Game Theory
Lecture #12 – Mixed Nash Equilibria

Focus of Lecture:

- Mixed Strategies
- Best Response Sets
- Mixed Nash Equilibria

1 Introduction

Last lecture focused on investigating strategic decision-making in finite strategic form games. We introduced the famous solution concept of *Nash equilibrium*, which can be viewed as an action profile where all of the players are acting as contingent optimizers. In cases where a dominant strategy does not exist, we viewed the solution concept of Nash equilibria as a reasonable description of strategic behavior and focused on analyzing critical questions associated with this modeling choice pertaining to both the existence and uniqueness of Nash equilibria. One of the challenges associated with this descriptive modeling choice is that a Nash equilibrium need not exist in a given game, hence this modeling choice is incomplete and unsatisfactory. This chapter will focus on resolving this issue by shifting our attention from pure strategies to mixed strategies.

2 Strategic Form Games with Mixed Strategies

Recall the framework of strategic form games introduced in the last lecture. The specific components of a strategic form game are as follows:

- **Decision-makers:** There are a collection of decision-makers, i.e., $N = \{1, 2, 3, \ldots, |N|\}$.
- **Choice Sets:** Each decision-maker $i \in N$ is associated with a given choice set $\mathcal{A}_i$.
- **Joint Choice Sets:** The set of joint choices is defined by $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$. We will denote a joint choice by the tuple $a = (a_1, a_2, \ldots, a_n) \in \mathcal{A}$ where $a_i \in \mathcal{A}_i$ denotes the choice of player i. Lastly, we will often express a joint choice profile a by (a_i, a_{-i}) where $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$ encodes the choice of all decision-makers $\neq i$. The set of joint choices for all agents $\neq i$ is given by $\mathcal{A}_{-i} = \prod_{j \neq i} \mathcal{A}_j$
- **Utility Function:** Each decision-maker $i \in N$ is associated with a given utility function $U_i : \mathcal{A} \to \mathbb{R}$ that defines her preference over the joint actions $\mathcal{A}$.

This lecture we will focus on extending this framework from pure strategies, i.e., choices $a_i \in \mathcal{A}$, to mixed strategies, i.e., $p_i \in \Delta(\mathcal{A}_i)$. Specifically, we will consider the following extension to the framework of strategic form games defined above:

- **Mixed Strategies:** Each decision-maker $i \in N$ is now able to employ a probabilistic mixed strategy $p_i \in \Delta(\mathcal{A}_i)$, where $\Delta(\mathcal{A}_i)$ denotes the simplex over the finite set $\mathcal{A}_i$. Given a mixed strategy p_i, let $p_i^{a_i} \geq 0$ denote the probability that agent i selects action $a_i \in \mathcal{A}_i$ given the mixed strategy p_i. Accordingly, by definition we have $\sum_{a_i \in \mathcal{A}_i} p_i^{a_i} = 1$.
- **vNM Utility Functions:** Given the reliance on mixed strategies, it is imperative that we extend the agents' utility functions to account for the randomness in the agents' strategies. Accordingly, the goal is to transform the original utility function $U_i : \mathcal{A} \to \mathbb{R}$ to a extension of the form $U_i : \Delta(\mathcal{A}) \to \mathbb{R}$, which associates a payoff with each possible lottery over the joint action profiles $\mathcal{A}$. To that end, we consider the vNM (Von Neumann and Morgenstern) utility function where this extension is derived under the belief that (i) each agent $i \in N$ is making a decision independently of the other agents and (ii) the payoff associated with a given lottery is defined by an expectation. Accordingly, given a mixed strategy profile $p = (p_1, \ldots, p_n)$ where $p_i \in \Delta(\mathcal{A}_i)$ for each agent $i \in N$, we have that

$$U_i(p_1, \ldots, p_n) = \sum_{a \in \mathcal{A}} U_i(a) \times p_1^{a_1} \times \cdots \times p_n^{a_n}. \tag{1}$$

  We will also refer to vNM preferences as Bernoulli payoffs.

