_1(p, q) = \begin{cases} -3(1-p) + p & \text{for } p < 0.6 \\ 3(1-p) - 3p & \text{for } p \geq 0.6 \end{cases}$$

We have

$$\begin{aligned} \max_{0 \leq p \leq 1} \min_{0 \leq q \leq 1} u_1(p, q) &= \min_{0 \leq q \leq 1} u_1(0.6, q) = -0.6 \\ \arg \max_{0 \leq p \leq 1} \left(\min_{0 \leq q \leq 1} u_1(p, q) \right) &= p^* = 0.6 \end{aligned}$$

Notation 1 Whenever a function $f(x)$ is to be maximized one writes

$$\arg \max_x f(x)$$


for the set of values x at which $f(x)$ is maximized, i.e.

$$x^* \in \arg \max_x f(x) \Leftrightarrow f(x^*) = \max_x f(x)$$

Thus the maximin strategy of player 1 is $p^* = 0.6$ and the value of the game is -0.6 .

We can identify the maximin strategy of player 2 similarly

$$\begin{aligned} u_2(p, q) &= -3(1-p)(1-q) + 3((1-p)q + p(1-q)) - pq \\ u_2(0, q) &= -3(1-q) + 3q \\ u_2(1, q) &= 3(1-q) - q \\ u_2(0, q) &= u_2(1, q) \Leftrightarrow -3(1-q) + 3q = 3(1-q) - q \Leftrightarrow q = 0.6 \end{aligned}$$



$$\min_{0 \leq p \leq 1} u_2(p, q) = \begin{cases} -3(1-q) + 3q & \text{for } q < 0.6 \\ 3(1-q) - q & \text{for } q \geq 0.6 \end{cases}$$


Thus the maximin strategy of the columns player is also given by $q^* = 0.6$.

Remark 1 It is not generally true for 2×2 games that the maximin strategies satisfy $p^* = q^*$. This is specific to the example.

Remark 2 We have identified the maximin strategy q^* of the column player by the condition $u_2(0, q^*) = u_2(1, q^*)$, i.e. the column player gets the same payoff regardless of which (pure) strategy the opponent is using. Since $u_1(p, q) = -u_2(p, q)$ this is equivalent to the condition $u_1(0, q^*) = u_1(1, q^*)$, i.e. **when the column player uses q^* then the row player is indifferent between his pure strategies T and B** . In a moment we will use this type of condition to find Nash equilibria in mixed strategies.


Remark 3 Let us verify that $p^* = q^* = 0.6$ is a saddle point. When the column player uses $q^* = 0.6$ the the row player's payoff is

$$\begin{aligned} u_1(0, 0.6) &= 3 \times 0.4 - 3 \times 0.6 = -0.6 \\ u_1(1, 0.6) &= -3 \times 0.4 + 0.6 = -0.6 \\ u_1(p, 0.6) &= -0.6 \end{aligned}$$



Thus the saddle point condition $u_1(p, q^*) \leq u_1(p^*, q^*)$ for all $0 \leq p \leq 1$ is trivially satisfied. Since $u_1(p^*, 0) = u_1(p^*, 1) = -0.6$ by construction we have $u_1(p^*, q^*) = u_1(p^*, q)$ for all $0 \leq q \leq 1$. Hence the other saddle-point condition $u_1(p^*, q^*) \leq u_1(p^*, q)$ for all $0 \leq q \leq 1$ is also satisfied.

Remark 4 For general 2×2 zero-sum games one finds the row player's maximin strategy condition by the condition $u_1(p, 0) = u_1(p, 1)$ only when the maximin strategy is mixed. Otherwise the equation has typically no solution in the interval $0 \leq p \leq 1$. In the following figure the graphs of the functions $u_1(p, 0) = p$ and $u_1(p, 1) = p + 1$ are parallel lines. The maximin strategy of player 1 would be $p = 1$.



