

Dominant strategies

Definition

A strategy s_i of some player i is *strictly dominated* by some other strategy s'_i if for all s_{-i} we have

$$u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$$

A strategy s_i is *weakly dominated* by s'_i if this inequality holds only with weak inequality, \geq .

We say that a strategy is *dominated* if there is some other strategy which dominates it; a strategy is *dominant* if it dominates all other strategies. Notice the asymmetry in these definitions: for a strategy to be dominated, we only need one strategy to dominate it; for a strategy to be dominant, it must dominate all other available strategies. Intuitively, a strategy is strictly dominated if there is some other strategy which offers a higher utility, *regardless of the strategies of the other players*.

	<i>L</i>	<i>C</i>	<i>R</i>
<i>T</i>	1, 1	2, 0	3, 2
<i>M</i>	-1, 2	3, 4	5, 3
<i>B</i>	0, 3	4, 5	2, 4

In this game, player 1 (rows) chooses from **T**op, **M**iddle, and **B**ottom, while player 2 (columns) chooses from **L**eft, **C**enter, and **R**ight. We can check to see that player 1 has no dominated strategy:

- If player 2 plays *L*, player 1 prefers *T* to both *M* and *B*. Therefore *T* is not dominated.
- If player 2 plays *C*, player 1 prefers *B* to both *T* and *M*. Therefore *B* is not dominated.
- If player 2 plays *R*, player 1 prefers *M* to both *T* and *B*. Therefore *M* is not dominated.

What is key here is that we check to make sure that there are circumstances under which each strategy is preferred to each other strategy. This is necessary to know that the strategy is not dominated; if we cannot, for example, find a circumstance under which *T* is preferred to *M* (unlike the present game), *T* would be dominated by *M*!

The same cannot be said of player 2. We can check this,

- If player 1 plays *T*, player 2 prefers *R* to *L*, receiving 2 over 1.
- If player 1 plays *M*, player 2 prefers *R* to *L*, receiving 3 over 2.
- If player 1 plays *B*, player 2 prefers *R* to *L*, receiving 4 over 3.

That is, regardless of what action player 1 takes player 2 prefers playing *R* to playing *L*; notice that this completely ignores strategy *C*, which neither dominates nor is dominated by *L* or *R*. Therefore, *L* is dominated by *R*. Intuitively, in practice it would be unreasonable to expect an agent to play *L* as player 2 in this game, as she can always do better by playing *R*. This allows us to make some predictions as to what a rational agent *will not* do.

Iterated elimination of strictly-dominated strategies

We can use dominance to reduce the set of strategies we might expect to be played in the game. Considering that we expect players to behave rationally, though, we may actually be able to say a lot more about what we expect to be played. That is, in the above case, since player 1 can see that player 2 will never play L he may be able to make some decisions about what he is not willing to play. We call this step of reasoning *iterated elimination of strictly-dominated strategies*, abbreviated IESDS.

Definition

Expected play in a game G can be pared down by the process of *iterated elimination of strictly-dominated strategies*. The algorithm for this process is

- (a) If G contains no strictly-dominated strategies, we are done; the resulting game is G .
 - (b) Eliminate *some* strictly-dominated strategy. If there are multiple strictly-dominated strategies, pick one arbitrarily.
 - (c) Consider the game G without the eliminated strategies as a new game, G' . Using G' , return to step 1 above.
-

Since this definition is a little abstract, we'll demonstrate this with an example. Consider our previous game,

	L	C	R
T	1, 1	2, 0	3, 2
M	-1, 2	3, 4	5, 3
B	0, 3	4, 5	2, 4

Since L is dominated for player 2, we can effectively remove it from the strategy space; player 2 never has any reason to play L , so it is irrelevant. We get the new game

	C	R
T	2, 0	3, 2
M	3, 4	5, 3
B	4, 5	2, 4

In this new game, T is dominated by M for player 1 (why?); notice that before L was removed, T was not dominated by M , so we're very heavily relying upon the assumption that L will not be played. We remove T from player 1's strategies and obtain the new game

	C	R
M	3, 4	5, 3
B	4, 5	2, 4

In this new game, R is dominated by C for player 2 (why?). So we remove R from player 2's strategies and obtain the new game

	C
M	3, 4
B	4, 5

Player 1's choice is now easy: he can play M and receive utility of 3, or play B and receive utility of 4. He will play B and the remaining game is

$$B \quad \begin{array}{c} C \\ \boxed{4,5} \end{array}$$

By applying IESDS, we have reduced reasonable play in this game to a unique strategy pair (B, C) . The rationale here is fairly straightforward: since player 2 will not play L , player 1 will never play T . Since player 1 will not play T , player 2 will never play R . We can iterate this logic forward to see that (B, C) is the only "reasonable" outcome of this game.

