

9

Mixed Strategies

The basic idea of Nash equilibria, that is, pairs of actions where each player is choosing *a particular one* of his possible actions, is an appealing one. Of course, we would like to put it to test: how will it fair with respect to our criteria of evaluating a solution concept?

Consider the following classic “zero sum” game called “Matching Pennies”.¹ Players 1 and 2 both put a penny on a table simultaneously. If the two pennies come up the same (heads or tails) then player 1 gets both, otherwise player 2 does. We can represent this in the following matrix:

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	1, -1

Clearly, the method we introduced above to find a pure strategy Nash equilibrium does not work – given a belief that player 1 has, he always wants to match it, and

¹A *zero sum game* is one in which the gains of one player are the losses of another, hence their payoffs sum to zero. The class of zero sum games was the main subject of analysis before Nash introduced his solution concept in the 1950's. These games have some very nice mathematical properties that we will not go into. An approachable reference to this older, yet mathematically appealing literature is ...

given a belief that player 2 has, he would like to choose the opposite orientation for his penny. Does this mean that a Nash equilibrium fails to exist? Not if we consider a richer set of possible behaviors, as we will soon see.

Matching pennies is not the only simple game that fails to have a pure-strategy Nash equilibrium. Consider the famous child's game "Rock-Paper-Scissors". Recall that rock beats scissors, scissors beats paper, and paper beats rock. If winning gives the player a payoff of 1 and the loser a payoff of -1 , we can describe this game by the following matrix:

	R	P	S
R	0,0	-1,1	1,-1
P	1,-1	0,0	-1,1
S	-1,1	1,-1	0,0

It is possible to write down, for example, the best response for player 1 when he believes that player 2 will play one of his pure strategies as:

$$s_1(s_2) = \begin{cases} P, & \text{when } s_2 = R \\ S, & \text{when } s_2 = P \\ R, & \text{when } s_2 = S \end{cases}$$

and a similar (symmetric) list would be the best response of player 2. What is the conclusion? There is no pure strategy equilibrium, just like in the matching pennies game.

9.1 Strategies and Beliefs

In what follows, we introduce the concept of *mixed strategies*, which captures the very plausible idea that there is uncertainty about the behavior of players, since they may choose to *randomly select* one of their pure strategies. This will give us flexibility in forming a much richer set of beliefs, and allowing players to choose from a richer set of strategies. This section will rely on you having some past experience with the ideas of randomness and probabilities over random events. That said, I will try to offer a decent amount of refreshing so that the notation and ideas will almost be self contained.

Definition 11 Let S_i be i 's pure strategy set. Define ΔS_i as the **simplex** of S_i or, the set of probability distributions over S_i . A **mixed strategy** for player i is an element $\sigma_i \in \Delta S_i$, so that σ_i is a probability distribution over S_i . We denote by $\sigma_i(s_i)$ the probability that player i plays s_i .

This definition of a mixed strategy can best be explained by a few simple examples. In the case of Matching Pennies, $S_i = \{H, T\}$, and the simplex of this two element strategy set can be represented by a single number $\sigma \in [0, 1]$, where σ is the probability that player i plays H , and $1 - \sigma$ is the probability that player i plays L . A more cumbersome, but maybe more revealing way of writing this can be as follows:

$$\Delta S_i = \{(\sigma(H), \sigma(T)) : \sigma(H) \geq 0, \sigma(T) \geq 0, \sigma(H) + \sigma(T) = 1\}.$$

The meaning of this notation is that the set of mixed strategies is the set of all pairs $(\sigma(H), \sigma(T))$ such that both are non-negative numbers, and they both sum up to one.² We use the notation $\sigma(H)$ to represent the probability that player i plays H , and $\sigma(T)$ as the probability that player i plays T .

In the example of Rock-Paper-Scissors, $S_i = \{R, P, S\}$ (for rock, paper and scissors respectively), and we can define

$$\Delta S_i = \{(\sigma(R), \sigma(P), \sigma(S)) : \sigma(R), \sigma(P), \sigma(S) \geq 0, \sigma(R) + \sigma(P) + \sigma(S) = 1\}$$

which is now three numbers, each defining the probability that the player plays one of his pure strategies.

Notice that a pure strategy is just a special case of a mixed strategy. For example, in the rock-paper-scissors game, we can represent the pure strategy of playing R with the mixed strategy, $\sigma(R) = 1, \sigma(P) = \sigma(S) = 0$.

In general, a finite probability distribution $p(\cdot)$ is a function on a finite *state space*. A finite state space has a finite number of *events*, and the probability distribution function $p(\cdot)$ determines the probability that any particular event in the

²This follows from the definition of a probability distribution over a two element set. This is why a single number p is enough to define a probability distribution over two actions. In general, the simplex of a k -dimensional strategy set (one with k pure strategies) will be in a $k - 1$ dimensional spaces, where each of the $k - 1$ numbers is in $[0, 1]$, and represent the probability of the first $k - 1$ pure strategies. All sum up to a number equal to or less than one so that the remainder is the probability of the k -th pure strategy.

state space occurs. For example, the state space may be the possible weather conditions, say, “rain” (R) “hail” (H) or “clear” (C), so the space is $S = \{R, H, C\}$. The probability of each event, or *state of nature* will be given by the associated probability function. In our case of a game, the finite state space is the finite strategy set of a player S_i .

Any probability distribution over a finite state space, in our case S_i , must satisfies two conditions:³

1. $p(s_i) \geq 0$ for all $s_i \in S_i$, and
2. $\sum_{s_i \in S_i} p(s_i) = 1$.

That is, the probability of every any event happening must be non negative, and the sum of the probabilities of all the possible events must add up to one.