3 O'Neals game

Two players simultaneously and independently have to name one of the four digits 0, 1, 2 or 3. The following table indicates the outcome. “+” means that player 1 (the row player) wins. She receives

payoff +1, player 2 (the column player) receives payoff -1. Correspondingly, a “-” means that player 1 loses. She receives payoff -1 and player 2 payoff +1.

1\2	0	1	2	3
0	+	-	-	-
1	-	-	+	+
2	-	+	-	+
3	-	+	+	-

Let us find the saddle point in O’Neals game. A mixed strategy for the row player now takes the form (p_0, p_1, p_2, p_3) with $p_0, p_1, p_2, p_3 \geq 0$ and $p_0 + p_1 + p_2 + p_3 = 1$. For the column player is takes the form (q_0, q_1, q_2, q_3) with $q_0, q_1, q_2, q_3 \geq 0$ and $q_0 + q_1 + q_2 + q_3 = 1$.

Symmetry arguments that both players should randomize with equal probability between their strategies 1, 2 and 3. We are hence interested in mixed strategies of the form

$$\left(1 - p, \frac{p}{3}, \frac{p}{3}, \frac{p}{3}\right) \text{ and } \left(1 - q, \frac{q}{3}, \frac{q}{3}, \frac{q}{3}\right).$$

We have the expected utilities

$$\begin{aligned} u_1((1, 0, 0, 0), (1, 0, 0, 0)) &= 1 \\ u_1\left((1, 0, 0, 0), \left(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)\right) &= -1 \\ u_1\left(\left(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), (1, 0, 0, 0)\right) &= -1 \\ u_1\left(\left(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \left(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)\right) &= \frac{1}{3} \end{aligned}$$

Let $T = L = (1, 0, 0, 0)$ and $B = R = \left(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$. Restricting attention to these two strategies gives us a 2×2 zero sum game

	L	R
T	+1	-1
B	-1	$+\frac{1}{3}$

which is, up to the multiplication of th payoffs by a scalar, the game we have analyzed. We know that the maximin strategies for this game are $p^* = q^* = 0.6$.

Hence it is tempting to conjecture that $\vec{p}^* = (p_0, p_1, p_2, p_3) = (0.4, 0.2, 0.2, 0.2)$ and $\vec{q}^* = (q_0, q_1, q_2, q_3) = (0.4, 0.2, 0.2, 0.2)$ are the maximin strategies in O’Neals game. To prove this we form that they form a saddle point. We have

$$\begin{aligned} u_1((1, 0, 0, 0), \vec{q}^*) &= 0.4 - 0.2 - 0.2 - 0.2 = -0.2 \\ u_1((0, 1, 0, 0), \vec{q}^*) &= -0.4 - 0.2 + 0.2 + 0.2 = -0.2 \\ u_1((0, 0, 1, 0), \vec{q}^*) &= -0.4 + 0.2 - 0.2 + 0.2 = -0.2 \\ u_1((0, 0, 0, 1), \vec{q}^*) &= -0.4 + 0.2 + 0.2 - 0.2 = -0.2 \end{aligned}$$

and hence

$$u_1((p_0, p_1, p_2, p_3), \vec{q}^*) = -0.2$$

for all mixed strategies (p_0, p_1, p_2, p_3) of the row player. This proves the first condition for a saddle point. Since

$$u_1(\vec{p}^*, (q_0, q_1, q_2, q_3)) = -0.2$$

for every mixed strategy (q_0, q_1, q_2, q_3) of the column player by a very similar calculation the second condition for a saddle-point is also satisfied.

4 Dominance


From now on we use “the other type of graph” where we graph the payoff of a player against the probability mixture of the opponent. Consider the 3×2 -game

$1 \setminus 2$	L	R
T	0	3
M	1	2
B	3	0
	1	1

Let q be the probability with which the column player plays R . We have

$$\begin{aligned}
 u_1(T, q) &= 3(1 - q) \\
 u_1(M, q) &= 3q \\
 u_1(B, q) &= (1 - q) + q = 1 \\
 u_1\left(\frac{1}{2}T + \frac{1}{2}M, q\right) &= \frac{1}{2} \times 3(1 - q) + \frac{1}{2} \times 3q = \frac{3}{2}
 \end{aligned}$$

which yields the graph



We see that the row player’s strategy B is *strictly dominated* by the mixed strategy $\frac{1}{2}T + \frac{1}{2}M$. Thus a pure strategy that is not strictly dominated by any pure strategy may be strictly dominated by a mixed strategy.

