Mixed Nash Equilibrium

IE 5545: Decision Analysis
Prof. Krishnamurthy Iyer*

May 2, 2026

We are now ready to extend the notion of Nash equilibrium to allow for mixed strategies. As we will see later, through this extension, we will be able to provide a solution/prediction even in games such as matching pennies for which our previous solution concepts have failed.

We begin with the following definition, exactly analogous to the definition of a pure Nash equilibrium, but uses mixed strategies in place of pure strategies:

Definition 1. *A mixed strategy profile $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ is a (mixed) Nash equilibrium if*

$$\pi_i(\sigma_i, \sigma_{-i}) \geq \pi_i(\sigma_i', \sigma_{-i}) \qquad \textit{for each player } i \textit{ and all } \sigma_i'.$$

In other words, in a mixed Nash equilibrium, each player plays a *best response* fixing all other players' strategies. This motivates us to extend the definition of the best response mapping to mixed strategies.

Definition 2. *The best response map $BR_i(\sigma_{-i})$ for a player i specifies the set of all mixed strategies that maximize the player's payoff when all other players play according to σ_{-i}:*

$$\begin{aligned} BR_i(\sigma_{-i}) &= \textit{Set of mixed strategies that maximize } \pi_i(\sigma_i, \sigma_{-i}) \textit{ over all } \sigma_i. \\ &= \{\sigma_i : \pi_i(\sigma_i, \sigma_{-i}) \geq \pi_i(\sigma_i', \sigma_{-i}) \textit{ for all } \sigma_i' \}. \end{aligned}$$

Observe that if a mixed strategy σ_i is a best response against σ_{-i}, then all pure strategy that are played with positive probability under σ_i must result in the same payoff to player i. This is because, otherwise, the player can increase her payoff by never choosing the pure strategy that yields strictly lower payoff, contradicting the fact that σ_i is a best response. Conversely, any mixed strategy σ_i that puts positive weight *only* on those pure strategies s_i that maximizes (across all pure strategies) her payoff against σ_{-i} must be a best response. Together, we can equivalently write $BR_i(\sigma_{-i})$ as

$$\begin{aligned} BR_i(\sigma_{-i}) &= \Delta\left(\begin{array}{c} \text{Set of pure strategies that maximize} \\ \pi_i(s_i, \sigma_{-i}) \text{ over all } s_i \end{array} \right), \\ &= \left\{\sigma_i : \text{if } \sigma_i(s_i) > 0, \text{ then } \pi_i(s_i, \sigma_{-i}) \geq \pi_i(s_i', \sigma_{-i}) \text{ for all } s_i' \right\} \end{aligned}$$

where $\Delta(A)$ of a set A denotes the set of all probability distributions on the set A.

Using the notion of best response mapping, we can equivalently define a mixed Nash equilibrium as follows:

Lemma 1. *A strategy $\sigma = (\sigma_1, \ldots, \sigma_N)$ is a mixed Nash equilibrium if and only if $\sigma_i \in BR_i(\sigma_{-i})$ for all i.*

Example (Matching pennies)**.** Consider the matching pennies game between two players, with payoffs as in the matrix below.

	H	T
H	$(1,-1)$	$(-1,1)$
T	$(-1,1)$	$(1,-1)$

It is easy to see that there is no pure Nash equilibrium in this game. We seek to find a mixed Nash equilibrium (if it exists). In other words, we want to find a strategy profile (σ_1, σ_2) such that $\sigma_1 \in BR_1(\sigma_2)$ and $\sigma_2 \in BR_2(\sigma_1)$. To do this, we will first compute the best response mapping for each player.

Denote the two players' strategies as follows:

$$\sigma_1 = \begin{bmatrix} p \\ 1-p \end{bmatrix} \qquad \text{and} \qquad \sigma_2 = \begin{bmatrix} q \\ 1-q \end{bmatrix} \qquad \text{for some } p, q \in [0,1].$$

To compute $BR_1(\sigma_2)$, we begin by finding player 1's payoff against σ_2 for each pure strategy. We have

$$\begin{aligned} \pi_1(H, \sigma_2) &= 2q - 1 \\ \pi_1(T, \sigma_2) &= 1 - 2q. \end{aligned}$$