2.1 Potential Issues with vNM Preferences?

Note that vNM utility functions represent one way to extend the original utility function $U_i : \mathcal{A} \to \mathbb{R}$ to a new utility function of the form $U_i : \Delta(\mathcal{A}) \to \mathbb{R}$. While this seems to be the most natural extension, it is important to highlight concerns regarding the significance of the specific payoff values. In particular, in our original setting we utilized payoffs as a convenient way to encode preferences, i.e., $a \succ a'$ if and only if $U_i(a) > U_i(a')$. Note though that the magnitude of the utility played no role in the resulting preferences. For example, consider the following payoff matrices for the Prisoner's Dilemma game

	C	D
C	$2, 2$	$0, 3$
D	$3, 0$	$1, 1$

	C	D
C	$3, 3$	$0, 4$
D	$4, 0$	$1, 1$

Note that each of these payoff matrices encodes exactly the same preferences, e.g., $(D, C) \succ_{\text{row}} (C, C) \succ_{\text{row}} (D, D) \succ_{\text{row}} (C, D)$. However, these payoff matrices can give rise to different preference relations over lotteries as the following example shows.

Example 2.1 *Consider the following lotteries for the Prisoner's Dilemma games highlighted above*

	C	D
C	2/5	3/5
D	0	0

	C	D
C	0	0
D	0	1

Here, the value in each cell corresponds to the probability that the joint action will be employed in the given lottery, e.g., for the distribution on the left the joint choice (C, C) will be chosen with probability $2/5$ and (C, D) will be chosen with probability $3/5$. Focusing on ROW *evaluating preference between the two lotteries above, if we employ the left payoff matrix we have*

- *Left Distribution:* $U_{\text{ROW}} = (2/5)2 + (3/5)0 + (0)3 + (0)1 = 4/5$
- *Right Distribution:* $U_{\text{ROW}} = (0)2 + (0)0 + (0)3 + (1)1 = 1$

Hence, ROW *prefers the lottery on the right when employing the payoff matrix on the left (i.e., $1 > 4/5$). If we employ the right payoff matrix we have*

- *Left Distribution:* $U_{\text{ROW}} = (2/5)3 + (3/5)0 + (0)4 + (0)1 = 6/5$
- *Right Distribution:* $U_{\text{ROW}} = (0)3 + (0)0 + (0)4 + (1)1 = 1$

Hence, ROW *now prefers the lottery on the left when using the payoff matrix on the right (i.e., $6/5 > 1$).*

The above example demonstrates that the actual payoff values now take on heightened importance when using these values to determine preferences over lotteries. This is clearly problematic, but unfortunately we have minimal options as it is impossible to directly a specify a payoff for all possible lotteries explicitly without relying on some form of an extension, like an expectation. While the importance of specific payoff values is problematic, it turns out that even using the expectation operator is also problematic as shown in the following example.

Example 2.2 (Allais Paradox) *In this example we will seek to understand whether one can design payoff values such that a desired preference relation over lotteries can be realized when using expected payoffs. To that end, consider the following two lotteries which involve the possibility of winning \$10 million, \$2 million, or \$0 million:*

\$10	\$2	\$0
0	1	0

A

vs

\$10	\$2	\$0
0.1	0.89	0.01

a

In lottery A, the individual wins \$2 million with certainty, while in lottery a the individual wins \$10 million with probability 0.1, \$2 million with probability 0.89, and \$0 million with probability 0.01. It turns out that most people prefer lottery A to a. Now, consider two alternative lotteries of the form

\$10	\$2	\$0
0.1	0	0.9

B

vs

\$10	\$2	\$0
0	0.11	0.89

b

which can be interpreted in a similar fashion. Here, it turns out that most people prefer lottery B to b. Accordingly, are there utility values $U(10)$, $U(2)$, and $U(0)$ such that when using the expected value we would get the preference $A \succ a$ and $B \succ b$. Alternatively, what are the utility values that ensure:

$$\mathbf{E}[A] = 1 \cdot U(2) \geq 0.1 \cdot U(10) + 0.89 \cdot U(2) + 0.01 \cdot U(0) = \mathbf{E}[a]$$
$$\mathbf{E}[B] = 0.1 \cdot U(10) + 0.9 \cdot U(0) \geq 0.11 \cdot U(2) + 0.89 \cdot U(0) = \mathbf{E}[b]$$

Focusing on the first expression, if we add the term $x = 0.89 \cdot U(0) - 0.89 \cdot U(2)$ to both the left and right side we have

$$\begin{aligned} 1 \cdot U(2) + x &= 0.11 \cdot U(2) + 0.89 \cdot U(0) \\ &\geq 0.1 \cdot U(10) + 0.89 \cdot U(2) + 0.01 \cdot U(0) + x \\ &= 0.1 \cdot U(10) + 0.9 \cdot U(0) \end{aligned}$$

which implies that $\mathbf{E}[b] \geq \mathbf{E}[B]$. Hence, it is impossible to assign payoff values $U(10)$, $U(2)$, and $U(0)$ that ensure the preference relations $A \succ a$ and $B \succ b$ when using expectations.