There are several important points about IESDS:

- We consider only strictly-dominated strategies, and not weakly-dominated strategies. This does make a difference.
- Because we consider only strictly-dominated strategies, the order of elimination doesn't matter (if we consider eliminating weakly-dominated strategies, order *may* matter, and so the process does not have a well-defined outcome). "Order" here refers not just to the order in which we eliminate the strategies of a particular player, but also the order in which apply this process to each player (i.e., we may look first at player 1 or player 2, it is of no consequence).
- We can pick dominated strategies to eliminate completely arbitrarily; we can also eliminate all dominated strategies in one fell swoop. This is because if a strategy is dominated in one step of IESDS, it will also be dominated in all subsequent steps; we won't prove that, but it is worth thinking about.

These considerations make IESDS a useful procedure, as it's not particularly sensitive to the details of how the elimination happens.

If application of IESDS results in a single undominated strategy for each player, we call the game *dominance solvable*; most games are not dominance solvable. Still, if we can use the notion of dominance to solve for the unique undominated play of the game we can make a fair prediction about what to expect if we were to actually watch agents play the game. In the case of 2-player games, if the game is dominance solvable the resulting undominated strategy is a Nash equilibrium.

Nash equilibrium

That provides a good segue to Nash equilibrium. To properly define a Nash equilibrium, we need to define a particular property of certain strategies in relation to others.

Definition

A strategy s_i for player i is a *best response* to the strategy profile s_{-i} of all other players if for all s'_i ,

$$u(s_i, s_{-i}) \geq u(s'_i, s_{-i})$$

That is, a strategy is a best response for a particular player if, given what all other players are doing, she can do no better. Notice that this definition is very close to the definition for dominance; the key here is that a strategy s_i is a best response in relation to some behaviour s_{-i} of the other players. In the case of dominance, the strategy s_i had to be preferred regardless of what the other players we doing, not just preferred in a specific instance.

Definition

A strategy profile s^* is a *Nash equilibrium* if s_i^* is a best response to s_{-i}^* for all i .

Simply put, a strategy profile is a Nash equilibrium if all players are best-responding to each other's play. If this is the case, a player will not want to deviate (she can do no better by playing any other strategy); since she's happy playing her particular strategy, and the other players are best-responding to her actions, they also have no incentive to deviate. This is the sense in which this is a notion of equilibrium.

In normal-form games (those with a payoff matrix) finding Nash equilibrium is not terribly difficult, *if it exists* (more on this later, in the section on MSNE). For a particular strategy of players $-i$, underline the payoff of player i at his preferred action(s). When this is done for all players, we will have underlined the best responses for each player; if all payoffs in a particular cell are underlined the players must be best-responding to each other, and so we have a Nash equilibrium.

Using the game above, we have

	<i>L</i>	<i>C</i>	<i>R</i>
<i>T</i>	<u>1</u> , 1	2, 0	3, <u>2</u>
<i>M</i>	-1, 2	3, <u>4</u>	<u>5</u> , 3
<i>B</i>	0, 3	<u>4</u> , <u>5</u>	2, 4

In words, this is

- If player 1 plays *T*, player 2 is best off playing *R*.
- If player 1 plays *M*, player 2 is best off playing *C*.
- If player 1 plays *B*, player 2 is best off playing *C*.
- If player 2 plays *L*, player 1 is best off playing *T*.
- If player 2 plays *C*, player 1 is best off playing *B*.
- If player 2 plays *R*, player 1 is best off playing *M*.

Since the only cell in this payoff matrix with co-occurrent underlines is (B, C) , (B, C) is the unique Nash equilibrium of the game; as mentioned above, this corresponds to the outcome of IESDS (but this is true only because this game is dominance solvable). Since player 1 is happiest playing *B* when player 2 plays *C*, and player 2 is happiest playing *C* when player 1 plays *B*, neither player wants to deviate and this is an equilibrium.

