Game Theory:
Dominance, Nash Equilibrium, Symmetry

Branislav L. Slantchev
Department of Political Science, University of California – San Diego

May 23, 2008

Contents.

---

1 Elimination of Dominated Strategies

1.1 Strict Dominance in Pure Strategies

In some games, a player's strategy is superior to all other strategies *regardless of what the other players do*. This strategy then *strictly dominates* the other strategies. Consider the Prisoner's Dilemma game in Fig. 1 (p. 2). Choosing D strictly dominates choosing C because it yields a better payoff regardless of what the other player chooses to do.

If one player is going to play D, then the other is better off by playing D as well. Also, if one player is going to play C, then the other is better off by playing D again. For each prisoner, choosing D is always better than C regardless of what the other prisoner does. We say that D strictly dominates C.

		Prisoner 2	
		C	D
Prisoner 1	C	2,2	0,3
	D	3,0	1,1

Figure 1: Prisoner's Dilemma.

DEFINITION 1. In the strategic form game G, let $s_i', s_i'' \in S_i$ be two strategies for player i. Strategy s_i' **strictly dominates** strategy s_i'' if

$$U_i(s_i', s_{-i}) > U_i(s_i'', s_{-i})$$

for every strategy profile $s_{-i} \in S_{-i}$.

In words, a strategy s_i' strictly dominates s_i'' if for each feasible combination of the other players' strategies, i's payoff from playing s_i' is strictly greater than the payoff from playing s_i''. Also, strategy s_i'' is **strictly dominated** by s_i'. In the PD game, D strictly dominates C, and C is strictly dominated by D. Observe that we are using expected utilities even for the pure-strategy profiles because they may involve chance moves.

Rational players never play strictly dominated strategies, because such strategies can never be best responses to any strategies of the other players. There is no belief that a rational player can have about the behavior of other players such that it would be optimal to choose a strictly dominated strategy. Thus, in PD a rational player would never choose C. We can use this concept to find solutions to some simple games. For example, since neither player will ever choose C in PD, we can **eliminate** this strategy from the strategy space, which means that now both players only have one strategy left to them: D. The solution is now trivial: It follows that the only possible rational outcome is $\langle D, D\rangle$.

Because players would never choose strictly dominated strategies, eliminating them from consideration should not affect the analysis of the game because this fact should be evident to all players in the game. In the PD example, eliminating strictly dominated strategies resulted in a unique prediction for how the game is going to be played. The concept is more general, however, because even in games with more strategies, eliminating a strictly dominated one may result in other strategies becoming strictly dominated in the game that remains.

Consider the abstract game depicted in Fig. 2 (p. 3). Player 1 does not have a strategy that is strictly dominated by another: playing U is better than M unless player 2 chooses C, in which case M is better. Playing D is better than playing U unless player 2 chooses R, in which case U is better. Finally, playing D instead of M is better unless player 2 chooses R, in which case M is better.

		Player 2		
		L	C	R
	U	4,3	5,1	6,2
Player 1	M	2,1	8,4	3,6
	D	5,9	9,6	2,8

Figure 2: A 3×3 Example Game.

For player 2, on the other hand, strategy C is strictly dominated by strategy R. Notice that whatever player 1 chooses, player 2 is better off playing R than playing C: she gets $2 > 1$ if player 1 chooses U; she gets $6 > 4$ if player 1 chooses M; and she gets $8 > 6$ if player 1 chooses D. Thus, a rational player 2 would never choose to play C when R is available. (Note here that R neither dominates, nor is dominated by, L.) If player 1 knows that player 2 is rational, then player 1 would play the game *as if* it were the game depicted in Fig. 3 (p. 3).

		Player 2	
		L	R
	U	4,3	6,2
Player 1	M	2,1	3,6
	D	5,9	2,8

Figure 3: The Reduced Example Game, Step I.

We examine player 1's strategies again. We now see that U strictly dominates M because player 1 gets $4 > 2$ if player 2 chooses L, and $6 > 3$ if player 2 chooses R. Thus, a rational player 1 would never choose M given that he knows player 2 is rational as well and consequently will never play C. (Note that U neither dominates, nor is dominated by, D.)

If player 2 knows that player 1 is rational and knows that player 1 knows that she is also rational, then player 2 would play the game *as if* it were the game depicted in Fig. 4 (p. 3).

		Player 2	
		L	R
Player 1	U	4,3	6,2
	D	5,9	2,8

Figure 4: The Reduced Example Game, Step II.

We examine player 2's strategies again and notice that L now strictly dominates R because player 2 would get $3 > 2$ if player 1 chooses U, and $9 > 8$ if player 1 chooses D. Thus, a rational player 2 would never choose R given that she knows that player 1 is rational, etc.

If player 1 knows that player 2 is rational, etc., then he would play the game *as if* it were the game depicted in Fig. 5 (p. 4).

But now, U is strictly dominated by D, so player 1 would never play U. Therefore, player 1's rational choice here is to play D. This means the outcome of this game will be $\langle D, L \rangle$, which yields player 1 a payoff of 5 and player 2 a payoff of 9.

		Player 2
		L
Player 1	U	4,3
	D	5,9

Figure 5: The Reduced Example Game, Step III.

The process described above is called **iterated elimination of strictly dominated strategies**. The solution of G is the equilibrium $\langle D, L\rangle$, and is sometimes called **iterated-dominance equilibrium**, or *iterated-dominant strategy equilibrium*. The game G is sometimes called **dominance-solvable**.

Although the process is intuitively appealing (after all, rational players would never play strictly dominated strategies), each step of elimination requires a further assumption about the other player's rationality. Recall that we started by assuming that player 1 knows that player 2 is rational and so she would not play C. This allowed the elimination of M. Next, we had to assume that player 2 knows that player 1 is rational *and* that she also knows that player 1 knows that she herself is rational as well. This allowed the elimination of R. Finally, we had to assume that player 1 knows that player 2 is rational and that he also knows that player 2 knows that player 1 is rational and that player 2 also knows that player 1 knows that player 2 is rational.