The two examples above demonstrate potential problems when trying to extend utility functions to lotteries by considering expected values. However, the purpose of this exercise is to recognize that this approach has limitations which we need to fully appreciate so as to appropriately ground the forthcoming results. Nonetheless, we will continue to use vNM utilities for this purposes as any particular design of utility functions would suffer from a similar set of concerns.

2.2 Best response sets and mixed strategy Nash equilibria

Now that we have committed to employing vNM utility function, we will shift our attention to recasting our definition of best response sets and Nash equilibria in terms of mixed strategies. We start with the definition of a best response set defined as follows:

Definition 2.1 (Best Response) *The best response of player i to the collective strategy of the other players $\alpha_{-i} \in \Delta(\mathcal{A}_{-i})$ is of the form*

$$B_i(\alpha_{-i}) = \{\alpha_i : U_i(\alpha_i, \alpha_{-i}) \geq U_i(\alpha_i', \alpha_{-i}) \text{ for all } \alpha_i' \in \Delta(\mathcal{A}_i)\} \quad (2)$$

Once again, note that the best response function is actually a set.

Here, it is important to recognize that the output of this best response function is a set of mixed strategies that maximize the player's expected payoff. The following example highlights the computation of such a best response set.

Example 2.3 *Consider a generic two player / two action game with payoff matrix*

	L	R
T	a, A	b, B
B	c, C	d, D

Assume mixed strategies are $(p, 1-p)$ *for* ROW *and* $(q, 1-q)$ *for* COL*. Focusing on* ROW*, the best response to a given* COL *strategy* q *takes the form*

$$B_{\text{ROW}}(q) = \arg\max_{p \in [0,1]} \left(p\Big(q \cdot a + (1-q) \cdot b\Big) + (1-p)\Big(q \cdot c + (1-q) \cdot d\Big) \right),$$

which simplifies to

$$B_{\text{ROW}}(q) = \begin{cases} 1 & (q \cdot a + (1-q) \cdot b) > (q \cdot c + (1-q) \cdot d) \\ 0 & (q \cdot a + (1-q) \cdot b) < (q \cdot c + (1-q) \cdot d) \\ [0,1] & (q \cdot a + (1-q) \cdot b) = (q \cdot c + (1-q) \cdot d) \end{cases}$$

Note that when $(q \cdot a + (1-q) \cdot b) = (q \cdot c + (1-q) \cdot d)$*, then anything is a best response for* ROW*. One could employ a similar analysis to derive* $B_{\text{COL}}(p)$*.*

We are now ready to state the definition of Nash equilibria over mixed strategies.

Definition 2.2 (Mixed Strategy Nash equilibrium) *A mixed strategy profile* $\alpha^* = (\alpha_1^*, ..., \alpha_n^*)$ *is a mixed strategy Nash equilibrium if for every player* $i \in N$

$$\alpha_i^* \in B_i(\alpha_{-i}^*)$$

Recall that a Nash equilibrium did not necessarily exist in any game, e.g., matching pennies. The following famous result by John Nash proves that every game has a Nash equilibrium when considering mixed strategies.

Theorem 2.1 (Nash, 1950) *Every strategic form game with vNM preferences in which each player has finitely many actions has a mixed strategy Nash equilibrium.*

Nash's proof, while beyond the scope of this class, relies on a branch of mathematics pertaining to advanced fixed point theory. Here, the goal is to find a mixed strategy profile $(\alpha_1^*, ..., \alpha_n^*)$ such that

$$\alpha^* \to (B_1(\cdot), ..., B_n(\cdot)) \to \alpha^*$$

Such a profile is known as a *fixed point*, and fixed point theorems are often employed to argue the existence of such a profile. One simple illustration of a fixed point point theorem focuses on continuous functions. Suppose you are given a continuous function on the closed interval $[0,1]$ that is bounded between $[0,1]$. Does there exist a point $x \in [0,1]$ such that $x = f(x)$, or alternatively does there exists a point x that is a fixed point associated with the function or operation $f(\cdot)$. While not immediately obvious, fixed point theory can be employed to answer there question in the affirmative, i.e., there always exists an $x \in [0,1]$ such that $x = f(x)$ regardless of the continuous function f.