To offer another example, we can introduce another “standard” game¹. This game is called *Battle of the Sexes*, and corresponds roughly to two individuals looking to go out on a date². They each have preference for a particular activity, but in the end would rather go on a date together than enjoy their favored activity alone. In this case, let's assume that the agents are going out to dinner; player 1 prefers Italian food, while player 2 prefers Mexican food. The normal form of this game is

¹Game theory has a number of standard, or well-known games. These include the prisoner's dilemma, matching pennies, the centipede game, battle of the sexes, etc. Although payoffs are not always the same in each version of these particular games, each version will retain a certain set of properties. In the case of the prisoner's dilemma, this property is that in equilibrium, agents do as poorly as possible; in the case of matching pennies, these properties are symmetry, zero-sum-ness, and lack of a pure-strategy Nash equilibrium.

²The word you may be looking for here is, “heteronormative.” Of course, considering when game theory came of age we can't be too surprised that we talk of dates between men and women; we could update the language, but that would involve fixing a lot of very old, very important papers and articles. This particular situation could just as well model the fact that my college roommate prefers *Call of Duty* while I prefer *Battlefield* even though we'd both rather play video games online together. Still, *Battle of the College Roommates* doesn't have quite the same ring to it.

	<i>I</i>	<i>M</i>
<i>I</i>	2, 1	0, 0
<i>M</i>	0, 0	1, 2

If it helps, think for a moment how this game form corresponds to the situation set forth in the previous paragraph.

We can find Nash equilibrium in the usual way,

	<i>I</i>	<i>M</i>
<i>I</i>	<u>2</u> , <u>1</u>	0, 0
<i>M</i>	0, 0	<u>1</u> , <u>2</u>

In this case, we have two distinct pure-strategy Nash equilibria: (I, I) and (M, M) . This should seem pretty reasonable on its face, since both players would rather go on a date than be alone. More power to Nash equilibrium!

However, we can alter the payoffs slightly to demonstrate how Nash equilibrium sometimes “breaks,” at least in terms of what a reasonable person would expect to see in the game. Suppose that Italian restaurants in Los Angeles are not date-friendly, so the players’ utility is diminished relative to what it was. We might state this new game as

	<i>I</i>	<i>M</i>
<i>I</i>	0, 0	0, 0
<i>M</i>	0, 0	1, 2

Notice that this game is not dominance-solvable (why?!). When we try to find Nash equilibrium, the following results:

	<i>I</i>	<i>M</i>
<i>I</i>	<u>0</u> , <u>0</u>	<u>0</u> , <u>0</u>
<i>M</i>	0, 0	<u>1</u> , <u>2</u>

According to our process for uncovering Nash equilibria, we retain our original set of equilibria, (I, I) and (M, M) . To a reasonable human being, this should seem counterintuitive: both agents clearly prefer going out together for Mexican food to going out together for Italian food, so why would they go out for Italian food? The simple answer is that Nash equilibrium just does not have the power to make this sort of prediction. If player 1 goes out for Italian food, player 2 is equally happy going out for Italian food or Mexican food, so might as well go out for Italian food.

Later in the quarter, we may discuss remedies to this situation, but for now it will stand as a reminder that Nash equilibrium is not necessarily an infallible descriptor of what we will see in the real world.

Sequential games

In real life, it is easy to imagine situations in which one agent moves first and another agent responds. We refer to this setup as a *sequential game*. Here, after the first agent moves the second agent will respond conditional on what the first player did. In the case of 2-player games, if player 1 moves first and player 2 responds, we express this as $a_2|s_1$, or the action that player 2 takes “given” player 1’s strategy. A strategy in this case for player 2 is a set of conditional strategies.

Continuing the original example of Battle of the Sexes, suppose that player 1 moves first and player 2 responds to her action. We update the normal form of the game to

	$I I, I M$	$I I, M M$	$M I, I M$	$M I, M M$
I	2, 1	2, 1	0, 0	0, 0
M	0, 0	1, 2	0, 0	1, 2

The strategy pair $[I, (I|I, I|M)]$ is read as, “**I**talian, **I**talian given **I**talian, and **I**talian given **M**exican,” and results in the payoff vector (2, 1). Here, player 2 essentially goes out for **I**talian food regardless of what player 1 does.

To find Nash equilibria, we use the usual method of underlining.