More generally, we want to be able to make this assumption for as many iterations as might be needed. That is, we must be able to assume not only that all players are rational, but also that all players know that all the players are rational, and that all the players know that all the players know that all players are rational, and so on, *ad infinitum*. This assumption is called **common knowledge** and is usually made in game theory.

Consider the following example. There are two bars in a city, with owners labeled A and B, who can charge \$2, \$4, or \$5 per drink. Each day, there are 6,000 tourists and 4,000 locals who decide which bar to visit. (Each person can only go to one bar and each must go to at least one bar, where each person has exactly one drink.) Since tourists have no idea about the bars, they randomize without reference to the pricing. The locals, however, always go to the cheapest bar (and randomize if the prices are the same). The question is: what prices should the owners set if they choose simultaneously? The strategies are $S_i = \{2, 4, 5\}$ and the payoffs can be computed as follows:

$$
\begin{aligned}
U_1(2,2) &= U_2(2,2) = (3,000 + 2,000) \times 2 = 10,000 \\
U_1(2,4) &= U_2(4,2) = (3,000 + 4,000) \times 2 = 14,000 \\
U_1(2,5) &= U_2(5,2) = (3,000 + 4,000) \times 2 = 14,000 \\
U_1(4,2) &= U_2(2,4) = 3,000 \times 4 = 12,000 \\
U_1(4,4) &= U_2(4,4) = (3,000 + 2,000) \times 4 = 20,000 \\
U_1(4,5) &= U_2(5,4) = (3,000 + 4,000) \times 4 = 28,000 \\
U_1(5,2) &= U_2(2,5) = 3,000 \times 5 = 15,000 \\
U_1(5,4) &= U_2(4,5) = 3,000 \times 5 = 15,000 \\
U_1(5,5) &= U_2(5,5) = (3,000 + 2,000) \times 5 = 25,000.
\end{aligned}
$$

These can be summarized in the strategic form game (payoffs are given in thousands of dollars for brevity) in Fig. 6 (p. 5). For either player, charging \$2 is strictly dominated by either of the other two pure strategies. Hence, we can eliminate this, which reduces the game to the intermediate form. We then notice that \$4 strictly dominates \$5 for either player, so we can eliminate \$5, leaving us with a unqiue prediction: $\langle 4, 4\rangle$. That is, each bar will charge \$4 for drinks, and they will split evenly both the tourist and the local populations. Note that \$4 does not strictly dominate \$5 initially: if the other player is expected to charge \$2, then one will do better by charging \$5 (which makes sense: since you are losing the locals anyway, you might as well jack up the prices for the tourists). It is only *after* a player establishes that his opponent will not choose \$2 that \$5 becomes strictly dominated.

		Player B: 2	4	5
Player A	2	10, 10	14, 12	14, 15
	4	12, 14	20, 20	28, 15
	5	15, 14	15, 28	25, 25

$\Rightarrow$

	4	5
4	20, 20	28, 15
5	15, 28	25, 25

$\Rightarrow$

	4
4	20, 20

Figure 6: The Bar Hopping Game.

1.2 Weak Dominance

Rational players would never play strictly dominated strategies, so eliminating these should not affect our analysis. There may be circumstances, however, where a strategy is "not worse" than another instead of being "always better" (as a strictly dominant one would be). To define this concept, we introduce the idea of **weakly dominated strategy**.

Definition 2. In the strategic form game G, let $s_i', s_i'' \in S_i$ be two strategies for player i. Strategy s_i' **weakly dominates** strategy s_i'' if

$$U_i(s_i', s_{-i}) \geq U_i(s_i'', s_{-i})$$

for every strategy profile $s_{-i} \in S_{-i}$, and there exists at least one s_{-i} such that the inequality is strict.

In other words, s_i' never does worse than s_i'', and sometimes does better. While iterated elimination of strictly dominated strategies seems to rest on rather firm foundation (except for the common knowledge requirement that might be a problem with more complicated situations), eliminating weakly dominated strategies is more controversial because it is harder to argue that it should not affect analysis. The reason is that by definition, a weakly dominated strategy can be a best response for the player. Furthermore, there are technical difficulties with eliminating weakly dominated strategies: the order of elimination can matter for the result!

		Player 2: L	R
Player 1	U	3, 2	2, 2
	M	1, 1	0, 0
	D	0, 0	1, 1

Figure 7: The Kohlberg and Mertens Game.

Consider the game in Fig. 7 (p. 5). Strategy D is strictly dominated by U, so if we remove it first, we are left with a game, in which L weakly dominates R. Eliminating R in turn results in a game where U strictly dominates M, so the prediction is $\langle U, L\rangle$. However, note that M is strictly dominated by U in the original game as well. If we begin by eliminating M, then R weakly dominates L in the resulting game. Eliminating L in turn results in a game where U strictly dominates D, so the prediction is $\langle U, R\rangle$. If we begin by eliminating M and D at the same time, then we are left with a game where neither of the strategies for player 2 weakly dominates the other. Thus, the order in which we eliminate the strictly dominated strategies for player 1 determines which of player 2's weakly dominated strategies will get eliminated in the iterative process.

This dependence on the order of elimination does not arise if we only eliminated strictly dominated strategies. If we perform the iterative process until no strictly dominated strategies remain, the resulting game will be the same regardless of the order in which we perform the elimination. Eliminating strategies for other players can never cause a strictly dominated strategy to cease to be dominated but it can cause a weakly dominated strategy to cease being dominated. Intuitively, you should see why the latter might be the case. For a strategy s_i to be weakly dominated, all that is required is that some other strategy s_i' is as good as s_i for all strategies s_{-i} and only better than s_i for one strategy of the opponent. If that particular strategy gets eliminated, then s_i and s_i' yield the same payoffs for all remaining strategies of the opponent, and neither weakly dominates the other.

1.3 Strict Dominance and Mixed Strategies

We now generalize the idea of dominance to mixed strategies. All that we have to do to decide whether a pure strategy is dominated is to check whether there exists some mixed-strategy that is a better response to all pure strategies of the opponents.