If we have such a finite state space with a probability distribution $p(\cdot)$, we will say that an event is *in the support of* $p(\cdot)$ if it occurs with positive probability. For example, for the game of rock-paper-scissors, a player can choose rock or paper, each with equal probability, so that $\sigma(R) = \sigma(P) = 0.5$ and $\sigma(S) = 0$. We will then say that R and P are in the support of $\sigma(\cdot)$, but S is not.

As we have seen with the Cournot and Bertrand duopoly examples, the strategy spaces need not be finite. In this case a mixed strategy will be given by a *cumulative distribution function* (CDF). A CDF replaces the role of a probability distribution function as a representation of randomness over a *continuous* state space.

For example, assume that there is a capacity constraint of 100 units of production in the Cournot duopoly game. Then the strategy space for each firm is $S_i = [0, 100]$, which will represent our continuous state space. In this case, we define a cumulative distribution function, $F : [0, 100] \rightarrow [0, 1]$ as follows,

$$F(x) = \text{Probability}\{s_i < x\}$$

that is, $F(x)$ would represent the probability that the *actual* chosen quantity is less than x .

³The notation $\sum_{s_i \in S_i} p(s_i)$ means the sum of $p(s_i)$ over all the $s_i \in S_i$.

If the CDF $F(\cdot)$ has a derivative $f(\cdot) = F'(\cdot)$, then for the case where $x \in [0, 100]$ we can rewrite the CDF as $F(x) = \int_0^x f(x)dx$. The derivative of a CDF is called a *density function*, and in loose terms represents the likelihood of an event.

When we have such a continuum of choices, then a mixed strategy will be represented by either a CDF $F(\cdot)$ or its density function $f(\cdot)$. We will say that a particular action is *in the support* of such a continuous mixed strategy if the density at that action is positive, that is, x is in the support of $F(\cdot)$ if $f(x) > 0$. This is analogous to the definition of a support for a finite mixed strategy.⁴

Treating a player's strategy set as such a state space, and introducing probability distributions over the strategy sets, lets us enrich the actions that players can take. For example, in the matching pennies game, a player is no longer restricted to *choose* between heads or tails, but he can flip the coin and thus choose a probability distribution over heads and tails. If the coin is a "fair" one, then by flipping it the probability of each side coming up is 0.5. However, using some play-clay on different sides of the coin the player can possibly choose other distributions.⁵

Interestingly, introducing probability distributions not only enriches the set of actions that a player can choose from pure to mixed strategies, but allows us to enrich the beliefs that players have in a natural way. Consider, for example, player i who plays against opponents $-i$. It may be that player i is uncertain about the behavior of his opponents for many reasons. For example, he may believe that his opponents are indeed choosing mixed strategies, which immediately implies that their behavior is not fixed but rather random. Maybe, more convincing is the case where player i is playing a game against an opponent that he does not know, and his opponent's background will determine how she will play. This interpretation will be revisited later, and is a very appealing justification for beliefs that are random, and behavior that justifies these beliefs. We introduce this idea formally

⁴The reason this is indeed loose is that for a continuous random variable, the probability that any particular event will occur is zero. If these terms are foreign, a short pause to read up on probability would be quite necessary. A good reference, for example, is...

⁵There will be many other ways to choose other distributions. For example, with a fair coin you can toss the coin twice, thus creating four events, (HH, HT, TH, TT) , where each event (combination of two draws) occurs with probability 0.25. If you choose Heads if it ever appears, then you choose heads with probability 0.75 because heads will come up at least once in three of the four events. You can even be more creative.

Definition 12 A *belief* for player i is given by a probability distribution $\pi_i \in \Delta S_{-i}$ over the strategies of his opponents.

Thus, a *belief* for player i is a probability distribution over the *actions* of his opponents. Notice that the belief of player i lies in the same set that represents the profiles of mixed strategies of player i 's opponents. For example, in the rock-paper-scissors game, we can represent the beliefs of player 1 as a triplet, $(\pi_1(R), \pi_1(P), \pi_1(S))$ where by definition, $\pi_1(R), \pi_1(P), \pi_1(S) \geq 0$, and $\pi_1(R) + \pi_1(P) + \pi_1(S) = 1$. The interpretation of $\pi_1(s_2)$ is the probability that player 1 assigns to player 2 playing some particular $s_2 \in S_2$. Recall that the strategy of player 2 is a triplet $\sigma_2(R), \sigma_2(P), \sigma_2(S) \geq 0$, with $\sigma_2(R) + \sigma_2(P) + \sigma_2(S) = 1$, so we can clearly see the analogy between π and σ .

This completes the enrichment of both the strategies that players can choose, and the corresponding beliefs that they will hold about their opponents' behavior. Now we have to address the issue of how players will evaluate their payoffs from playing a mixed strategy, as well as the effect of random behavior of their opponents on their own payoffs. For this, we resort to the well established ideas of utility over random payoffs.

9.2 Expected Utility

Consider the matching pennies game described above, and assume for the moment that player 2 chose to randomize when she plays. In particular, assume she has a device that plays H with probability $\frac{1}{3}$ and T with probability $\frac{2}{3}$. If player 1 plays H for sure, he will win with probability $\frac{1}{3}$ while he will lose with probability $\frac{2}{3}$. If, however, he plays T then with probability $\frac{2}{3}$ he will win and with probability $\frac{1}{3}$ he will lose. It is quite intuitive that the second option, playing T , should be more attractive. But how much should each option be "worth" in terms of payoffs for player 1?

If we think of the choice of H or T for player 1, it is like choosing between two *lotteries*, where a lottery is exactly a random payoff. Indeed, if we think of the outcomes of games as the payoffs that they imply, then when some players are playing mixed strategies, the players' payoffs are *random variables*. When the

lotteries are as simple as the one we just described, then the choice seems obvious: winning with a higher probability must be better. However, when the lotteries are not as simple, say winning and losing have different final payoffs in the different lotteries, then a more thoughtful approach is needed.