Exercise 1 Does the game have a unique Nash equilibrium?

BEEM105 – Game Theory and Applications	Dieter Balkenborg
Lecture	Departments of Economics
2 × 2-games	University of Exeter

5 Best replies

We analyze the game


1\2	L	R
T	0	-1
B	-1	-c

for variable c . For sufficiently large c we get a “chicken” game. Our aim is to find all Nash equilibria of the game. We start with $c = 5$.

We have

$$\begin{aligned}
 u_1(T, q) &= -q & u_1(B, q) &= (1 - q) - qc = 1 - (c + 1)q = 1 - 6q \text{ for } c = 5 \\
 u_2(p, L) &= -p & u_2(p, B) &= (1 - p) - pc = 1 - (c + 1)p = 1 - 6p \text{ for } c = 5
 \end{aligned}$$

We obtain the following graphs



The figure on the left tells us the following: When the column player chooses his right strategy R with a sufficiently low probability it is optimal for the row player to play his top strategy. Any other strategy, pure or mixed, gives a strictly lower payoff. For high values of q the unique *best reply* (or *best response*) for the row player is to play his bottom strategy. If and only if the two lines in the graph intersect, i.e. when


$$u_1(T, q) = u_1(B, q) \Leftrightarrow -q = 1 - (c + 1)q \Leftrightarrow cq = 1 \Leftrightarrow q = \frac{1}{c}$$

and so for $c = 5$ when $q = \frac{1}{5}$ the two pure strategies of the row player yield the same expected payoff, namely $-\frac{1}{5}$. Against $q = \frac{1}{5}$ all mixed strategies of the row player are best responses. The information from the graph is summarized in the *best reply correspondence for the row player*

$$\beta_1((1 - q)L + qR) = \begin{cases} \{B\} & \text{for } 0 \leq q < \frac{1}{c} \\ \{(1 - p)T + pB \mid 0 \leq p \leq 1\} & \text{for } q = \frac{1}{c} \\ \{T\} & \text{for } \frac{1}{c} < q \leq 1 \end{cases}$$

This is the function which assigns to every strategy of the column player the *set of best replies* of the row player against the strategy. One speaks of a *correspondence* rather than a function because the values are sets rather than individual strategies. Otherwise it is very similar to the *reaction function* we had in

the Cournot game. We can graph the best reply correspondence as below on the left:




Similarly we have a best reply correspondence for the column player

$$\beta_2((1-p)T + pB) = \begin{cases} \{R\} & \text{for } 0 \leq p < \frac{1}{c} \\ \{(1-q)L + qR \mid 0 \leq q \leq 1\} & \text{for } p = \frac{1}{c} \\ \{L\} & \text{for } \frac{1}{c} < p \leq 1 \end{cases}$$

who's graph is given on the right.

6 Mixed Nash equilibria

Per definition, the Nash equilibria are the points of intersection of the graphs of the two best reply correspondences.



We have three Nash equilibria for $c = 5$: (B, L) , (T, R) and $(\frac{1}{5}T + \frac{4}{5}B, \frac{1}{5}L + \frac{4}{5}R)$. Notice that the equilibrium in mixed strategies is identified by the following conditions: Player 1 must mix such that