Definition 3. In a strategic form game G with vNM preferences, the pure strategy s_i is **strictly dominated for player i** if there exists a mixed strategy $\sigma_i \in \Sigma_i$ such that

$$U_i(\sigma_i, s_{-i}) > U_i(s_i, s_{-i}) \text{ for every } s_{-i} \in S_{-i}. \tag{1}$$

The strategy s_i is **weakly dominated** if there exists a σ_i such that inequality (1) holds with weak inequality, and the inequality is strict for at least one s_{-i}.

Also note that when checking if a pure strategy is dominated by a mixed strategy, we only consider pure strategies for the rest of the players. This is because for a given s_i, the strategy σ_i satisfies (1) for all pure strategies of the opponents if, and only if, it satisfies it for all mixed strategies σ_{-i} as well because player i's payoff when his opponents play mixed strategies is a convex combination of his payoffs when they play pure strategies.

As a first example of a pure strategy dominated by a mixed strategy, consider our favorite card game, whose strategic form, reproduced in Fig. 8 (p. 7), we have derived before. Consider strategy $s_1 = Ff$ for player 1 and the mixed strategy $\sigma_1 = (0.5)[Rr] + (0.5)[Fr]$. We now have:

$$U_1(\sigma_1, m) = (0.5)(0) + (0.5)(0.5) = 0.25 > 0 = U_1(s_1, m)$$
$$U_1(\sigma_1, p) = (0.5)(1) + (0.5)(0) = 0.5 > 0 = U_1(s_1, p).$$

In other words, playing σ_1 yields a higher expected payoff than s_1 does against any possible strategy for player 2. Therefore, s_1 is strictly dominated by σ_1, and we should not expect player 1 to play s_1. On the other hand, the strategy Fr only weakly dominates Ff because it yields a strictly better payoff against m but the same payoff against p. Eliminating weakly dominated strategies is much more controversial than eliminating strictly dominated ones (we shall see why in the homework).

In general, if σ_i strictly dominates s_i and $\sigma_i(s_i) = 0$, then we can eliminate s_i. Note that in addition to strict dominance, we also require that the strictly dominant mixed strategy assigns zero probability to the strictly dominated pure strategy before we can eliminate that pure strategy. The reason for that should be clear: if this were not the case, then we would be eliminating a pure strategy with a mixed strategy, which assumes that this pure strategy would actually be played. Of course, if we eliminate s_i, then this can no longer be the case—we are, in effect, eliminating all mixed strategies that have s_i in their supports as well.

		Player 2	
		m	p
Player 1	Rr	$0, 0$	$1, -1$
	Rf	$-0.5, 0.5$	$1, -1$
	Fr	$0.5, -0.5$	$0, 0$
	Ff	$0, 0$	$0, 0$

Figure 8: The Strategic Form of the Myerson Card Game.

As an example of iterated elimination of strictly dominated strategies that can involve mixed strategies, consider the game in Fig. 9 (p. 7). Obviously, we cannot eliminate any of the pure strategies for player 1 with a mixed strategy (because it would have to include both pure strategies in its support). Also, no pure strategy for player 1 strictly dominates the other: whereas U does better than D against L and C, it does worse against R. Furthermore, none of the pure strategies strictly dominates any other strategy for player 2: L does better than C (or R) against U but worse against D; similarly, R does better than C against U but worse against D.

		Player 2		
		L	C	R
Player 1	U	$2, 3$	$3, 0$	$0, 1$
	D	$0, 0$	$1, 6$	$4, 2$

$\Rightarrow$

	L	C
U	$2, 3$	$3, 0$
D	$0, 0$	$1, 6$

$\Rightarrow$

	L	C
U	$2, 3$	$3, 0$

$\Rightarrow$

	L
U	$2, 3$

Figure 9: Another Game from Myerson (p. 58).

Let's see if we can find a mixed strategy for player 2 that would strictly dominate one of her pure strategies. Which pure strategy should we try to eliminate? It cannot be C because any mixture between L and R would involve a convex combination of 0 and 2 against D which can never exceed 6, which is what C would yield in this case. It cannot be L either because it yields 3 against U, and any mixture between C and R against U would yield at most 1. Hence, let's try to eliminate R: one can imagine mixtures between L and C that would yield a payoff higher than 1 against U and higher than 2 against D. Which mixtures would work? To eliminate R if player 1 chooses U, player 2 should put more than $1/3$ on L. To see this, observe that $3 \times \sigma_2(L) + 0 \times (1 - \sigma_2(L)) > 1$ is required, and this immediately gives $\sigma_2(L) > 1/3$. To eliminate R if player 1 chooses D, it must be the case that $0 \times \sigma_2(L) + 6 \times (1 - \sigma_2(L)) > 2$, which implies $\sigma_2(L) < 2/3$. Hence, any $\sigma_2(L) \in (1/3, 2/3)$ would do the trick. One such mixture would be $\sigma_2(L) = 1/2$. The mixed strategy $\sigma_2 = (0.5)[L] + (0.5)[C] = (0.5, 0.5, 0)$ strictly dominates the pure strategy R. To see this, note that

$$U_2(\sigma_2, U) = (0.5)(3) + (0.5)(0) = 1.5 > 1 = U_2(R, U)$$
$$U_2(\sigma_2, D) = (0.5)(0) + (0.5)(6) = 3 > 2 = U_2(R, D).$$

We can therefore eliminate R, which produces the intermediate game in Fig. 9 (p. 7). In this game, strategy U strictly dominates D because $2 > 0$ against L and $3 > 1$ against C. Therefore, because player 1 knows that player 2 is rational and would never choose R, he can eliminate D from his own choice set. But now player 2's choice is also simple because in the resulting game L strictly dominates C because $3 > 0$. She therefore eliminates this strategy from her choice set. The iterated elimination of strictly dominated strategies leads to a unique prediction as to what rational players should do in this game: $\langle U, L\rangle$.

The following several remarks are useful observations about the relationship between dominance and mixed strategies. Each is easily verifiable by example.

Remark 1. A mixed strategy that assigns positive probability to a dominated pure strategy is itself dominated (by any other mixed strategy that assigns less probability to the dominated pure strategy).

Remark 2. A mixed strategy may be strictly dominated even though it assigns positive probability only to pure strategies that are not even weakly dominated.