3 Illustrative Example – Hawk versus Dove

We conclude this section by looking at the classic game of Hawk versus Dove. This game is used is used to model aggressive versus passive behavior, such as in the game of "chicken" or driver behavior at a traffic intersection. Specifically, we have a two player game with payoff matrix

	H	D
H	$0,0$	$6,1$
D	$1,6$	$3,3$

where the action H, or hawk, represents aggressive behavior and D, or dove, represents passive behavior. We begin by characterizing the best response functions over pure strategies, which take the form

$$B_{\text{ROW}}(H) = D \qquad \& \qquad B_{\text{ROW}}(D) = H$$

and COL has a similar form. Accordingly, there are two Nash equilibria which are (H, D) and (D, H).

Are there additional mixed strategy Nash equilibria that are not represented by either (H, D) or (D, H)? To answer this question, we begin by deriving the best response function for *mixed* strategies. To that end, for ROW let $\mathbf{Pr}[H] = p$ and $\mathbf{Pr}[D] = 1-p$ and for COL let $\mathbf{Pr}[H] = q$ and $\mathbf{Pr}[D] = 1-q$. Given that the players select independently, the best response for ROW is

$$B_{\text{ROW}}(q) = \arg\max_{p\in[0,1]} \Big(p\Big(0\cdot q + 6\cdot(1-q)\Big) + (1-p)\Big(1\cdot q + 3\cdot(1-q)\Big)\Big) \tag{3}$$

which takes on the form

$$B_{\text{ROW}}(q) = \begin{cases} 1 & \Big(0\cdot q + 6\cdot(1-q)\Big) > \Big(1\cdot q + 3\cdot(1-q)\Big) \\ [0,1] & \Big(0\cdot q + 6\cdot(1-q)\Big) = \Big(1\cdot q + 3\cdot(1-q)\Big) \\ 0 & \Big(0\cdot q + 6\cdot(1-q)\Big) < \Big(1\cdot q + 3\cdot(1-q)\Big) \end{cases} = \begin{cases} 1 & q < 3/4 \\ [0,1] & q = 3/4 \\ 0 & q > 3/4 \end{cases}$$

Since the two players are symmetric we also have that

$$B_{\text{COL}}(p) = \begin{cases} 1 & p < 3/4 \\ [0,1] & p = 3/4 \\ 0 & p > 3/4 \end{cases}$$

Our goal is to determine a mixed strategy profile (p^*, q^*) such that $p^* \in B_{row}(q^*)$ and $q^* \in B_{col}(p^*)$. As done in the previous lecture, we plot the best response curves to aid us in this process

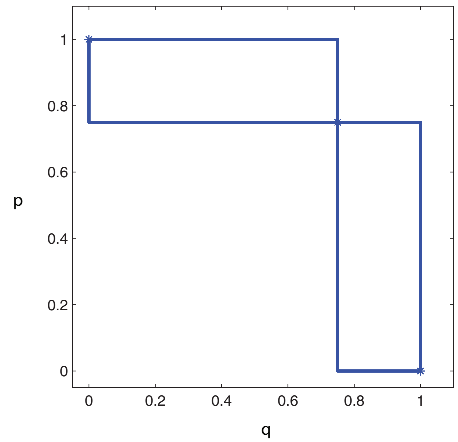

Recall that a Nash equilibrium occurs at an intersection of these best response plots. Note that there are three intersections on this plot, namely $(p = 1, q = 0)$, $(p = 0, q = 1)$, and $(p = 3/4, q = 3/4)$. The first two intersections highlighted above correspond to the pure Nash equilibria identified above, e.g., $(p = 1, q = 0)$ captures the pure strategy Nash equilibrium (H, D). The last intersection covers the new "mixed strategy" Nash equilibrium $(p^*, q^*) = (3/4, 3/4)$. Here, it is important to notice a key peculiarity about this mixed Nash equilibria – That is, both players are *indifferent*, i.e.,

$$B_{\text{ROW}}(3/4) = [0, 1] \;\;\&\;\; B_{\text{COL}}(3/4) = [0, 1]$$

i.e., at this mixed strategy Nash equilibria, the best response for either player is to play (H, D) with any probability combination. This is known as the indifference phenomena, and is true for any mixed Nash equilibria in any game.