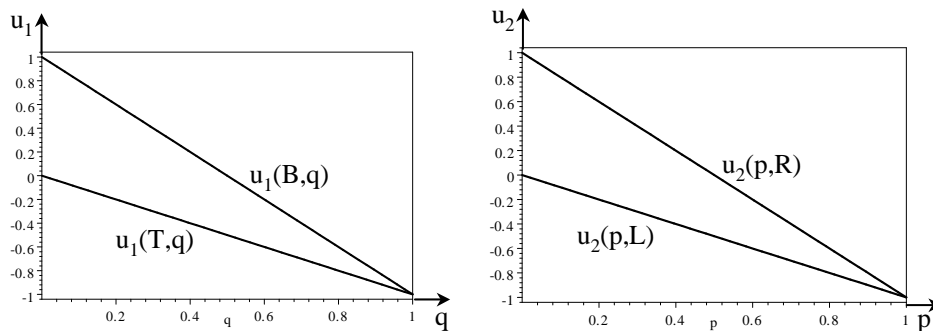
player 2 is indifferent between his pure strategies (p must solve $u_2(p, L) = u_2(p, R)$) and conversely player 2 must mix such that player 1 is indifferent between his two strategies (q must solve $u_1(T, q) = u_1(B, q)$).

In general, for sufficiently large c we have three Nash equilibria: The two pure ones (B, L) and (T, R) and the mixed one $(\frac{c-1}{c}T + \frac{1}{c}B, \frac{c-1}{c}L + \frac{1}{c}R)$. When the costs of an accident increases ($c \rightarrow +\infty$) the probability of an accident $\frac{1}{2}$ in the mixed strategy Nash equilibrium goes to zero. The outcomes of the three Nash equilibria are $(1, -1)$, $(-1, 1)$ and $(-\frac{1}{c}, -\frac{1}{c})$. (Since in the mixed strategy Nash equilibrium the row player gets the same utility with both pure strategies, when he takes T he gets $\frac{c-1}{c} \times 0 + \frac{1}{c} \times (-1) = -\frac{1}{c}$.) The three equilibrium outcomes are not Pareto-comparable. It is not possible to make both players better off by switching from one equilibrium to another. When we use a utilitarian welfare function and sum the utilities the pure Nash equilibria look slightly better because they yield the sum 0 while the mixed strategy Nash equilibrium yields the sum $-\frac{2}{c}$. Still, for large c the difference is minor.


By c “sufficiently large” I mean hereby that $c > 1$.

For $c = 1$ the strategies T and L become weakly dominated.

$1 \setminus 2$	L	R
T	0	-1
B	1	-c




The best reply correspondences of both players have identical graphs.




This graph also gives the set of all Nash equilibria: (p, q) is a Nash equilibrium iff $(1-p)(1-q) = 1$. There is a continuum of Nash equilibria consisting of two line segments.

When $c < 1$, e.g. when $c = 0$, the strategies T and L are strictly dominated.



The unique Nash equilibrium is (B, R) .



For varying c the Nash equilibrium correspondence has roughly the shape of a pitchfork.




BEEM105 – Game Theory and Applications	Dieter Balkenborg
Lecture	Departments of Economics
3×3-games	University of Exeter

We present a graphical method to determine all Nash equilibria of a 3×3 -game. As an example we use the game

		<i>a</i>	<i>b</i>	<i>c</i>
<i>A</i>		1	−3	0
	1	−3	−3	−3
<i>B</i>		−3	1	0
	−3	−3	1	−3
<i>C</i>		−3	−3	0
	0	0	0	0

7 Mixed Strategies

We denote a mixed strategy of the row player by (p_A, p_B, p_C) , whereby $p_A, p_B, p_C \geq 0$ and $p_A + p_B + p_C = 1$, and a mixed strategy of the column player by (q_a, q_b, q_c) , whereby $q_a, q_b, q_c \geq 0$ and $q_a + q_b + q_c = 1$. Each strategy (p_A, p_B, p_C) of the row player defines a point (p_B, p_C) in the triangle defined by the inequalities $p_B \geq 0$, $p_C \geq 0$ and $p_B + p_C \leq 1$. Conversely, any such point (p_B, p_C) defines a mixed strategy (p_A, p_B, p_C) by setting $p_A = 1 - p_B - p_C$. In particular, the point $(0, 0)$ corresponds to the mixed strategy $(1, 0, 0)$ and hence to the pure strategy *A*. Similarly, $(1, 0)$ corresponds to $(0, 1, 0)$ or *B* and $(0, 1)$ to $(0, 0, 1)$ or *C*. Corresponding statement hold for the mixed strategies of the column player.