		Player 2	
		L	R
Player 1	U	1, 3	−2, 0
	M	−2, 0	1, 3
	D	0, 1	0, 1

Figure 10: Mixed Strategy Dominated by a Pure Strategy.

Consider the example in Fig. 10 (p. 8). Playing U and M with probability $1/2$ each gives player 1 an expected payoff of $-1/2$ regardless of what player 2 does. This is strictly dominated by the pure strategy D, which gives him a payoff of 0, even though neither U nor M is (even weakly) dominated.

Remark 3. A strategy not strictly dominated by any other pure strategy may be strictly dominated by a mixed strategy.

		Player 2	
		L	R
Player 1	U	1, 3	1, 0
	M	4, 0	0, 3
	D	0, 1	3, 1

$\Rightarrow$

	L	R
M	4, 0	0, 3
D	0, 1	3, 1

$\Rightarrow$

	R
M	0, 3
D	3, 1

$\Rightarrow$

	R
D	3, 1

Figure 11: Pure Strategy Dominated by a Mixed Strategy.

Consider the example in Fig. 11 (p. 8). Playing U is not strictly dominated by either M or D and gives player 1 a payoff of 1 regardless of what player 2 does. This is strictly dominated by the mixed strategy in which player 1 chooses M and D with probability $1/2$ each, which would yield 2 if player 2 chooses L and $3/2$ if player 2 chooses R, so it would yield at least $3/2 > 1$ regardless of what player 2 does. Note that there are many possible mixed strategies that would do the trick. To see this, observe that to do better than U against L, it has to be the case that $\sigma_1(M) > 1/4$, and, similarly, to do better than U against R, $\sigma_1(D) > 1/3$. The latter implies that $\sigma_1(M) < 2/3$. Putting the two requirements together yields: $\sigma_1(M) \in (1/4, 2/3)$. Any mixed strategy that satisfies this and includes only M and D in its support will strictly dominate U too. For our purposes, however, all that we need to do is find one mixture that works. Continuing with this example, observe that eliminating U has made R weakly dominant for player 2. If we eliminate L because of that, D will strictly dominate M, and we end up with a unique solution, $\langle D, R \rangle$. In this case, eliminating a weakly dominated strategy will not cause any problems.

The iterated elimination of strictly dominated strategies is quite intuitive but it has a very important drawback. Even though the dominant strategy equilibrium is unique if it exists, for most games that we wish to analyze, all strategies (or too many of them) will survive iterated elimination, and there will be no such equilibrium. Thus, this solution concept will leave many games "unsolvable" in the sense that we shall not be able predict how rational players will play them. In contrast, the concept of Nash equilibrium, to which we turn now, has the advantage that it exists in a very broad class of games.

2 Nash Equilibrium

2.1 Pure-Strategy Nash Equilibrium

Rational players think about actions that the other players might take. In other words, players form beliefs about one another's behavior. For example, in the BoS game, if the man believed the woman would go to the ballet, it would be prudent for him to go to the ballet as well. Conversely, if he believed that the woman would go to the fight, it is probably best if he went to the fight as well. So, to maximize his payoff, he would select the strategy that yields the greatest expected payoff given his belief. Such a strategy is called a **best response** (or *best reply*).

Definition 4. Suppose player i has some belief $s_{-i} \in S_{-i}$ about the strategies played by the other players. Player i's strategy $s_i \in S_i$ is a **best response** if

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for every } s_i' \in S_i.$$

We now define the **best response correspondence**), $BR_i(s_{-i})$, as the set of best responses player i has to s_{-i}. It is important to note that the best response correspondence is set-valued. That is, there may be more than one best response for any given belief of player i. If the other players stick to s_{-i}, then player i can do no better than using any of the strategies in the set $BR_i(s_{-i})$. In the BoS game, the set consists of a single member: $BR_m(F) = \{F\}$ and $BR_m(B) = \{B\}$. Thus, here the players have a single optimal strategy for every belief. In other games, like the one in Fig. 12 (p. 10), $BR_i(s_{-i})$ can contain more than one strategy.

In this game, $BR_1(L) = \{M\}$, $BR_1(C) = \{U, M\}$, and $BR_1(R) = \{U\}$. Also, $BR_2(U) = \{C, R\}$, $BR_2(M) = \{R\}$, and $BR_2(D) = \{C\}$. You should get used to thinking of the best response correspondence as a set of strategies, one for each combination of the other players' strategies. (This is why we enclose the values of the correspondence in braces even when there is only one element.)

		Player 2		
		L	C	R
	U	2,2	1,4	4,4
Player 1	M	3,3	1,0	1,5
	D	1,1	0,5	2,3

Figure 12: The Best Response Game.

We can now use the concept of best responses to define Nash equilibrium: a Nash equilibrium is a strategy profile such that each player's strategy is a best response to the other players' strategies:

DEFINITION 5 (Nash Equilibrium). The strategy profile $(s_i^*, s_{-i}^*) \in S$ is a **pure-strategy Nash equilibrium** if, and only if, $s_i^* \in BR_i(s_{-i}^*)$ for each player $i \in \mathcal{I}$.

An equivalent useful way of defining Nash equilibrium is in terms of the payoffs players receive from various strategy profiles.

DEFINITION 6. The strategy profile (s_i^*, s_{-i}^*) is a **pure-strategy Nash equilibrium** if, and only if, $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$ for each player $i \in \mathcal{I}$ and each $s_i \in S_i$.

That is, for every player i and every strategy s_i of that player, (s_i^*, s_{-i}^*) is at least as good as the profile (s_i, s_{-i}^*) in which player i chooses s_i and every other player chooses s_{-i}^*. In a Nash equilibrium, no player i has an incentive to choose a different strategy when everyone else plays the strategies prescribed by the equilibrium. It is quite important to understand that a **strategy profile is a Nash equilibrium if no player has incentive to deviate from his strategy given that the other players do not deviate**. When examining a strategy for a candidate to be part of a Nash equilibrium (strategy profile), we always hold the strategies of all other players constant.[1]

To understand the definition of Nash equilibrium a little better, suppose there is some player i, for whom s_i is not a best response to s_{-i}. Then, there exists some s_i' such that $u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$. Then this (at least one) player has an incentive to deviate from the theory's prediction and these strategies are not Nash equilibrium.