Thus, (i) when $\pi_1(H,\sigma_2) > \pi_1(T,\sigma_2)$ (i.e., when $q > \frac{1}{2}$) then H is the unique best response for player 1; (ii) when $\pi_1(H,\sigma_2) < \pi_1(T,\sigma_2)$ (i.e., when $q < \frac{1}{2}$) then T is the unique best response for player 1; and (iii) when $\pi_1(H,\sigma_2) = \pi_1(T,\sigma_2)$ (i.e., when $q = \frac{1}{2}$) then any mixed strategy that mixes between H and T is a best response for player 1. Thus, we obtain

$$BR_1(\sigma_2) = \begin{cases} \{H\} = \left\{[1 \;\; 0]^\mathsf{T}\right\}, & \text{if } q > \frac{1}{2}, \\ \{T\} = \left\{[0 \;\; 1]^\mathsf{T}\right\}, & \text{if } q < \frac{1}{2}, \\ \Delta(\{H,T\}) = \left\{ \begin{bmatrix} p \\ 1-p \end{bmatrix} : p \in [0,1] \right\}, & \text{if } q = \frac{1}{2}. \end{cases}$$

Similarly, we have

$$\begin{aligned} \pi_2(H, \sigma_1) &= 1 - 2p \\ \pi_2(T, \sigma_1) &= 2p - 1, \end{aligned}$$

from which we conclude that

$$BR_2(\sigma_1) = \begin{cases} \{T\} = \left\{[0 \;\; 1]^\mathsf{T}\right\}, & \text{if } p > \frac{1}{2}, \\ \{H\} = \left\{[1 \;\; 0]^\mathsf{T}\right\}, & \text{if } p < \frac{1}{2}, \\ \Delta(\{H,T\}) = \left\{ \begin{bmatrix} q \\ 1-q \end{bmatrix} : q \in [0,1] \right\}, & \text{if } p = \frac{1}{2}. \end{cases}$$

From these sets, we observe that the only value of $(\sigma_1, \sigma_2$ for which we have $\sigma_1 \in BR_1(\sigma_2)$ and $\sigma_2 \in BR_2(\sigma_1)$ corresponds to $p = q = \frac{1}{2}$. Hence

$$(\sigma_1, \sigma_2) = \left(\begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}, \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix} \right)$$

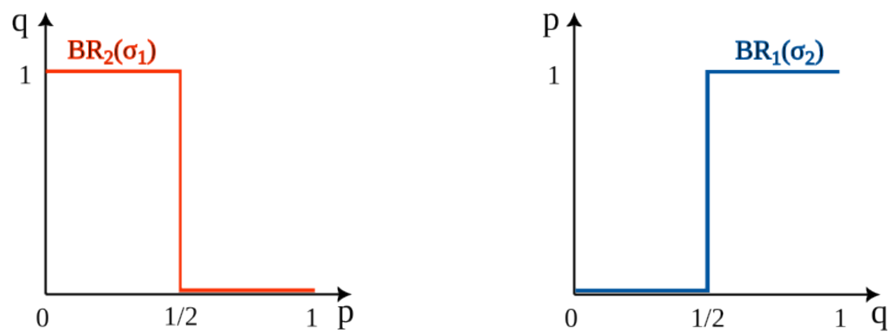

Figure 1: Best response maps for each player in matching pennies. (Notice that $BR_2(\sigma_1)$ is not a function as plotted, because for $p = \frac{1}{2}$, it takes multiple values. Similarly for $BR_1(\sigma_2)$.)