4 Conclusion

This lecture covered one of the most famous results in game theory. In particular, we discussed John Nash's groundbreaking result which ensures the existence of a mixed strategy Nash equilibrium in any finite strategic form game when using vNM utility functions. Note that this guarantee is in stark contrast to the existence of pure strategy Nash equilibria, which might not exist in a given game. Accordingly, from an engineering perspective we can now rely on mixed strategy Nash equilibria as a reasonable description of societal behavior without having to worry about whether one exists.

Lastly, we analyzed the hawk dove game which seek to model strategic interaction involving either passive or aggressive play, such as chicken and driver interactions at an intersections. Are the pure Nash equilibrium desirable from a societal perspective? Is the mixed strategy Nash equilibria socially desirable? Are there other outcomes that could lead to more desirable behavior? If so, how could we justify such outcomes from a game theoretic perspective?

5 Exercises

1. Watch the movie A Beautiful Mind. Explain why the infamous bar scene is not an example of a Nash equilibrium.

2. Consider the following set of games

   - BoS:

	B	S
B	$2, 1$	$0, 0$
S	$0, 0$	$1, 2$

   - Stag hunt:

	Stag	Hare
Stag	$2, 2$	$0, 1$
Hare	$1, 0$	$1, 1$

   - Typewriter:

	Alt	Std
Alt	$3, 3$	$0, 0$
Std	$0, 0$	$1, 1$

   Answer the following questions for each game.

   (a) Determine the best response of the row player as a function of strategy of the column player $(q, 1-q)$.

   (b) For what value of q is the row player indifferent between its "top" action versus its "bottom" action?

   (c) Compute all Nash equilibria (both mixed and pure) for each game.

3. **Bribes:** Two inventors find themselves in a legal battle over a patent. The patent is worth 20 to each player, so the winner would receive 20 and the loser 0. Given the norms of the country, it is common to bribe the judge hearing a case. Each player can offer a bribe secretly, and the one whose bribe is the highest will be awarded the patent. If both choose not to bribe, or if the bribes are the same amount, then each has an equal chance of being awarded the patent. If a player does bribe, then they can either give a bribe of either 9 or 20. Any other number is considered very unlucky, and the judge would surely rule against a party who offered a different number.

   (a) Find the unique pure-strategy Nash equilibrium for this game.

   (b) If the norms were different, so that a bribe of 15 was also acceptable, is there a pure strategy Nash equilibrium?

   (c) Find the symmetric mixed-strategy Nash equilibrium for the game with the possible bribes of 9, 15, and 20.

   Notes: Not giving a bribe, i.e., giving a bribe of 0, is an option for parts (a) and (b). However, to simplify the analysis, it is not an option for part (c) to simplify the analysis. Also, when you give a bribe you lose that amount even if you are not awarded the patent.

6 Solutions

1. Watch the movie A Beautiful Mind. Explain was the infamous bar scene is not an example of a Nash equilibrium.

   In Nash's proposed joint action, each man chooses to dance with one of the blonde's friends. However this leaves the blonde without a dance partner, and one man could choose to *unilaterally-deviate* to dance with the blonde. The definition of a Nash equilibrium states that no player will be incentivized to unilaterally deviate, so this is not a Nash equilibrium.

2. Consider the following set of games

   - BoS:

	B	S
B	$2,1$	$0,0$
S	$0,0$	$1,2$

   - Stag hunt:

	Stag	Hare
Stag	$2,2$	$0,1$
Hare	$1,0$	$1,1$

   - Typewriter:

	Alt	Std
Alt	3, 3	0, 0
Std	0, 0	1, 1

Answer the following questions for each game.

(a) Determine the best response of the row player as a function of strategy of the column player $(q, 1-q)$.

(b) For what value of q is the row player indifferent between its "top" action versus its "bottom" action?

(c) Compute all Nash equilibria (both mixed and pure) for each game.

The following is for stag hunt. The other two are similar.