There is a well developed methodology to evaluate how much a lottery is worth for a player, and how different lotteries compare to others, and to “sure” payoffs. This methodology, evaluating *expected utility*, is the framework we will adopt and use throughout this text. First, we introduce the following definition:

Definition 13 Let $u_i(x)$ be player i 's utility function over monetary payoffs $x \in X$. If $X = \{x_1, x_2, \dots, x_L\}$ is finite, and $p = (p_1, p_2, \dots, p_L)$ is a **lottery over monetary payoffs** such that $p_k = \Pr\{x = x_k\}$ then we define i 's **expected utility** as

$$E[u(x)] = \sum_{k=1}^L p_k \cdot u(x_k) .$$

If X is a continuous space with a lottery given by the cumulative distribution $F(x)$, with density $f(x) = F'(x)$, then we define i 's **expected utility** as:⁶

$$E[u(x)] = \int_{x \in X} u(x) f(x) dx$$

The idea of expected utility is quite intuitive: if we interpret a lottery as a list of “weights” on monetary values, so that numbers that appear with higher probability have more weight, then the expected utility of that lottery is nothing other than the weighted average of utilities for each realization of the lottery. It turns out that there are some important foundations that make such a definition valid, and they were developed by John von-Neuman and Oscar Morgenstern, two of the founding fathers of game theory. These foundations are beyond the scope of this text, and a nice treatment of the subject appears in Kreps (1990), chapter 3.

What does it mean for player i to have a *utility function over monetary payoffs*? This can be best explained with an example: Suppose that a player i has a utility from money that is increasing in the amount of money he has, but at a diminishing

⁶More generally, if there are continuous distributions that do not have a density because $F(\cdot)$ is not differentiable, then the expected utility is given by $E[u(x)] = \int_{x \in X} u(x) dF(x)$.

rate. Arguably, this is a rather reasonable assumption which means that he values the first thousand dollars more than the second thousand, and so on. Suppose, for example, that the utility is given by the function $u(x) = 2\sqrt{x}$, which is increasing ($u'(x) = \frac{1}{\sqrt{x}} > 0$) and concave ($u''(x) = -\frac{1}{2x\sqrt{x}} < 0$).⁷ The concavity of $u(x)$ represents the diminishing marginal value of money: the increase of utility is smaller for each additional dollar.

Now assume that this player is offered a lottery X that can take one of two values, $x \in \{4, 9\}$, with $\Pr\{x = 4\} = \frac{1}{3}$, that is, with probability $\frac{1}{3}$ the outcome is 4, while with probability $\frac{2}{3}$ it is 9. The expected utility of this agent is then, by definition,

$$E[u(x)] = \frac{1}{3} \cdot 2\sqrt{4} + \frac{2}{3} \cdot 2\sqrt{9} = 5\frac{1}{3}.$$

We now ask, what is the most that this player would be willing to pay for this lottery? The answer must come from the utility of the player: how much money, c , gives him a utility of $5\frac{1}{3}$? The answer comes from solving the equation $2\sqrt{c} = 5\frac{1}{3}$, which yields $c = 7\frac{1}{9}$. That is, if this player follows the rules of expected utility as we defined them above, then he *should be indifferent* between a sure amount of $7\frac{1}{9}$ and the proposed lottery X .⁸

Notice that this sure amount is smaller than the *expected amount of money* that this lottery promises, which is $E(x) = \frac{1}{3} \cdot 4 + \frac{2}{3} \cdot 9 = 7\frac{1}{3}$. The fact that the sure amount c is smaller than the expected payment of X is a consequence of the shape of $u(\cdot)$, which in this case is concave. The concavity of $u(\cdot)$ implies that for any lottery, the player would be willing to accept *less than the mean* of the lottery instead of the lottery, a property we know as *risk aversion*.⁹ This implies that to be able to say how much a lottery is worth for a player, we need to know something about his risk preferences.

⁷A function $f(x)$ is said to be *concave* over an interval (a, b) if for every pair $x, y \in (a, b)$ and every $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)y) \geq \lambda f(x) + (1 - \lambda)f(y)$$

and is *strictly concave* if for every $\lambda \in (0, 1)$ this inequality is strict.

⁸In the economics literature this certain amount c is called the *certainty equivalent* to the lottery X . Think of it as the highest amount the player will pay to get the lottery X , or the lowest amount the player will accept to give X up.

⁹This property is often referred to as Jensen's inequality. In its most general formalization, it says that if $f(\cdot)$ is concave and X is a random variable, then $E[f(X)] \leq f(E[X])$, where $E[\cdot]$ is the expectations operator.

However, we will circumvent this problem in a rather nifty way. When we write down a game, we do not write down the payoffs in their monetary terms, but rather we write them in some abstract way as “utils”. Namely, in the example above replace the payoffs of the lottery $\{x_1, x_2\} = \{4, 9\}$ with the associated utils $\{u(x_1), u(x_2)\} = \{4, 6\}$, so that the expected utility is immediately given by its definition of the weighted sum of utils,

$$E[u] = \frac{1}{3} \cdot 4 + \frac{2}{3} \cdot 6 = 5\frac{1}{3}.$$

This implies that when we have the payoffs of a game determined in utils, we can calculate his expected utility from a lottery over outcomes as the expectation over his payoffs from the outcomes. That is, the payoffs of a game are *not* defined as monetary payoffs but instead are defined as the utility values from the outcomes of play. This allows us to define the expected payoff of a player from mixed strategies as follows:

Definition 14 *The **expected payoff** of player i when he chooses $s_i \in S_i$ and his opponents play the mixed strategy $\sigma_{-i} \in \Delta S_{-i}$ is*

$$u_i(s_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} \overbrace{\sigma_{-i}(s_{-i})}^{\text{probability}} \cdot \overbrace{u_i(s_i, s_{-i})}^{\text{utils}}$$