Another important thing to keep in mind: Nash equilibrium is a strategy profile. Finding a solution to a game involves finding strategy profiles that meet certain rationality requirements. In strict dominance we required that none of the players' equilibrium strategy is strictly dominated. In Nash equilibrium, we require that each player's strategy is a best response to the strategies of the other players.

THE PRISONER'S DILEMMA. By examining all four possible strategy profiles, we see that (D, D) is the unique Nash equilibrium (NE). It is NE because (a) given that player 2 chooses D, then player 1 can do no better than chose D himself $(1 > 0)$; and (b) given that player 1 chooses D, player 2 can do no better than choose D himself. No other strategy profile is NE:

- (C, C) is not NE because if player 2 chooses C, then player 1 can profitably deviate by choosing D $(3 > 2)$. Although this is enough to establish the claim, also note that the

[1]There are several ways to motivate Nash equilibrium. Osborne offers the idea of social convention and Gibbons justifies it on the basis of self-enforcing predictions. Each has its merits and there are others (e.g. steady state in an evolutionary game). You should become familiar with these.

profile is not NE for another sufficient reason: if player 1 chooses C, then player 2 can profitably deviate by playing D instead. (Note that it is enough to show that one player can deviate profitably for a profile to be eliminated.)

- (C, D) is not NE because if player 2 chooses D, then player 1 can get a better payoff by choosing D as well.

- (D, C) is not NE because if player 1 chooses D, then player 2 can get a better payoff by choosing D as well.

Since this exhausts all possible strategy profiles, (D, D) is the unique Nash equilibrium of the game. It is no coincidence that the Nash equilibrium is the same as the strict dominance equilibrium we found before. In fact, as you will have to prove in your homework, a player will never use a strictly dominated strategy in a Nash equilibrium. Further, if a game is dominance solvable, then its solution is the unique Nash equilibrium.

How do we use best responses to find Nash equilibria? We proceed in two steps: First, we determine the best responses of each player, and second, we find the strategy profiles where strategies are best responses to each other.

For example, consider again the game in Fig. 12 (p. 10). We have already determined the best responses for both players, so we only need to find the profiles where each is best response to the other. An easy way to do this in the bi-matrix is by going through the list of best responses and marking the payoffs with a '*' for the relevant player where a profile involves a best response. Thus, we mark player 1's payoffs in (U, C), (U, R), (M, L), and (M, C). We also mark player 2's payoffs in (U, C), (U, R), (M, R), and (D, C). This yields the matrix in Fig. 13 (p. 11).

		Player 2		
		L	C	R
	U	2,2	1*,4*	4*,4*
Player 1	M	3*,3	1*,0	1,5*
	D	1,1	0,5*	2,3

Figure 13: The Best Response Game Marked.

There are two profiles with stars for both players, (U, C) and (U, R), which means these profiles meet the requirements for NE. Thus, we conclude this game has two pure-strategy Nash equilibria.

2.1.1 Diving Money

(Osborne, 38.2) Two players have \$10 to divide. Each names an integer $0 \le k \le 10$. If $k_1 + k_2 \le 10$, each gets k_i. If $k_1 + k_2 > 10$, then (a) if $k_1 < k_2$, player 1 gets k_1 and player 2 gets $10 - k_1$; (b) if $k_1 > k_2$, player 1 gets $10 - k_2$ and player 2 gets k_2; and (c) if $k_1 = k_2$, each player gets \$5.

Instead of constructing 11×11 matrix and using the procedure above, we shall employ an alternative, less cumbersome notation. We draw a coordinate system with 11 marks on each of the abscissa and the ordinate. We then identify the best responses for each player given any of the 11 possible strategies of his opponent. We mark the best responses for player 1 with a circle, and the best responses for player 2 with a smaller disc.

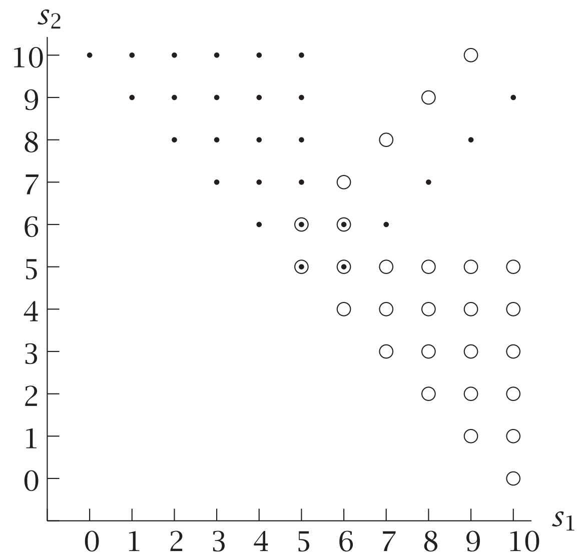

Figure 14: Best Responses in the Dividing Money Game.

Looking at the plot makes clear which strategies are mutual best responses. This game has 4 Nash equilibria in pure strategies: $(5,5)$, $(5,6)$, $(6,5)$, and $(6,6)$. The payoffs in all of these are the same: each player gets \$5.

Alternatively, we know that players never use strictly dominated strategies. Observe now that playing any number less than 5 is strictly dominated by playing 5. To see that, suppose $0 \leq k_1 \leq 4$. There are several cases to consider:

- if $k_2 \leq k_1$, then $k_1 + k_2 < 10$ and player 1 gets k_1; if he plays 5 instead, $5 + k_2 < 10$ and he gets 5, which is better;
- if $k_2 > k_1$ and $k_1 + k_2 > 10$ (which implies $k_2 > 6$), then he gets k_1; if he plays 5 instead, $5 + k_2 > 10$ as well and since $k_2 > k_1$ he gets 5, which is better;
- if $k_2 > k_1$ and $k_1 + k_2 \leq 10$, then he gets k_1; if he plays 5 instead, then:
  - if $5 + k_2 \leq 10$, he gets 5, which is better;
  - if $5 + k_2 > 10$, then $k_1 < k_2$, so he also gets 5, which is better.