8 Best Reply Diagrams

In order to determine all Nash equilibria in mixed strategies we have to determine for each mixed strategy (p_A, p_B, p_C) of the row player the possible best responses of the column player against it. This information can be encoded in a graph representing the row player's mixed strategy set. It suffices to determine which pure strategies of the column player are best responses against strategies of the row player because a mixed strategy is a best response if and only if all pure strategies it uses with positive probability are themselves best responses.¹

¹The reason is roughly follows. In whichever way the column player randomizes between his various pure strategies, he can never get a higher expected payoff than he can get with any of the pure strategies he is using. For instance, if he is randomizing between one pure strategies giving him the expected payoff 2 and another giving him 1, then he cannot get an

To obtain the Nash equilibria of the game given in the exercise we calculate as the first step the expected payoffs that the column player would get with each of his pure strategies a , b , c if his opponent uses a mixed strategy (p_A, p_B, p_C) . These are, corresponding to his payoffs in the three rows of the normal form,

$$\begin{aligned} v_a &= p_A - 3p_B - 3p_C \\ v_b &= -3p_A + p_B - 3p_C \\ v_c &= 0. \end{aligned}$$

The condition $v_a = v_b$ identifies a line in the triangle representing the row player's mixed strategies. When the row player plays any mixed strategy corresponding to a point on this line, the column player is indifferent between playing a or b or any mixture of the two because he will always get the same expected payoff. However, it could be better for him to play c . The indifference condition $v_a = v_b$ is equivalent to

$$p_A - 3p_B - 3p_C = -3p_A + p_B - 3p_C \quad \text{or} \quad 4p_A = 4p_B \quad \text{or} \quad p_A = p_B$$

To draw the line we identify the points where the line intersects the boundary of the triangle. These we obtain by setting one probability zero and then determine the other two from the indifference condition and the condition that the sum of the remaining two probabilities must be 1. If we set $p_C = 0$ we obtain in this way $p_A = p_B$ and $p_A + p_B = 1$ and hence that $p_A = p_B = \frac{1}{2}$. Thus we obtain the mixed strategy $(p_A, p_B, p_C) = (\frac{1}{2}, \frac{1}{2}, 0)$ and the corresponding point $(\frac{1}{2}, 0)$ in the diagram. Setting $p_A = 0$ we obtain from $p_A = p_B$ and $p_A + p_B = 1$ that $p_B = 0$ and $p_C = 1$, i.e., we obtain the mixed strategy $(p_A, p_B, p_C) = (0, 0, 1)$ and the corresponding point $(0, 1)$ in the diagram. We can now draw the indifference line because it must pass through the two points, as indicated in Figure 1.




Figure 1

The indifference line splits the triangle into two regions. One region corresponds to the set of mixed strategies of the row player such that when she plays any of these, the column player prefers to play a rather than b against it because a yields the higher expected payoff. Against strategies corresponding to points in the other region the column player would prefer to play b rather than a . To identify the regions, we take a vertex of the triangle which is not on the indifference line and determine to which of the two regions it belongs. For instance, the vertex $(0, 0)$ is not on the indifference line and corresponds to the pure strategy A . From the column player's payoffs in the normal form we see that if the row player plays A , the column player would prefer to play a (which would yield payoff 1) rather than b (which would yield expected payoff higher than 2. (The behaviour of the opponent is assumed to be fixed in this argument.) Consequently, he must have at least one pure strategy which gives him the highest expected payoff among all his pure or mixed strategies, i.e., there must be pure best response. If there are several pure best responses yielding the same highest payoff then, however he randomizes between them, he is sure to get the highest payoff. However, if he randomizes such that with positive probability a pure strategy is used which does not yield the highest payoff, then this mix cannot be optimal. With positive probability he will then use this suboptimal strategy and get a corresponding low payoff although he could clearly avoid this by using a better pure strategy instead.

yield payoff -3). Thus the dark shaded region in the above diagram is the one where a is preferred to b because it contains the point corresponding to A . We indicate this by writing “ $a > b$ ” in the shaded area. The other region must be the one where b is preferred to a .