Similarly, the expected payoff of player i when he chooses $\sigma_i \in \Delta S_i$ and his opponents play the mixed strategy $\sigma_{-i} \in \Delta S_{-i}$ is

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) \cdot u_i(s_i, \sigma_{-i}) = \sum_{s_i \in S_i} \left(\sum_{s_{-i} \in S_{-i}} \sigma_i(s_i) \cdot \sigma_{-i}(s_{-i}) \cdot u_i(s_i, s_{-i}) \right)$$

Example: Rock-Paper-Scissors

Recall the rock-paper-scissors example above,

	R	P	S
R	0, 0	−1, 1	1, −1
P	1, −1	0, 0	−1, 1
S	−1, 1	1, −1	0, 0

and assume that player 2 plays: $\sigma_2(R) = \sigma_2(P) = \frac{1}{2}$; $\sigma_2(S) = 0$. We can now calculate the expected payoff for player 1 from any of his pure strategies,

$$\begin{aligned} u_1(R, \sigma_2) &= \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot (-1) + 0 \cdot 1 = -\frac{1}{2} \\ u_1(P, \sigma_2) &= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 + 0 \cdot (-1) = \frac{1}{2} \\ u_1(S, \sigma_2) &= \frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot 1 + 0 \cdot 0 = 0 \end{aligned}$$

It is easy to see that player 1 has a unique best response to this mixed strategy of player 2: if he plays P , he wins or ties with equal probability, while his other two pure strategies are worse: with R he either loses or ties and with S he either loses or wins.

Example: Bidding for a Dollar

Imagine the following game in which two players can bid for a dollar. Each can submit a bid that is a real number, so that $S_i = [0, \infty)$, $i \in \{1, 2\}$. The person with the highest bid gets the dollar, but the twist is that both bidders have to pay their bids. If there is a tie then both pay and the dollar is awarded to each player with an equal probability of 0.5. Thus, if player i bids s_i and player $j \neq i$ bids s_j then player i 's payoff is

$$u_i(s_i, s_{-i}) = \begin{cases} -s_i & \text{if } s_i < s_j \\ \frac{1}{2} - s_i & \text{if } s_i = s_j \\ 1 - s_i & \text{if } s_i > s_j \end{cases}.$$

Now imagine that player 2 is playing a mixed strategy in which he is *uniformly* choosing a bid between 0 and 1. That is, player 2's mixed strategy σ_2 is a uniform random variable between 0 and 1, which has a CDF $F(x) = x$ and a density $f(x) = 1$ for all $x \in [0, 1]$. The expected payoff of player 1 from offering a bid $s_i > 1$ is $1 - s_i < 0$ since he will win for sure, but this would not be wise. The

expected payoff from bidding $s_i < 1$ is

$$\begin{aligned}
 Eu_1(s_1, \sigma_2) &= \Pr\{s_1 < s_2\} \cdot (-s_1) + \Pr\{s_1 = s_2\} \cdot \left(\frac{1}{2} - s_1\right) + \Pr\{s_1 > s_2\} \cdot (1 - s_1) \\
 &= (1 - F(s_1)) \cdot (-s_1) + 0 \cdot \left(\frac{1}{2} - s_1\right) + F(s_1) \cdot (1 - s_1) \\
 &= 0
 \end{aligned}$$

Thus, when player 2 is using a uniform distribution between 0 and 1 for his bid, then player 1 cannot get any positive *expected* payoff from any bid he offers: any bid less than one offers an expected payoff of 0, and any bid above 1 guarantees getting the dollar at an inflated price. This cute game is one to which we will return later, since it has several interesting features and twists.

9.3 Mixed Strategy Nash Equilibrium

Now that we are equipped with a richer strategy space for our players, we are ready to restate the definition of a Nash equilibrium—which we have already introduced in pure strategies—to this more general setup where players can choose mixed strategies:

Definition 15 *The mixed strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$ is a Nash Equilibrium if for each player σ_i^* is a best response to σ_{-i}^* . That is, for all $i \in N$,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \quad \forall \sigma_i \in \Delta S_i.$$

This definition is the natural generalization of what we defined previously. We require that each player is choosing a strategy $\sigma_i^* \in \Delta S_i$ that is the best thing he can do when his opponents are choosing some profile $\sigma_{-i}^* \in \Delta S_{-i}$.

There is another interesting interpretation of the definition of Nash equilibrium. We can think of σ_{-i}^* as the belief of player i about his opponents, π_i , which captures the idea that player i is uncertain of his opponents' behavior. The profile of mixed strategies σ_{-i}^* thus captures this uncertain belief over all of the pure strategies that player i 's opponents can play. Clearly, rationality requires that a player play a best response given his beliefs (which now extends the notion of rationalizability

to allow for uncertain beliefs). A Nash equilibrium requires that these beliefs be correct.

Turning back to the explicit interpretation of a player actually mixing between several pure strategies, the following definition is useful:

Definition 16 *Let σ_i be a mixed strategy played by player i . We say that the pure strategy $s_i \in S_i$ is in the support of σ_i if $\sigma_i(s_i) > 0$, that is, if s_i is played with positive probability.*

That is, we will say that a pure strategy s_i is in the support of σ_i if when playing σ_i , player i chooses s_i with some positive probability $\sigma_i(s_i) > 0$.

Now imagine that in the Nash equilibrium profile σ^* , the support of i 's mixed strategy σ_i^* contains more than one pure strategy, say s_i and s'_i are both in the support of σ_i^* . What must we conclude about player i if σ_i^* is indeed part of a Nash equilibrium $(\sigma_i^*, \sigma_{-i}^*)$? By definition, σ_i^* is a best response against σ_{-i}^* , which means that given σ_{-i}^* , player i cannot do better than to randomize between more than one of his pure strategies, in this case, s_i and s'_i . But, when would a player be willing to randomize between two alternative pure strategies?