In other words, player 1 can guarantee itself a payoff of 5 by playing 5, and any of the strategies that involve choosing a lower number give a strictly lower payoff regardless of what player 2 chooses. A symmetric argument for player 2 establishes that $0 \leq k_2 \leq 4$ is also strictly dominated by choosing $k_2 = 5$. We eliminate these strategies, which leaves a 6×6 payoff matrix to consider (not a bad improvement, we've gone from 121 cells to "only" 36). At this point, we can re-do the plot by restricting it to the numbers above 4 or we can continue the elimination. Observe that $k_i = 10$ is weakly dominated by $k_i = 9$: playing 10 against 10 yields 5 but playing 9 against 10 yields 9; playing 10 against 9 yields 1, but playing 9 against 9 yields 5; playing 10 against any number between 5 and 8 yields the same payoff as playing 9 against that number. If we eliminate 10 because it is weakly dominated by 9, then 9 itself becomes weakly dominated by 8 (that's because the only case where 9 gets a better payoff than 8 is when it's played against 10). Eliminating 9 makes 8 weakly dominated by 7, and eliminating 8 makes 7 weakly dominated by 6. At this point, we've reached a stage where no more elimination can be done. The game is a simple 2×2 shown in Fig. 15 (p. 13).

	\$5	\$6
\$5	5, 5	5, 5
\$6	5, 5	5, 5

Figure 15: The Game after Elimination of Strictly and Weakly Dominated Strategies.

It should be clear from inspection that all four strategy profiles are Nash equilibria. It may appear that IEWDS is not problematic here because we end up with the same solution. However, (unfortunately) this is not the case. Observe that once we eliminate the strictly dominated strategies, we could have also noted that 6 weakly dominates 5. To see this, observe that playing 5 always guarantees a payoff of 5. Playing 6 also gives a payoff of 5 against either 5 or 6 but then gives a payoff of 6 against anything between 7 and 10. Using this argument, we can eliminate 5. We can then apply the IEWDS as before, starting from 10 and working our way down the list until we reach 6. At this point, we are left with a *unique* prediction: $\langle 6, 6 \rangle$. In other words, if we started in this way, we would have missed three of the PSNE. This happens because starting IEWDS at 10 eventually causes 5 to cease to be weakly dominated by 6, so we cannot eliminate it. This also shows that it's quite possible to use weakly dominated strategies in a Nash equilibrium (unlike strictly dominated ones).

Still, the point should be clear even when we restrict ourselves to the safe IESDS: by reducing the game from one with 121 outcomes to one with 36, can save ourselves a lot of analysis with a little bit of thought. Always simplify games (if you can) by finding at least strictly dominated strategies. Going into weakly dominated strategies may or may not be a problem, and you will have to be much more careful there. Usually, it would be too dangerous to do IEWDS because you are likely to miss PSNEs.[2] In this case, you could re-do Fig. 14 (p. 12) with only $s_i \geq 5$ to get all four solutions.

2.1.2 The Partnership Game

There is a firm with two partners. The firm's profit depends on the effort each partner expends on the job and is given by $\pi(x, y) = 4(x + y + cxy)$, where x is the amount of effort expended by partner 1 and y is the amount of effort expended by partner 2. Assume that $x, y \in [0, 4]$. The value $c \in [0, 1/4]$ measures how complementary the tasks of the partners are. Partner 1 incurs a personal cost x^2 of expending effort, and partner 2 incurs cost y^2. Each partner selects the level of his effort independently of the other, and both do so simultaneously. Each partner seeks to maximize their share of the firm's profit (which is split equally) net of the cost of effort. That is, the payoff function for partner 1 is $u_1(x, y) = \pi(x, y)/2 - x^2$, and that for partner 2 is $u_2(x, y) = \pi(x, y)/2 - y^2$.

The strategy spaces here are continuous and we cannot construct a payoff matrix. (Mathematically, $S_1 = S_2 = [0, 4]$ and $\triangle S = [0, 4] \times [0, 4]$.) We can, however, analyze this game using best response functions. Let $\hat{y}$ represent some belief partner 1 has about the other partner's effort. In this case, partner 1's payoff will be $2(x + \hat{y} + cx\hat{y}) - x^2$. We need to maximize this expression with respect to x (recall that we are holding partner's two strategy constant and trying to find the optimal response for partner 1 to that strategy). Taking the

[2] As I did when I improvised IEWDS in this example in class.

derivative yields $2 + 2c\hat{y} - 2x$. Setting the derivative to 0 and solving for x yields the best response $BR_1(\hat{y}) = \{1 + c\hat{y}\}$. Going through the equivalent calculations for the other partner yields his best response function $BR_2(\hat{x}) = \{1 + c\hat{x}\}$.

We are now looking for a strategy profile (x^*, y^*) such that $x^* = BR_1(y^*)$ and $y^* = BR_2(x^*)$. (We can use equalities here because the best response functions produce single values!) To find this profile, we solve the system of equations:

$$x^* = 1 + cy^*$$
$$y^* = 1 + cx^*.$$

The solution is $x^* = y^* = 1/(1-c)$. Thus, this game has a unique Nash equilibrium in pure strategies, in which both partners expend $1/(1-c)$ worth of effort.

2.1.3 Modified Partnership Game

Consider now a game similar to that in the preceding example. Let effort be restricted to the interval $[0, 1]$. Let $p = 4xy$, and let the personal costs be x and y respectively. Thus, $u_1(x, y) = 2xy - x = x(2y - 1)$ and $u_2(x, y) = y(2x - 1)$. We find the best response functions for partner 1 (the other one is the same). If $y < 1/2$, then, since $2y - 1 < 0$, partner 1's best response is 0. If $y = 1/2$, then $2y - 1 = 0$, and partner 1 can choose any level of effort. If $y > 1/2$, then $2y - 1 > 0$, so partner 1's optimal response is to choose 1. This is summarized below:

$$BR_1(y) = \begin{cases} 0 & \text{if } y < 1/2 \\ [0, 1] & \text{if } y = 1/2 \\ 1 & \text{if } y > 1/2 \end{cases}$$

Since $BR_2(x)$ is the same, we can immediately see that there are three Nash equilibria in pure strategies: $(0, 0)$, $(1, 1)$, and $(1/2, 1/2)$ with payoffs $(0, 0)$, $(1, 1)$, and $(0, 0)$ respectively. Let's plot the best response functions, just to see this result graphically in Fig. 16 (p. 15). The three discs at the points where the best response functions intersect represent the three pure-strategy Nash equilibria we found above.