	Stag	Hare
Stag	2, 2	0, 1
Hare	1, 0	1, 1

- Suppose the row and column players uses a mixed strategies $(p, 1-p)$ and $(q, 1-q)$, respectively.
- The best response of the row player is:

$$B_{\text{ROW}}(q) = \begin{cases} 1 & 2q > q + (1-q) \\ [0,1] & 2q = q + (1-q) \\ 0 & 2q < q + (1-q) \end{cases}$$

  which can be rewritten as

$$B_{\text{ROW}}(q) = \begin{cases} 1 & q > 1/2 \\ [0,1] & q = 1/2 \\ 0 & q < 1/2 \end{cases}$$

- Because of the symmetry between players, the best response for the column player is

$$B_{\text{COL}}(p) = \begin{cases} 1 & p > 1/2 \\ [0,1] & p = 1/2 \\ 0 & p < 1/2 \end{cases}$$

- The resulting Nash equilibria in terms of (p^*, q^*) pairs are

$$(1,1), (0,0), \&(1/2, 1/2)$$

  This is illustrated in the best response plot:

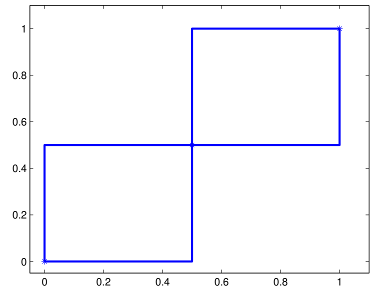

3. **Bribes:** Two players find themselves in a legal battle over a patent. The patent is worth 20 to each player, so the winner would receive 20 and the loser 0. Given the norms of the country, it is common to bribe the judge hearing a case. Each player can offer a bribe secretly, and the one whose bribe is the highest will be awarded the patent. If both choose not to bribe, or if the bribes are the same amount, then each has an equal chance of being awarded the patent. If a player does bribe, then the bribe can be valued at either 9 or 20. Any other number is considered very unlucky, and the judge would surely rule against a party who offered a different number.

   (a) Find the unique pure-strategy Nash equilibrium for this game.

   (b) If the norms were different, so that a bribe of 15 was also acceptable, is there a pure strategy Nash equilibrium?

   (c) Find the symmetric mixed-strategy Nash equilibrium for the game with the possible bribes of 9, 15, and 20.

   Notes: Not giving a bribe, i.e., giving a bribe of 0, is an option for parts (a) and (b). However, to simplify the analysis, it is not an option for part (c) to simplify the analysis. Also, when you give a bribe you lose that amount even if you are not awarded the patent.

   (a) In setting up the payoff matrix for this game, we note that when the players each bribe the same amount, they get an expected payoff of 10 minus their bid. The resulting payoff matrix is therefore

	0	9	20
0	$10, 10$	$0, 11$	$0, 0$
9	$11, 0$	$1, 1$	$-9, 0$
20	$0, 0$	$0, -9$	$-10, -10$

   Looking through each joint action, we can see that there is a Nash equilibrium at

$$(a_1 = 9, a_2 = 9).$$

   (b) If each player now has the option to bid 15, we need to add an extra action for each player to the payoff matrix:

	0	9	15	20
0	$10, 10$	$0, 11$	$0, 5$	$0, 0$
9	$11, 0$	$1, 1$	$-9, 15$	$-9, 0$
15	$5, 0$	$5, -9$	$-5, -5$	$-15, 0$
20	$0, 0$	$0, -9$	$0, -15$	$-10, -10$

With the addition of this new action, there is no longer a pure Nash equilibrium.

(c) For this part of the problem, we take out the option of bidding zero, and get the new payoff matrix

		9	15	20
(p)	9	$1, 1$	$-9, 15$	$-9, 0$
(q)	15	$5, -9$	$-5, -5$	$-15, 0$
(1-p-q)	20	$0, -9$	$0, -15$	$-10, -10$

To find a Nash equilibrium, we now look for a mixed strategy $(p, q, 1-p-q)$ which makes the other player indifferent over their actions, i.e.,

$$\begin{aligned} U_2(9,(p,q,1-p-q)) &= U_2(15,(p,q,1-p-q)) = U_2(20,(p,q,1-p-q)), \\ p(1)+q(-9)+(1-p-q)(-9) &= p(15)+q(-5)+(1-p-q)(-15) \\ &= p(0)+q(0)+(1-p-q)(-10) \end{aligned}$$

Solving this system of equations gives an equilibrium strategy of $(1/2, 1/10, 2/5)$.