Proposition 7 *If σ^* is a Nash equilibrium, and both s_i and s'_i are in the support of σ_i^* , then*

$$u_i(s_i, \sigma_{-i}^*) = u_i(s'_i, \sigma_{-i}^*) = u_i(\sigma_i^*, \sigma_{-i}^*).$$

The proof is quite straightforward and follows from the observation that if a player is randomizing between two alternatives, then he must be indifferent between them. If this were not the case, say $u_i(s_i, \sigma_{-i}^*) > u_i(s'_i, \sigma_{-i}^*)$ with both s_i and s'_i in the support of σ_i^* , then by reducing the probability of playing s'_i from $\sigma_i^*(s'_i)$ to zero, and increasing the probability of playing s_i from $\sigma_i^*(s_i)$ to $\sigma_i^*(s_i) + \sigma_i^*(s'_i)$, player i 's expected utility must go up, implying that σ_i^* could not have been a best response to σ_{-i}^* .

This simple observation will play an important role for finding mixed strategy Nash equilibria. In particular, we know that if a player is playing a mixed strategy, he must be indifferent between the actions he is choosing with positive probability, that is, the actions that are in the support of his mixed strategy. One player's indifference will impose restrictions on the behavior of other players, and these

restrictions will help us find the mixed strategy Nash equilibrium. An example will be helpful.

Example: Matching Pennies

Consider the matching pennies game,

	H	T
H	1,-1	-1, 1
T	-1, 1	1,-1

and recall that we showed that this game does not have a pure strategy Nash equilibrium. We now ask, does it have a mixed strategy Nash equilibrium? To answer this, we have to find mixed strategies for both players that are mutual best responses.

Define mixed strategies for players 1 and 2 as follows: Let p be the probability that player 1 plays H and $1 - p$ the probability that he plays T . Similarly, let q be the probability that player 2 plays H and $1 - q$ the probability that he plays T .

Using the formulae for expected utility in this game, we can write player 1's expected utility from each of his two pure actions as follows:

$$\begin{aligned} u_1(H, q) &= q \cdot (1) + (1 - q) \cdot (-1) = 2q - 1 \\ u_1(T, q) &= q \cdot (-1) + (1 - q) \cdot 1 = 1 - 2q \end{aligned} \tag{9.1}$$

With these equalities in hand, we can calculate the best response of player 1 for any choice q of player 2. In particular, playing H will be strictly better than playing T for player 1 if and only if $u_1(H, q) > u_1(T, q)$, and using (9.1) above this will be true if and only if


$$2q - 1 > 1 - 2q,$$

which is equivalent to $q > \frac{1}{2}$. Similarly, playing T will be strictly better than playing H for player 1 if and only if $q < \frac{1}{2}$. Finally, when $q = \frac{1}{2}$ player 1 will be

indifferent between playing H or T . This simple analysis gives us the best response correspondence¹⁰ of player 1, which is:

$$BR_1(q) = \begin{cases} p = 0 & \text{if } q < \frac{1}{2} \\ p \in [0, 1] & \text{if } q = \frac{1}{2} \\ p = 1 & \text{if } q > \frac{1}{2} \end{cases}.$$

It may be insightful to graph the expected utility of player 1 from choosing either H or T as a function of q , the choice of player 2, as shown in figure 2.2.



Expected utility in the Matching Pennies game

The expected utility of player 1 from playing H was given by the function $u_1(H, q) = 2q - 1$ as described in (9.1) above. This is the rising linear function in the figure. Similarly, $u_1(T, q) = 1 - 2q$ is the declining function. Now, it is easy to see where the best response of player 1 is coming from. The “upper envelope” of the graph will show the highest utility that player 1 can achieve when player 2 plays q . When $q < \frac{1}{2}$ this is achieved by playing T , when $q > \frac{1}{2}$ this is achieved by playing H , and when $q = \frac{1}{2}$ both H and T are equally good for player one.

In a similar way we can calculate the utilities of player 2 given a mixed strategy p of player 1 to be,

¹⁰Unlike a function, which maps each value from the domain onto a unique value in the range, a correspondence can map values in the domain to *several* values in the range. Here, for example, when the domain value is $q = \frac{1}{2}$, then the correspondence maps this into a set, $BR(\frac{1}{2}) = [0, 1]$.

$$\begin{aligned} u_2(p, H) &= p \cdot (-1) + (1 - p) \cdot 1 = 1 - 2p \\ u_2(p, T) &= p \cdot 1 + (1 - p) \cdot (-1) = 2p - 1 \end{aligned}$$

and this implies that 2's best response is,

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < \frac{1}{2} \\ q \in [0, 1] & \text{if } p = \frac{1}{2} \\ q = 0 & \text{if } p > \frac{1}{2} \end{cases} .$$


We know from the proposition above that when player 1 is mixing between H and T , both with positive probability, then his payoff from H and from T must be identical. This, it turns out, imposes a restriction on the behavior of player 2, given by the choice of q . Namely, player 1 is willing to mix between H and T if and only if

$$u_1(H, q) = u_1(T, q) ,$$

which is true, from our analysis above, only when $q = \frac{1}{2}$. This is the way in which the indifference of player 1 imposes a restriction on player 2: only when player 2 is playing $q = \frac{1}{2}$, will player 1 be willing to mix between his actions H and T . Similarly, player 2 is willing to mix between H and T only when $u_2(p, H) = u_2(p, T)$, which is true only when $p = \frac{1}{2}$.

At this stage we have come to the conclusion of our quest for a Nash equilibrium in this game. We can see that there is indeed a pair of mixed strategies that form a Nash Equilibrium, and these are precisely when $(p, q) = (\frac{1}{2}, \frac{1}{2})$.