2.2 Strict Nash Equilibrium

Consider the game in Fig. 17 (p. 15). (Its story goes like this. The setting is the South Pacific in 1943. Admiral Kimura has to transport Japanese troops across the Bismarck Sea to New Guinea, and Admiral Kenney wants to bomb the transports. Kimura must choose between a shorter Northern route or a longer Southern route, and Kenney must decide where to send his planes to look for the transports. If Kenney sends the plans to the wrong route, he can recall them, but the number of days of bombing is reduced.)

This game has a unique Nash equilibrium, in which both choose the northern route, (N, N). Note, however, that if Kenney plays N, then Kimura is indifferent between N and S (because the advantage of the shorter route is offset by the disadvantage of longer bombing raids). Still, the strategy profile (N, N) meets the requirements of NE. This equilibrium is not strict.

More generally, an equilibrium is strict if, and only if, each player has a unique best response to the other players' strategies:

DEFINITION 7. A strategy profile (s_i^*, s_{-i}^*) is a **strict Nash equilibrium** if for every player i, $u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_{-i}^*)$ for every strategy $s_i \neq s_i^*$.

The difference from the original definition of NE is only in the strict inequality sign.

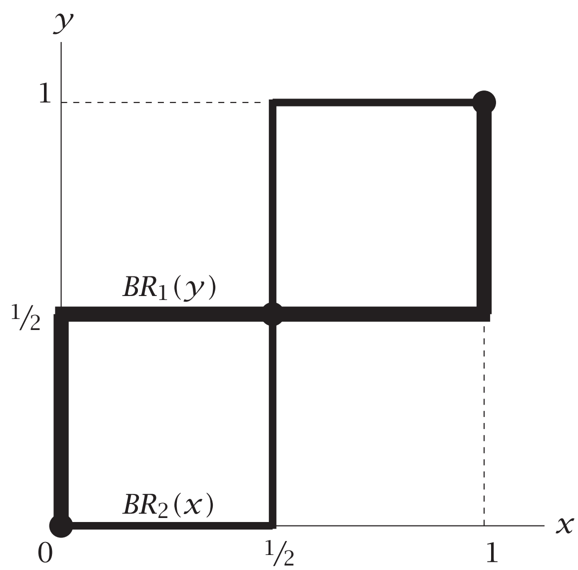

Figure 16: Best Responses in the Modified Partnership Game.

		Kimura N	Kimura S
Kenney	N	$2, -2$	$2, -2$
Kenney	S	$1, -1$	$3, -3$

Figure 17: The Battle of Bismarck Sea.

2.3 Mixed Strategy Nash Equilibrium

The most common example of a game with no Nash equilibrium in pure strategies is MATCHING PENNIES, which is given in Fig. 18 (p. 15).

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

Figure 18: Matching Pennies.

This is a strictly competitive (zero-sum) situation, in which the gain for one player is the loss of the other.[3] This game has no Nash equilibrium in pure strategies. Let's consider mixed strategies.

We first extend the idea of best responses to mixed strategies: Let $BR_i(\sigma_{-i})$ denote player i's best response correspondence when the others play σ_{-i}. The definition of Nash equilibrium is analogous to the pure-strategy case:

DEFINITION 8. A mixed strategy profile σ^* is a **mixed-strategy Nash equilibrium** if, and

[3] It is these zero-sum games that von Neumann and Morgenstern studied and found solutions for. However, Nash's solution can be used in non-zero-sum games, and is thus far more general and useful.

only if, $\sigma_i^* \in BR_i(\sigma_{-i}^*)$.

As before, a strategy profile is a Nash equilibrium whenever all players' strategies are best responses to each other. For a mixed strategy to be a best response, it must put positive probabilities only on pure strategies that are best responses. Mixed strategy equilibria, like pure strategy equilibria, never use dominated strategies.

Turning now to Matching Pennies, let $\sigma_1 = (p, 1-p)$ denote a mixed strategy for player 1 where he chooses H with probability p, and T with probability $1-p$. Similarly, let $\sigma_2 = (q, 1-q)$ denote a mixed strategy for player 2 where she chooses H with probability q, and T with probability $1-q$. We now derive the best response correspondence for player 1 as a function of player 2's mixed strategy.

Player 1's expected payoffs from his pure strategies given player 2's mixed strategy are:

$$\begin{aligned} U_1(H, \sigma_2) &= (1)q + (-1)(1-q) = 2q - 1 \\ U_1(T, \sigma_2) &= (-1)q + (1)(1-q) = 1 - 2q. \end{aligned}$$

Playing H is a best response if, and only if:

$$\begin{aligned} U_1(H, \sigma_2) &\geq U_1(T, \sigma_2) \\ 2q - 1 &\geq 1 - 2q \\ q &\geq 1/2. \end{aligned}$$

Analogously, T is a best response if, and only if, $q \leq 1/2$. Thus, player 1 should choose $p = 1$ if $q \geq 1/2$ and $p = 0$ if $q \leq 1/2$. Note now that whenever $q = 1/2$, player 1 is indifferent between his two pure strategies: choosing either one yields the same expected payoff of 0. Thus, both strategies are best responses, which implies that any mixed strategy that includes both of them in its support is a best response as well. Again, the reason is that if the player is getting the same expected payoff from his two pure strategies, he will get the same expected payoff from any mixed strategy whose support they are.