	$I I, I M$	$I I, M M$	$M I, I M$	$M I, M M$
I	<u>2</u> , <u>1</u>	2, <u>1</u>	0, 0	0, 0
M	0, 0	1, <u>2</u>	<u>0</u> , 0	1, <u>2</u>

So this game has three Nash equilibria: $[I, (I|I, I|M)]$, $[I, (I|I, M|M)]$, and $[M, (M|I, M|M)]$. These can be read as

- $[I, (I|I, I|M)]$: player 1 goes out for **I**talian food, and player 2 goes out for **I**talian food regardless of what player 1 does.
- $[I, (I|I, M|M)]$: player 1 goes out for **I**talian food, and player 2 goes wherever player 1 goes.
- $[M, (M|I, M|M)]$: player 2 goes out for **M**exican food, and player 2 goes out for **M**exican food regardless of what player 1 does.

The last equilibrium, in particular, seems strange: by in effect threatening to go out for **M**exican food no matter what, player 2 forces player 1 to go out for **M**exican food. But this is an empty threat! If player 1 did go out for **I**talian food, player 2 would rather follow than go out for **M**exican food. Nash equilibrium does not give us the power to think about this effect in this game (we will see how to accommodate this in lecture before too long). One way to consider the machinery in the background here is that player 2 can commit to going out for **M**exican food regardless, and player 1 knows that this cannot be changed and so chooses to go out for **M**exican food. If necessary, we sometimes say that player 2, having stated this strategy and being committed to it, will be shot if he does not follow it; this offers a pretty sound incentive to follow the stated strategy, but keep in mind that shooting of players is outside of the game we specified, so we should rely too heavily on this justification. The takeaway here is that agents commit to actions, and other agents believe them.

Of course, there are technically the same issues in the first equilibrium listed, but they do not set off our intuition: since player 1 would rather go out for **I**talian, there is no incentive to test player 2’s strategy of going out for **I**talian regardless.

Mixed-strategy Nash equilibrium

Switching gears, we can broaden our set of play by allowing players to randomly select the strategy they apply.

Definition

A *mixed strategy* σ_i for player i is a probability distribution over i ’s pure strategies. When i has a finite (more generally, a countable) set of actions, we can think of this as

$$\sigma_i^j = \Pr(s_i = a_i^j)$$

where a_i^j is the j^{th} action available to agent i .

Being a little obtuse, we can see this through an example. In the original, nonsequential Battle of the Sexes, suppose that player 1 randomly goes out for **I**talian food or **M**exican food, assigning a 50% probability to each. In the notation of a mixed strategy, this is $\sigma_1 = (\frac{1}{2}, \frac{1}{2})$. What is player 2's best response?

To answer this question, we need to move into the realm of expected utility. The argument here is that since player 2 doesn't know which action player 1 will take, he must account for the inherent randomness somehow; if he expects more utility from one action than another, he will take this action as his strategy. We then compute

$$\begin{aligned} E[u_2(I, \sigma_1)] &= \Pr(s_1 = I)u_2(I, I) + \Pr(s_1 = M)u_2(I, M) \\ &= \frac{1}{2}(1) + \frac{1}{2}(0) \\ &= \frac{1}{2} \end{aligned}$$

$$\begin{aligned} E[u_2(M, \sigma_1)] &= \Pr(s_1 = I)u_2(M, I) + \Pr(s_1 = M)u_2(M, M) \\ &= \frac{1}{2}(0) + \frac{1}{2}(2) \\ &= 1 \end{aligned}$$

So player 2's expected utility directly takes into account how player 1 is randomizing among strategies. When he does this, his expected utility from going out for **I**talian food is $\frac{1}{2}$ while his expected utility from going out for **M**exican food is 1. With this in mind, since his expected utility is higher if he goes out for **M**exican food, this is the strategy he will employ.

Knowing that player 2 goes out for **M**exican food, player 1 will strictly prefer going out for **M**exican food to going out for **I**talian food, and will not want to randomize. So $\sigma_1 = (\frac{1}{2}, \frac{1}{2})$ cannot be part of an equilibrium.

Definition

A mixed strategy profile σ^* is a *mixed strategy Nash equilibrium* if for all players i , and all alternate strategies σ'_i we have

$$E[u_i(\sigma_i^*, \sigma_{-i}^*)] \geq E[u_i(\sigma'_i, \sigma_{-i}^*)]$$

In words, σ^* is a mixed strategy Nash equilibrium if each player is best-responding to the mixed strategies of other players, as compared by expected utility. As a quick point, every Nash equilibrium in pure strategies is also a mixed strategy Nash equilibrium: a pure strategy is a (degenerate) probability distribution which assigns all probability to a particular action. A mixed strategy in which each pure strategy is played with some positive probability is called a *totally-mixed strategy*.