We return now to observe that the best response correspondence for each player is a function the other player's strategy, which in this case is a probability between 0 and 1, namely, the opponent's mixed strategy (probability of playing H). Thus, we can graph the best response correspondence of each player in a similar way that we did for the Cournot duopoly game since each strategy belongs to a well defined interval, $[0, 1]$. For the matching pennies example, player 2's best response $q(p)$ can be graphed in figure X.X. ($q(p)$ is on the x -axis, as a function of p on the y -axis.) Similarly, we can graph player 1's best response, $p(q)$, and these two correspondences will indeed intersect only at one point: $(p, q) = (\frac{1}{2}, \frac{1}{2})$. By definition, when these two correspondences meet, the point of intersection is a Nash equilibrium.



Best Response correspondences in the Matching Pennies game

There is a simple logic that we can derive from the Matching Pennies example that is behind the general way of finding mixed strategy equilibria in games. The logic relies on a fact that we had already discussed, which says that if a player is mixing between several strategies then he *must be indifferent* between them. What a particular player i is willing to do depends on the strategies of his opponents. Therefore, to find out when player i is willing to mix between some of his pure strategies, we must find strategies of his opponents, $-i$, that make him indifferent between some of his pure actions.

For the matching pennies game this can be easily illustrated as follows. First, we ask *which strategy of player 2 will make player 1 indifferent* between playing H and T . The answer to this question (assuming it is unique) *must be player 2's strategy in equilibrium*. The reason is simple: if player 1 is to mix in equilibrium, then player 2 must be playing a strategy for which player 1's best response is mixing, and this strategy is the one that makes player 1 indifferent between playing H and T . Similarly, we ask which strategy of player 1 will make player 2 indifferent between playing H and T , and this must be player 1's equilibrium strategy.

When we have games with more than 2 strategies for each player, then coming up with quick ways to solve mixed strategy equilibria is not as straightforward, and will usually involve more tedious algebra that solves several equations with

several unknowns. If we take the game of rock-paper-scissors, for example, then there are many mixing combinations for each player, and we can't simply check things the way we did for the matching pennies game. We can try to see if there is, however, an equilibrium in which each player mixes between *all* of his strategies. This is done in the appendix of this chapter, and is simply a solution to a system of linear equations. There are computer algorithms that have been developed to solve for all the Nash equilibria of games, and some are quite successful.

L6

Example: Advertising Game

Imagine that a current market segment is monopolized by one incumbent firm. All is well until a potential entrant is considering entry into the market. The firms are now playing a game, in which each must simultaneously decide their business strategy. The potential entrant must commit to set up for entry or to forgo the option of entry, and the incumbent must decide on a possible advertising campaign that will affect the potential gains of a threatening entrant. This kind of interaction is not uncommon, and the question of course is whether we can shed light on the outcomes we expect to see.

The story can be translated into a game once we formalize the actions that each player can choose, and assigns payoffs to the different outcomes. Assume that the incumbent firm (player 1) currently monopolizes the market, and that this market offers a net value of 15 to the sellers in it. The entrant (player 2) must commit to one of two actions:

- Enter (E), which costs 7
- Stay out (O), which costs 0

The incumbent firm can try to deter entry by choosing one of three ad campaign strategies:

- No ad (N), which costs 0
- Low level ad (L), which costs 2
- High level ad (H), which costs 4

Assume that if the entrant does not enter and stays out, then all the market value of 15 will accrue to the incumbent. If, however, the entrant enters the market, then the two firms will split the market value of 15 according to the ad campaign that the incumbent chose as follows:

- If 1 chose N then 2 gets all the market
- If 1 chose L then 2 gets 0.6 of the market
- if 1 chose H then 2 gets 0.4 of the market

Now that we have taken the informal story and formalized it with some assumptions on the actions and the associated payoffs, we can write this game down in the following matrix form,

		Player 2	
		E	O
Player 1	H	5, -1	11, 0
	L	4, 2	13, 0
	N	0, 8	15, 0

First note that there is no pure strategy Nash equilibrium of this game, which implies that we need to look for a mixed strategy Nash equilibrium. To do this, it is easier to start with finding which strategies of player 2 would make player 1 indifferent between some of his actions. The reason is that player 2's mixed strategy can be represented by a single probability (say of playing E), which is a consequence of player 2 having only two pure strategies. In contrast, we need two probabilities to fully represent a mixed strategy for player 1.

Let q be the probability that player 2 plays E . Given any choice of $q \in [0, 1]$, the expected payoff of player 1 from each of his three pure strategies can be written as follows:

$$\begin{aligned}
 u_1(H, q) &= q \cdot 5 + (1 - q) \cdot 11 = 11 - 6q \\
 u_1(L, q) &= q \cdot 4 + (1 - q) \cdot 13 = 13 - 9q \\
 u_1(N, q) &= q \cdot 0 + (1 - q) \cdot 15 = 15 - 15q
 \end{aligned}$$




FIGURE 9.1.

As with the matching pennies game, we will use these expected utility values to find the best response correspondence of player 1. Playing H will dominate L if and only if $u_1(H, q) > u_1(L, q)$, which is true when $11 - 6q > 13 - 9q$, or $q > \frac{2}{3}$. Playing H will dominate N if and only if $u_1(H, q) > u_1(N, q)$, which is true when $11 - 6q > 15 - 15q$, or $q > \frac{4}{9}$. Thus, from these two observations we know that H is the unique best response if $q > \frac{2}{3}$.

From the analysis above we already know that playing L will dominate H if and only if $q < \frac{2}{3}$. Playing L will dominate N if and only if $u_1(L, q) > u_1(N, q)$, which is true when $13 - 9q > 15 - 15q$, or $q > \frac{1}{3}$. Thus, from these two observations we know that L is the unique best response if $\frac{2}{3} > q > \frac{1}{3}$. Finally, from these observations and the ones above we know that N is the unique best response if $q < \frac{1}{3}$.

If we draw these expected utilities, $u_1(H, q)$, $u_1(L, q)$ and $u_1(N, q)$ as a function of q , we can see the graphical representation of each region of q to find the best response correspondence of player 1 as a function of q .