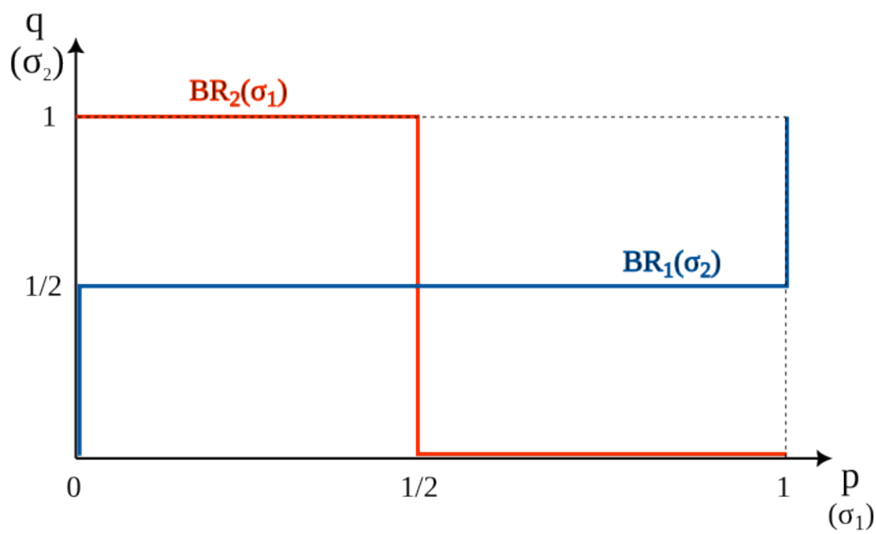

Figure 2: Best response maps for each player in matching pennies, shown on the same plot. Here, we have taken the earlier plot for $BR_1(\sigma_2)$, rotated it along the 45° diagonal (to swap the p and q axes), and superimposed it over the plot for $BR_2(\sigma_1)$.

is the unique mixed Nash Equilibrium of matching pennies.

Alternatively, we can find the NE (and see its uniqueness in this case) by plotting the best responses for the two players: we vary p (i.e., σ_1) and plot q for any $[q \quad 1-q]^\top \in BR_2(\sigma_1)$ (and vice versa). We obtain the plots shown in Figure 1. In Figure 2, we have shown these two maps on the same plot. An intersection (σ_1, σ_2) of these two maps satisfies $\sigma_1 \in BR_1(\sigma_2)$ and $\sigma_2 \in BR(\sigma_1)$, and thus yields a Nash equilibrium. From Figure 2, we see that there is a unique intersection at $p = q = \frac{1}{2}$, again yielding $([\frac{1}{2} \quad \frac{1}{2}]^\top, [\frac{1}{2} \quad \frac{1}{2}]^\top)$ as the unique NE.

Example (Co-ordination game)**.** Consider a coordination game where two players, Alice and Bob, must choose where to go out for the evening. The choices are a Salsa dancing workshop (Sa) or a Star Trek convention (ST). Suppose Alice prefers the Star Trek convention, whereas Bob prefers the Salsa dancing workshop, but both would prefer going to the same venue together over going to separate venues. A payoff matrix that represents these references is shown below, with Alice as the

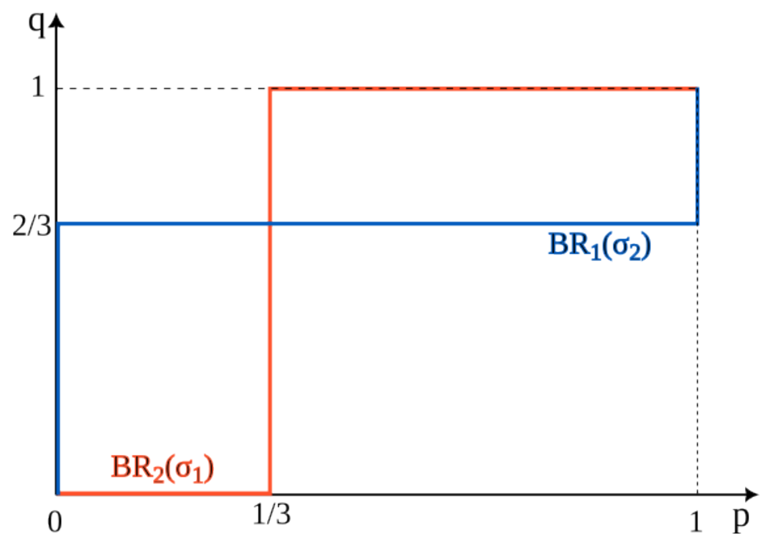

Figure 3: Best response maps $BR_1(\cdot)$ and $BR_2(\cdot)$ for the two players, shown on the same plot.

row player (player 1) and Bob as the column player (player 2).