Now, let's consider what mixed strategy Nash equilibria might exist in the Battle of the Sexes. Since a rational agent will always choose the strategy with the highest payoff, she will only mix if she is indifferent between some set of strategies. That is, she is willing to randomize among some set of strategies which yield the same expected utility.

So let's consider the situation in which player 1 is indifferent between playing I and playing M . We represent this as

$$\begin{aligned} E[u_1(I, \sigma_2)] &= E[u_1(M, \sigma_2)] \\ \Pr(s_2 = I)u_1(I, I) + \Pr(s_2 = M)u_1(I, M) &= \Pr(s_2 = I)u_1(M, I) + \Pr(s_2 = M)u_1(M, M) \\ \Pr(s_2 = I)(2) + \Pr(s_2 = M)(0) &= \Pr(s_2 = I)(0) + \Pr(s_2 = M)(1) \\ 2\Pr(s_2 = I) &= \Pr(s_2 = M) \end{aligned}$$

Notice that player 1's expected utility is stated in terms of the probability that player 2 assigns to particular strategies! This makes mathematical sense, but can seem a little backwards: we are using the indifference condition of player 1 to determine the probability that player 2 assigns in his mixture. We can then restate this as

$$\begin{aligned} 2\Pr(s_2 = I) &= \Pr(s_2 = M) \\ 2\sigma_2^I &= \sigma_2^M \end{aligned}$$

By definition, σ_2 is a probability distribution over the strategies of player 2. From the rules of probability, this tells us $\sigma_2^I + \sigma_2^M = 1$, or $\sigma_2^M = \sigma_2^I$. We then continue our derivation,

$$\begin{aligned} 2\sigma_2^I &= \sigma_2^M \\ 2\sigma_2^I &= 1 - \sigma_2^I \\ 3\sigma_2^I &= 1 \\ \implies \sigma_2^I &= \frac{1}{3} \end{aligned}$$

So if player 2 assigns probability $\frac{1}{3}$ to going out for Italian food and $\frac{2}{3}$ to going out for Mexican food, player 1 is indifferent between her strategies. This mixed strategy is denoted $\sigma_2 = (\frac{1}{3}, \frac{2}{3})$. We can check this,

$$\begin{aligned} E[u_1(I, \sigma_2)] &= \frac{1}{3}(2) + \frac{2}{3}(0) \\ &= \frac{2}{3} \end{aligned}$$

$$\begin{aligned} E[u_1(M, \sigma_2)] &= \frac{1}{3}(0) + \frac{2}{3}(1) \\ &= \frac{2}{3} \end{aligned}$$

Player 1 expects utility of $\frac{2}{3}$ regardless of the strategy she plays.

With this, we have the randomization probabilities that player 2 must use to cause player 1 to randomize. But when will player 2 randomize? Again, he is willing to randomize when he is indifferent between his strategies. We solve this just as for player 1 (but without the commentary),

$$\begin{aligned} E[u_2(I, \sigma_1)] &= E[u_2(M, \sigma_1)] \\ \Pr(s_1 = I)u_2(I, I) + \Pr(s_1 = M)u_2(I, M) &= \Pr(s_1 = I)u_2(M, I) + \Pr(s_1 = M)u_2(M, M) \\ \sigma_1^I(1) + \sigma_1^M(0) &= \sigma_1^I(0) + \sigma_1^M(2) \\ \sigma_1^I &= 2\sigma_1^M \\ 1 - \sigma_1^M &= 2\sigma_1^M \\ \implies \sigma_1^M &= \frac{1}{3} \end{aligned}$$

So player 2 is indifferent between his strategies when player 1 mixes, assigning probability $\frac{2}{3}$ to going out for Italian food and probability $\frac{1}{3}$ to going out for Mexican food. This mixed strategy is denoted $\sigma_1 = (\frac{2}{3}, \frac{1}{3})$.

Since each player is perfectly happy to randomize between strategies — this is precisely what we have solved for — we see that $\sigma^* = [(\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3})]$ is a mixed strategy Nash equilibrium. This holds because neither player can do better by switching probabilities or by changing to a particular pure strategy.