From the graph we can see that for $q < q_1 = \frac{1}{3}$ (low probability of entry) player 1 prefers no ad (N), for $q > q_2 = \frac{2}{3}$ (high probabilities of entry) player 1 prefers high

ad (H), and for $q \in (q_1, q_2)$ (intermediate probabilities of entry) player 1 prefers low ad (L).

It should be now be evident that when $q = q_1$, player 1 would be willing to mix between N and L , while when $q = q_2$, player 1 would be willing to mix between H and L . Furthermore, player 1 would not be willing to mix between N and H because for the value of q for which $u_1(N, q) = u_1(H, q)$, we can see that playing L yields a higher utility.¹¹

It follows that player 1's best response is,

$$BR_1(q) = \begin{cases} N & \text{if } q \leq \frac{1}{3} \\ L & \text{if } q \in [\frac{1}{3}, \frac{2}{3}] \\ H & \text{if } q \geq \frac{2}{3} \end{cases}.$$

By writing the best response in this way it follows that for $q = \frac{1}{3}$, player 1 has two pure strategy best responses, N and L , which means that any mixture between them is also a best response. Similarly for L and H when $q = \frac{2}{3}$.

Since we have already established by looking at the matrix that there is no pure strategy Nash equilibrium, it must be the case that player 2 must also be mixing in equilibrium. Thus, there can be two candidates for a Nash equilibrium as follows:

1. player 1 will mix between N and L and player 2 plays $q = \frac{1}{3}$;
2. player 1 will mix between H and L and player 2 plays $q = \frac{2}{3}$

To check if either of the above can be a mixed strategy equilibrium, we need to look at two separate cases. First, see if we can find a mixed strategy for player 1 consisting of N and L with positive probability that makes player 2 indifferent, so that he in turn is willing to play $q = \frac{1}{3}$. Second, see if we can find a mixed strategy for player 1 consisting of L and H with positive probability that would make player 2 indifferent, so that he can play $q = \frac{2}{3}$.

¹¹To see this in a graph the graph has to be carefully drawn and accurate. It is safer to check for the possible mixing using math as we did above. To see what combinations of mixtures are possible, we need to first calculate the q that makes two strategies equally good, and then check that the third strategy is not better for that q . If we have more than three strategies then the idea is the same, but we need to check all the other strategies against the two indifferent ones.

However, if we are observant and look back at this particular game then we can notice a shortcut. If player 2 chooses $q = \frac{1}{3}$ as part of a mixed strategy Nash equilibrium, then player 1 will play a mixed strategy of N and L but not H . cannot be part of a Nash equilibrium. However, if player 1 is *not playing* H and only mixing between N and L then effectively the players are playing the following game:

	E	O
L	4, 2	13, 0
N	0, 8	15, 0

Notice now that in this game player 2 has a dominant strategy, which is E , and therefore we will not be able to find *any mixing* of player 1 between N and L that would make player 2 indifferent between E and O , so he would not be willing to mix and choose $q = \frac{1}{3}$. This implies that candidate 1 above is inconsistent with the requirement of Nash equilibrium, and we can now check for candidate 2.

Since for this candidate equilibrium player 1 is *not playing* N and only mixing between H and L , then effectively the players are playing the following game:

	E	O
H	5, -1	11, 0
L	4, 2	13, 0

Consider the case where player 1 mixes between L and H , and let player 1 play H with probability p and L with probability $1 - p$. We will skip the process of finding the best response for player 2 (which would be a good exercise) and proceed directly to find the mixed strategy Nash equilibrium of this reduced game. We need to find the value of p for which player 2 is indifferent between E and O , that is, we solve for $u_2(E, p) = u_2(O, p)$, or,

$$p \cdot (-1) + (1 - p) \cdot 2 = 0$$

and this yields $p = \frac{2}{3}$. Thus, if player 2 mixes with $q = \frac{2}{3}$ then player 1 is willing to mix between H and L , and if player 1 mixes with $p = \frac{2}{3}$ then player 2 will be willing to mix between E and O . Thus the unique Nash equilibrium is,

- Player 1 chooses $(\sigma_1(H), \sigma_1(L), \sigma_1(N)) = (\frac{2}{3}, \frac{1}{3}, 0)$
 - Player 2 chooses $(\sigma_2(E), \sigma_2(O)) = (\frac{2}{3}, \frac{1}{3})$
-

9.4 IESDS in Mixed Strategies[♠]

As we have seen above, by introducing mixed strategies we offered two advancements: First, players can have richer beliefs, and second, players can choose a richer set of actions. This second advancement can be useful when we reconsider the idea of IESDS. In particular, we can now offer the following definition:

Definition 17 Let $\sigma_i \in \Delta S_i$ and $s'_i \in S_i$ be feasible strategies for player i . We say that s'_i is **strictly dominated** by σ_i if

$$u_i(\sigma_i, s_{-i}) > u_i(s'_i, s_{-i}) \quad \forall s_{-i} \in S_{-i}.$$

That is, to consider a strategy as strictly dominated, we no longer require that some other *pure strategy* dominate it, but allow for mixtures to dominate it as well. It turns out that this allows IESDS to have more bite. For example, consider the following game,

	L	C	R		$\frac{1}{2}C + \frac{1}{2}R$
U	5,1	1,4	1,0	player 2's payoff from mixing C and R	2
M	3,2	0,0	3,5		2.5
D	4,3	4,4	0,3		3.5

and denote mixed strategies for players 1 and 2 as triplets, $(\sigma_1(U), \sigma_1(M), \sigma_1(D))$ and $(\sigma_2(L), \sigma_2(C), \sigma_2(R))$ respectively. It is easy to see that no pure strategy is strictly dominated by another pure strategy for any player. However, we can do the following sequence of IESDS with mixed strategies:

1. $(0, \frac{1}{2}, \frac{1}{2}) \succ_2 L$
2. in the resulting game, $(0, \frac{1}{2}, \frac{1}{2}) \succ_1 U$

Thus, after two stages of IESDS we have reduced the game above to,

	<i>C</i>	<i>R</i>
<i>M</i>	0,0	3,5
<i>D</i>	4,4	0,3

How can we find these dominated strategies? Well, a good eye for the numbers is what it takes, short of a computer program or brute force. Also, notice that there are other mixed strategies that would work because strict dominance implies that if we add a small $\varepsilon > 0$ to one of the probabilities, and subtract it from another, then the resulting expected utility from the new mixed strategies can be made arbitrarily close to that of the original one, thus it too would dominate.