	Sa	ST
Sa	$(1,2)$	$(0,0)$
ST	$(0,0)$	$(2,1)$

It is easy to see that this game has two pure Nash equilibria (ST, ST) and (Sa, Sa). (Since pure strategies are also mixed, these are also mixed Nash equilibria.) Are there additional mixed NE in this game? To answer this question, we begin by finding the best response maps of each player. As before, let $\sigma_1 = [p \;\; 1-p]^\top$ and $\sigma_2 = [q \;\; 1-q]^\top$. The payoffs of the two players are given by

$$\begin{aligned} \pi_1(Sa, \sigma_2) &= q & \pi_1(ST, \sigma_2) &= 2(1-q) \\ \pi_2(Sa, \sigma_1) &= 2p, & \pi_2(ST, \sigma_1) &= (1-p). \end{aligned}$$

From these expressions, we obtain the best response maps as

$$BR_1(\sigma_2) = \begin{cases} \left\{[0 \;\; 1]^\top\right\}, & \text{if } q < \frac{2}{3}, \\ \left\{[1 \;\; 0]^\top\right\}, & \text{if } q > \frac{2}{3}, \\ \left\{[p \;\; 1-p]^\top : p \in [0,1]\right\}, & \text{if } q = \frac{2}{3}, \end{cases}$$

and

$$BR_2(\sigma_1) = \begin{cases} \left\{[0 \;\; 1]^\top\right\}, & \text{if } p < \frac{1}{3}, \\ \left\{[1 \;\; 0]^\top\right\}, & \text{if } p > \frac{1}{3}, \\ \left\{[q \;\; 1-q]^\top : q \in [0,1]\right\}, & \text{if } p = \frac{1}{3}. \end{cases}$$

We show these best response maps on the same plot in Figure 3. We observe that the two best response curves have three intersections:

1. $p = 0, q = 0$: This corresponds to the pure NE where both go to the Star Trek convention.

2. $p = 1, q = 1$: This corresponds to the pure NE where both go to Salsa dancing workshop.

3. $p = \frac{1}{3}, q = \frac{2}{3}$: This is a *fully mixed* NE where both Alice & Bob randomize between the two choices.

We note that in the fully mixed NE, there is a positive probability that the two players end up at different venues. Furthermore, we see that in this NE each player chooses their preferred venue with 2/3 probability; this is the most likely outcome (Alice choosing ST and Bob choosing Sa) with probability 4/9. Finally, with 1/9 probability, the two players both end up choosing their less preferred option. Because of this, the fully mixed NE results in significantly reduced payoff than the two pure NEs in this game.

Existence of mixed Nash equilibrium

In the previous section, we analyzed two games, each with 2 players and 2 strategies. For both of them, we were able to find a mixed Nash equilibrium. Does it hold for other 2-player 2-strategy games? From our approach for analyzing these games, it is easy to see that this question is equivalent to asking whether for such games, the two best response curves will always intersect. By considering different geometric shapes the best response curves can take (there are only a few possible shapes), one can indeed show that these curves will always intersect, and hence such games always have (at least one) mixed Nash equilibrium. (Exercise: try implementing this proof approach.)

Does a mixed Nash equilibrium always exist in general games, with more players and/or more strategies? The answer is Yes! This is an important result established by John Nash Jr, in his seminal work introducing the concept of the equilibrium notion that bears his name. The result allowed for analyzing any finite game (and even some nice infinite games), and has been extremely influential, generating significant advancement (and development) of many different fields of economics and operations research.

While Nash's proof of existence is short, it uses advanced results from topology. Specifically, the proof uses a fixed point theorem to show a certain mapping has a fixed point. While the complete proof is beyond our scope, we will seek to understand it at a high level.

We begin with the notion of the *best response correspondence*[1] $BR(\cdot)$, which is a function that takes as input a mixed strategy profile, and gives as output a set of mixed strategy profiles. For any mixed strategy profile σ, the value of the best response correspondence $BR(\sigma)$ is the set of all mixed strategy profiles that the player could possibly deviate to, if each player i believes the other players will play according to σ_{-i} and she herself plays a best response. We have

$$BR(\sigma) = \{(\hat{\sigma}_1, \hat{\sigma}_2, \ldots, \hat{\sigma}_N) : \hat{\sigma}_i \in BR_i(\sigma_{-i}) \text{ for each } i\}.$$

The starting point of the proof is the observation that a Nash equilibrium corresponds to a *fixed point*[2] of the best response correspondence, i.e., a strategy profile σ satisfying $\sigma \in BR(\sigma)$.

Lemma 2. *A (mixed) strategy profile $\sigma = (\sigma_1, \ldots, \sigma_N)$ is a mixed Nash equilibrium if and only if $\sigma \in BR(\sigma)$; i.e. if and only if, σ is a fixed point of $BR(\cdot)$.*

Proof. A strategy profile σ is Nash equilibrium if and only if $\sigma_i \in BR_i(\sigma_{-i})$ for each i. The latter condition is just the condition for membership in the set $BR(\sigma)$, and hence is equivalent to requiring $\sigma \in BR(\sigma)$. $\square$

Thus, to show that a Nash equilibrium exists is equivalent to showing that the best response correspondence of any finite game always has a fixed point. Using the *Kakutani fixed point theorem*, Nash established the following result.