Next, the indifference condition $v_b = v_c$ yields $-3p_A + p_B - 3p_C = 0$ or

$$p_B = 3p_A + 3p_C.$$

Setting $p_A = 0$ we obtain from the indifference condition and $p_B + p_C = 1$ the mixed strategy $(p_A, p_B, p_C) = (0, \frac{3}{4}, \frac{1}{4})$. Setting $p_C = 0$ and using $p_A + p_B = 1$ we obtain the mixed strategy $(p_A, p_B, p_C) = (\frac{1}{4}, \frac{3}{4}, 0)$. The points $(\frac{3}{4}, \frac{1}{4})$ and $(\frac{3}{4}, 0)$ identify the indifference line in Figure 2 below. Since against B the column player prefers b to c , b is preferred to c in the shaded region.

Finally, the indifference condition $v_c = v_a$ yields $0 = p_A - 3p_B - 3p_C = 0$ or

$$p_A = 3p_B + 3p_C.$$

Setting $p_B = 0$ and using $p_A + p_C = 1$ we obtain the mixed strategy $(p_A, p_B, p_C) = (\frac{3}{4}, 0, \frac{1}{4})$. Setting $p_C = 0$ and using $p_A + p_B = 1$ we obtain the mixed strategy $(p_A, p_B, p_C) = (\frac{3}{4}, \frac{1}{4}, 0)$. The points $(0, \frac{1}{4})$ and $(\frac{1}{4}, 0)$ identify the indifference line in Figure 3 on the right below. Since against A the column player prefers a to c , a is preferred to c in the shaded region.




Figure 2




Figure 3

Figures 1 - 3 allow us to identify quickly against which strategies of the column player a pure strategy of the row player is a best response. a is better than b against strategies of the row player in the shaded region of Figure 1 while a is better than c against strategies in the shaded region of Figure 3. Overall a is better than b and c and hence a best response against strategies in the region of Figure 4 with the label a encircled. Similarly, b is better than a in the white region of Figure 1 and better than c in the shaded region of Figure 2. Overall b is better than both a and c and hence a best reply against strategies in the region of Figure 4 with the label b encircled. In the remaining region c must be a best response. Indeed, c is better than b in the white region of Figure 2 and better than a in the white region of Figure 3. Overall, it is a best reply against strategies in the intersection of the two regions, which is the region with label c encircled in Figure 4. Figure 4 is completed by labelling the side of the triangle opposite the vertex A with A encircled and by labelling the other two sides correspondingly.

Luckily the game is symmetric and hence no further derivations are required to obtain the best response diagram with the triangle representing the column player's strategy space. All we need to do is to replace p by q and to replace small letters by capital letter and vice versa. This follows because the expected payoffs u_A , u_B and u_C of the row player with the pure strategies A, B, C against a mixed strategy (q_a, q_b, q_c) are those of the column player after you have made the corresponding replacements.

Hence we obtain the diagram in Figure 5.




Figure 4




Figure 5

9 Nash Equilibria

A mixed strategy pair $(\mathbf{p}, \mathbf{q}) = ((p_A, p_B, p_C), (q_a, q_b, q_c))$ is a *Nash equilibrium* if each player acts optimal given his belief and if his belief about his opponent's behaviour is correct. If the column player believes that her opponent is using the strategy \mathbf{q} then playing the strategy \mathbf{p} is optimal for her (i.e., it is a *best reply / response* to \mathbf{q}) if and only if she expects not to gain a lower payoff with \mathbf{p} against \mathbf{q} than she could gain with any of her other strategies. Often her best response will not be unique and there will be many strategies yielding the same optimal payoff.

As remarked above, a mixed strategy \mathbf{p} will be a best response to \mathbf{q} if and only if all pure strategies used in \mathbf{p} with positive probability are themselves best responses to \mathbf{q} .