9.5 Multiple Equilibria: Pure and Mixed

When we have games with multiple pure strategy Nash equilibria, it turns out that they will often have other Nash equilibria in mixed strategies. Let us consider the 2x2 game,

	<i>C</i>	<i>R</i>
<i>M</i>	0,0	3,5
<i>D</i>	4,4	0,3

It is easy to check that (M, R) and (D, C) are both pure strategy Nash equilibria. It turns out that in cases like this, when there are two distinct pure strategy Nash equilibria, there will generally be a third one in mixed strategies. For this game, let p be player 1's mixed strategy where $p = \sigma_1(M)$, and let q be player 2's strategy where $q = \sigma_2(C)$.

Player 1 will mix when $u_1(M, q) = u_1(D, q)$, or,

$$q \cdot 0 + (1 - q) \cdot 3 = q \cdot 4 + (1 - q) \cdot 0 \Rightarrow q = \frac{3}{7},$$

and player 2 will mix when $u_2(p, C) = u_2(p, R)$, or,

$$p \cdot 0 + (1 - p) \cdot 4 = p \cdot 5 + (1 - p) \cdot 3 \Rightarrow p = \frac{1}{6}.$$

This yields our third Nash equilibrium: $(p, q) = (\frac{1}{6}, \frac{3}{7})$.




FIGURE 9.2.

It is interesting to see that all three equilibria would show up in a careful drawing of the best response functions. Using the utility functions $u_1(M, q)$ and $u_1(D, q)$ we have,

$$BR_1(q) = \begin{cases} p = 1 & \text{if } q < \frac{3}{7} \\ p \in [0, 1] & \text{if } q = \frac{3}{7} \\ p = 0 & \text{if } q > \frac{3}{7} \end{cases}.$$

Similarly, using the utility functions $u_2(p, C)$ and $u_2(p, R)$ we have,

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < \frac{1}{6} \\ q \in [0, 1] & \text{if } p = \frac{1}{6} \\ q = 0 & \text{if } p > \frac{1}{6} \end{cases}.$$

We can draw these two best response correspondences, which appear in figure X.X,

and the three Nash equilibria are revealed: $(p, q) \in \{(1, 0), (\frac{1}{6}, \frac{3}{7}), (0, 1)\}$ are Nash equilibria, where $(p, q) = (1, 0)$ corresponds to the pure strategy (M, C) , and $(p, q) = (0, 1)$ corresponds to the pure strategy (D, R) .

Remark 6 *It may be interesting to know that there are generically (a form of “almost always”) an odd number of equilibria. To prove this one must resort to rather heavy techniques.*




FIGURE 9.3.

9.6 Existence: Nash's Theorem♠♠

What can we say about Nash equilibrium? Unlike IESDS for example, the Bertrand game with different marginal costs did not give rise to a well defined best response, which in turn failed to give a Nash equilibrium. In his seminal dissertation, Nash offered conditions under which this pathology does not arise, and as the Bertrand example suggested, it has to do with continuity of the payoff functions. We now state Nash's theorem:

Theorem 8 (Nash's Theorem) *If Γ is an n -player game with finite strategy spaces S_i , and continuous utility function $u_i : S \rightarrow \mathbb{R}$, $i \in N$, then there exists a Nash Equilibrium in mixed strategies.*

Proving this theorem is well beyond the scope of this text, but it is illuminating to provide some intuition. The idea of Nash's proof builds on a Fixed-Point Theorem. Consider a function $f : [0, 1] \rightarrow [0, 1]$. Brauer's Fixed-Point Theorem states that if f is continuous then there exists some $x^* \in [0, 1]$ that satisfies $f(x^*) = x^*$. The intuition for this theorem can be captured by the graph in figure X.X

How does this relate to game theory? Nash showed that if the utility functions are continuous, then something like continuity is satisfied for the *best response correspondence* of each player. That is, if we have a sequence of strategies $\{\sigma_{-i}^k\}_{k=1}^{\infty}$ for i 's opponents that converges to $\hat{\sigma}_{-i}$, then there is a converging subsequence of player i 's best response $BR_i(\sigma_{-i}^k)$ that converges to an element of his best response $BR_i(\hat{\sigma}_{-i})$.

We can then proceed by considering the *collection* of best response correspondences, $BR \equiv BR_1 \times BR_2 \times \cdots \times BR_n$, which operates from ΔS to itself. That is, $BR : \Delta S \rightrightarrows \Delta S$ takes every element $\sigma' \in \Delta S$, and converts it into a subset $BR(\sigma') \subset \Delta S$. Since ΔS_i is a simplex, it must be compact (closed and bounded, just like $[0,1]$).

Nash then applied a more powerful extension of Brauer's theorem, called Kakutani's fixed point theorem, that says that under these conditions, there exists some σ^* such that $\sigma^* \in BR(\sigma^*)$, that is, $BR : \Delta S \rightrightarrows \Delta S$ has a fixed point. This precisely means that there is some σ^* for which the operator of the collection of best response correspondences, when operating on σ^* , includes σ^* . This fixed point means that for every i , σ_i^* is a best response to σ_{-i}^* , which is the definition of a Nash equilibrium.