Theorem 1. *Suppose a game with finitely many players satisfies the following conditions:*

1. *the set of strategies Σ_i of each player is closed and bounded, and convex;*
2. *each player's payoff function $\pi_i(\sigma_1, \ldots, \sigma_N)$ is continuous in $\sigma = (\sigma_1, \ldots, \sigma_N)$*

[1] A correspondence f from $\mathcal{X}$ to $\mathcal{Y}$ is a set-valued function that maps each element in $x \in \mathcal{X}$ to a subset $f(x)$ of $\mathcal{Y}$. This is denoted by $f : \mathcal{X} \rightrightarrows \mathcal{Y}$. For example, the function $f(x) = [0, x]$ for all $x \in [0, 1]$ is a correspondence $f : [0, 1] \rightrightarrows [0, 1]$.

[2] A fixed point of a correspondence $f : \mathcal{X} \rightrightarrows \mathcal{X}$ is an element $x \in \mathcal{X}$ such that $x \in f(x)$.

3. *fixing the strategies σ_{-i} of other players, each player's payoff is concave*[3] *in their own strategy σ_i.*

Then, the best response correspondence $BR(\cdot)$ has a fixed point, and hence the game has a Nash equilibrium.

Let us verify that a finite game (i.e., a game with finitely many players each with finitely many pure strategies) satisfies the three conditions in the theorem.

1. For a finite game, each player's mixed strategy σ_i is a column vector whose entries are between 0 and 1 adding up to 1. This is clearly a closed and bounded set. Hence, the set of mixed strategy profiles is closed and bounded. Moreover, it is easy to check the set of mixed strategies Σ_i is convex, i.e., if $\sigma_i^1 \in \Sigma_i$ and $\sigma_i^2 \in \Sigma_i$ then $\alpha\sigma_i^1 + (1-\alpha)\sigma_i^2 \in \Sigma_i$ for all $\alpha \in [0,1]$.

2. In a finite game, the payoff of a player is given by

$$\pi_i(\sigma_1, \ldots, \sigma_N) = \sum_{s_1}\sum_{s_2}\cdots\sum_{s_N}\sigma_1(s_1)\sigma_2(s_2)\cdots\sigma_N(s_N)\,\pi_i(s_1, \ldots, s_N).$$

   This is a polynomial in $(\sigma_1, \ldots, \sigma_N)$ and hence continuous.

3. Further, fixing σ_{-i}, we have

$$\pi_i(\sigma_i, \sigma_{-i}) = \sum_{s_i}\sigma_i(s_i)\,\pi_i(s_i, \sigma_{-i}).$$

   This is a linear function in σ_i, and hence also concave.

Thus, in finite games, all three conditions hold, and hence the BR correspondence has a fixed point. This fixed point is a Nash equilibrium. Thus, we have established the following result.

Theorem 2 (Nash)**.** *All finite games have a mixed Nash equilibrium.*

The importance of this theorem is significant: it allows us to model interactions as games and to generate predictions of what will happen in a broad range of settings.

All concepts of equilibrium that are studied later are extensions of Nash equilibrium, with additional constraints and properties to select among (multiple) NE. Thus Nash equilibrium sits at the core of any such analysis.

One drawback of this result is that it is *non-constructive* — it shows the existence of a Nash equilibrium but does not provide a way to find it. In particular, we do not have an algorithm at the end of the proof that gives an NE as the output. It turns out that the problem of finding a Nash equilibrium in general games is quite hard, and is a central topic of active study in the field of *algorithmic game theory*.

For certain subclasses of games with additional structure, there exists simple computational procedures that lead to a Nash equilibrium. Some of the procedures have a natural *learning* interpretation, where one can think of the players as repeatedly playing the game and learning by observing past plays. We will look at some of these procedures next.

[3] A function $f : \mathbb{R}^n \to \mathbb{R}$ is concave if for any x and y we have

$$f(\alpha x + (1-\alpha)y) \geq \alpha f(x) + (1-\alpha)f(y) \qquad \text{for all } \alpha \in [0,1].$$