The question should be asked, though: what are agents' expected utilities? We have already seen $E[u_1(\sigma^*)] = \frac{2}{3}$ (we did not show this precise equality, but we showed something which turns out to be identical; it may help to think about this). Similarly,

$$\begin{aligned} E[u_2(\sigma^*)] &= \Pr(s_1 = I, s_2 = I)u_2(I, I) + \Pr(s_1 = I, s_2 = M)u_2(M, I) + \\ &\quad \Pr(s_1 = M, s_2 = I)u_2(I, M) + \Pr(s_1 = M, s_2 = M)u_2(M, M) \\ &= \Pr(s_1 = I, s_2 = I) + 2\Pr(s_1 = M, s_2 = M) \\ &= \Pr(s_1 = I)\Pr(s_2 = I) + 2\Pr(s_1 = M)\Pr(s_2 = M) \\ &= \sigma_1^I \sigma_2^I + 2\sigma_1^M \sigma_2^M \\ &= \left(\frac{2}{3}\right) \left(\frac{1}{3}\right) + 2 \left(\frac{1}{3}\right) \left(\frac{2}{3}\right) \\ &= \frac{2}{9} + 2 \left(\frac{2}{9}\right) \\ &= \frac{2}{3} \end{aligned}$$

There are a couple of points here: firstly, the joint probability of particular strategies is the product of the marginal probability of the particular strategies. This is because agents' randomization is assumed to be independent. Secondly, the expected overall utility is precisely equal to the expected utility from a particular strategy within the mixed strategy. This makes sense: since the player is indifferent between the strategies she is randomizing amongst, her utility from picking any one is identical, and hence her utility from randomizing among them must be identical to the utility from any particular strategy in her randomization.

In this particular example, notice that the expected utility in the mixed strategy Nash equilibrium — $\frac{2}{3}$ — is strictly less than the utility from either of the pure strategy Nash equilibria (in which neither player ever makes less than 2). This is another example of Nash equilibrium giving us “funny” equilibria; in this sense Battle of the Sexes is not the most motivating example for mixed strategy Nash equilibrium, but solving for a mixed strategy Nash equilibrium in a game in which it does make sense will follow the exact same solution pattern.

As practice, let's assume that player 1 *really* prefers Italian food. We'll change her payoff slightly to generate the new game form

	<i>I</i>	<i>M</i>
<i>I</i>	3, 1	0, 0
<i>M</i>	0, 0	1, 2

Notice that we retain the same pure strategy Nash equilibria (identical except for the payoff that player 1 receives at (I, I)). Can we still find mixed strategy Nash equilibrium? We'll use the standard conditions.

$$\begin{aligned}
 E[u_2(I, \sigma_1)] &= E[u_2(M, \sigma_1)] \\
 \Pr(s_1 = I)u_2(I, I) + \Pr(s_1 = M)u_2(I, M) &= \Pr(s_1 = I)u_2(M, I) + \Pr(s_1 = M)u_2(M, M) \\
 \sigma_1^I(1) + \sigma_1^M(0) &= \sigma_1^I(0) + \sigma_1^M(2) \\
 \sigma_1^I &= 2\sigma_1^M \\
 1 - \sigma_1^M &= 2\sigma_1^M \\
 \implies \sigma_1^M &= \frac{1}{3}
 \end{aligned}$$

$$\begin{aligned}
 E[u_1(I, \sigma_2)] &= E[u_1(M, \sigma_2)] \\
 \Pr(s_2 = I)u_1(I, I) + \Pr(s_2 = M)u_1(I, M) &= \Pr(s_2 = I)u_1(M, I) + \Pr(s_2 = M)u_1(M, M) \\
 \sigma_2^I(3) + \sigma_2^M(0) &= \sigma_2^I(0) + \sigma_2^M(1) \\
 3\sigma_2^I &= \sigma_2^M \\
 3\sigma_2^I &= 1 - \sigma_2^I \\
 \implies \sigma_2^I &= \frac{1}{4}
 \end{aligned}$$

The mixed strategy Nash equilibrium in this case is $\sigma^* = [(\frac{2}{3}, \frac{1}{3}), (\frac{1}{4}, \frac{3}{4})]$. Interestingly, by changing the payoffs for player 1, we change the mixing probabilities of player 2! The mixing probabilities of player 1 are the same. This is because mixing probabilities are solved as a function of indifference for the *other* player. Hopefully this makes it a little more apparent how solving this way will work.

As a final note, we cannot always find pure strategy Nash equilibria (think of Matching Pennies). However, when we allow mixing, we will *always* have a Nash equilibrium of some sort, possibly in mixed strategies. For a bit of practice, show that there are no pure strategy Nash equilibria in matching pennies, and find the mixed strategy Nash equilibrium. Recall that matching pennies is represented by the normal form payoff matrix

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-1, 1	1, -1