This observation gives us the following alternative characterization of a Nash equilibrium. A strategy pair (\mathbf{p}, \mathbf{q}) is a Nash equilibrium if and only if each player randomizes only between pure strategies that are best responses against the strategy of his opponent.

To illustrate this graphically fix a mixed strategy pair (\mathbf{p}, \mathbf{q}) and a pure strategy of one of the players, say strategy A of the row player. \mathbf{p} is represented by a point in the triangle representing the row player's mixed strategy set and \mathbf{q} is represented by a point in the triangle representing the column player's mixed strategy set. For the pair to be a Nash equilibrium it must be true that the row player is using A in his mixed strategy \mathbf{p} with positive probability only if A is a best response again \mathbf{q} . Now A is a best response to \mathbf{q} if \mathbf{q} is in the region with the encircled label A in Figure 5. (It may be on the boundary of that region.) We say then that label A is adjacent to \mathbf{q} . If A is not a best reply to \mathbf{q} then (\mathbf{p}, \mathbf{q}) can be a Nash equilibrium only if the row player is not using A with positive probability in her mixed strategy \mathbf{p} . In Figure 4 the mixed strategies where A is not used with positive probability are on the side of the triangle opposite the vertex A , i.e., on the line with the encircled label A adjacent. Since \mathbf{p} must be on this line we say that label A must be adjacent to \mathbf{p} .

Thus, for the strategy pair (\mathbf{p}, \mathbf{q}) to be Nash equilibrium label A must be adjacent to \mathbf{p} in Figure 4 or to \mathbf{q} in Figure 5 and the same must hold for the labels corresponding to the other five pure strategies B, C, a, b and c . If this condition holds, then indeed we have a Nash equilibrium.

This gives us finally the following method to find all Nash equilibria: Pick a point \mathbf{p} in the triangle of Figure 4 and write down the adjacent labels. Some of the six labels will not be adjacent. If you can find a point \mathbf{q} in the triangle of Figure 5 for which these missing labels are adjacent, the (\mathbf{p}, \mathbf{q}) is a Nash equilibrium.

Pick, for instance, the point $U = (\frac{1}{4}, 0, \frac{1}{4})$ in Figure 4. It has the labels a, B, c adjacent and hence we seek for a point in Figure 5 with the labels A, b, C adjacent. The only point with this property is $u = (\frac{3}{4}, 0, \frac{1}{4})$ and so $(U, u) = ((\frac{3}{4}, 0, \frac{1}{4}), (\frac{3}{4}, 0, \frac{1}{4}))$ is a Nash equilibrium.

The point $X = (\frac{3}{4}, \frac{1}{4}, 0)$ in Figure 4 has only the labels a, c, C adjacent. Since there is no point in Figure 5 with the labels A, B, b adjacent, X is not part of a Nash equilibrium.

Notice that there are at most three labels adjacent to any point in one of the two triangles of Figures 4 and 5. Since we need six distinct labels adjacent to a Nash equilibrium, a Nash equilibrium must be a pair of points with three labels adjacent to each point. This reduces our search to the strategies A, B, C, U, V, W, X for the row player and to the strategies a, b, c, u, v, w, x for the column player.

Hence the search proceeds as indicated in the following table:

strategy	adjacent labels	missing labels	fitting strategy
A	a, B, C	A, b, c	a
B	A, b, C	a, B, c	b
C	A, B, c	a, b, C	c
U	a, B, c	A, b, C	u
V	A, b, c	a, B, C	v
W	a, c, C	A, b, B	–
X	b, c, C	A, a, B	–

Thus, the game has five Nash equilibria, namely the three pure equilibria (A, a) , (B, b) and (C, c) and two partially mixed equilibria $(U, u) = ((\frac{3}{4}, 0, \frac{1}{4}), (\frac{3}{4}, 0, \frac{1}{4}))$ and $(V, v) = ((0, \frac{3}{4}, \frac{1}{4}), (0, \frac{3}{4}, \frac{1}{4}))$.