Analogous calculations yield the best response correspondence for player 2 as a function of σ_1. Putting these together yields:

$$BR_1(q) = \begin{cases} 0 & \text{if } q < 1/2 \\ [0,1] & \text{if } q = 1/2 \\ 1 & \text{if } q > 1/2 \end{cases} \qquad BR_2(p) = \begin{cases} 0 & \text{if } p > 1/2 \\ [0,1] & \text{if } p = 1/2 \\ 1 & \text{if } p < 1/2 \end{cases}$$

The graphical representation of the best response correspondences is in Fig. 19 (p. 17). The only place where the randomizing strategies are best responses to each other is at the intersection point, where each player randomizes between the two strategies with probability $1/2$. Thus, the Matching Pennies game has a unique Nash equilibrium in mixed strategies $\langle \sigma_1^*, \sigma_2^* \rangle$, where $\sigma_1^* = (1/2, 1/2)$, and $\sigma_2^* = (1/2, 1/2)$. That is, where $p = q = 1/2$.

As before, the alternative definition of Nash equilibrium is in terms of the payoff functions. We require that no player can do better by using any other strategy than the one he uses in the equilibrium mixed strategy profile given that all other players stick to their mixed strategies. In other words, the player's expected payoff of the MSNE profile is at least as good as the expected payoff of using any other strategy.

Definition 9. A mixed strategy profile σ^* is a **mixed-strategy Nash equilibrium** if, for all players i,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*) \text{ for all } s_i \in S_i.$$

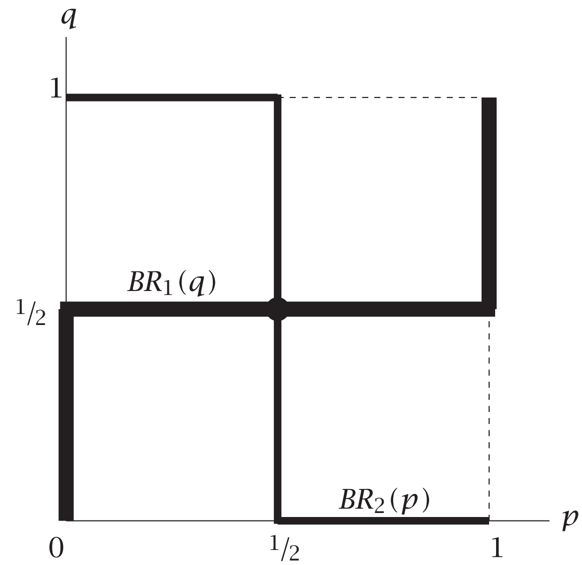

Figure 19: Best Responses in Matching Pennies.

Since expected utilities are linear in the probabilities, if a player uses a non-degenerate mixed strategy in a Nash equilibrium, then he must be indifferent between all pure strategies to which he assigns positive probability. This is why we only need to check for a profitable pure strategy deviation. (Note that this differs from Osborne's definition, which involves checking against profitable mixed strategy deviations.)

2.3.1 Battle of the Sexes

We now analyze the Battle of the Sexes game, reproduced in Fig. 20 (p. 17).

		Player 2	
		F	B
Player 1	F	2,1	0,0
	B	0,0	1,2

Figure 20: Battle of the Sexes.

As a first step, we plot each player's expected payoff from each of the pure strategies as a function of the other player's mixed strategy. Let p denote the probability that player 1 chooses F, and let q denote the probability that player 2 chooses F. Player 1's expected payoff from F is then $2q + 0(1 - q) = 2q$, and his payoff from B is $0q + 1(1 - q) = 1 - q$. Since $2q = 1 - q$ whenever $q = 1/3$, the two lines intersect there.

Looking at the plot in Fig. 21 (p. 18) makes it obvious that for any $q < 1/3$, player 1 has a unique best response in playing the pure strategy B, for $q > 1/3$, his best response is again unique and it is the pure strategy F, while at $q = 1/3$, he is indifferent between his two pure strategies, which also implies he will be indifferent between any mixing of them. Thus, we

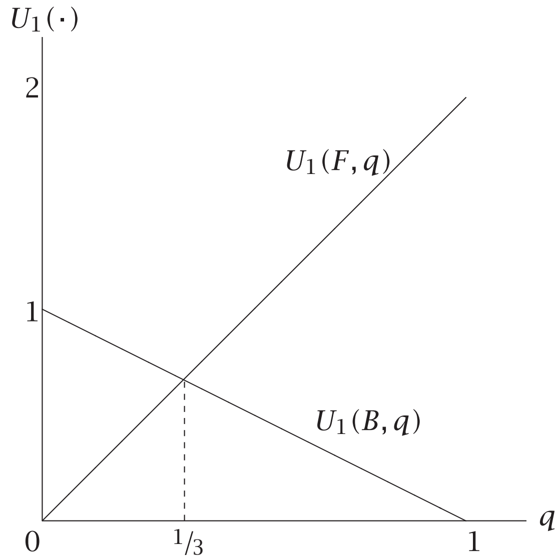

Figure 21: Player 1's Expected Payoffs as a Function of Player 2's Mixed Strategy.

can specify player 1's best response (in terms of p):

$$BR_1(q) = \begin{cases} 0 & \text{if } q < {}^1/_3 \\ [0,1] & \text{if } q = {}^1/_3 \\ 1 & \text{if } q > {}^1/_3 \end{cases}$$

We now do the same for the expected payoffs of player 2's pure strategies as a function of player 1's mixed strategy. Her expected payoff from F is $1p + 0(1-p) = p$ and her expected payoff from B is $0p + 2(1-p) = 2(1-p)$. Noting that $p = 2(1-p)$ whenever $p = {}^2/_3$, we should expect that the plots of her expected payoffs from the pure strategies will intersect at $p = {}^2/_3$. Indeed, Fig. 22 (p. 19) shows that this is the case.

Looking at the plot reveals that player 2 strictly prefers playing B whenever $p < {}^2/_3$, strictly prefers playing F whenever $p > {}^2/_3$, and is indifferent between the two (and any mixture of them) whenever $p = {}^2/_3$. This allows us to specify her best response (in terms of q):

$$BR_2(p) = \begin{cases} 0 & \text{if } p < {}^2/_3 \\ [0,1] & \text{if } p = {}^2/_3 \\ 1 & \text{if } p > {}^2/_3 \end{cases}$$

Having derived the best response correspondences, we can plot them in the $p \times q$ space, which is done in Fig. 23 (p. 19). The best response correspondences intersect in three places, which means there are three mixed strategy profiles in which the two strategies are best responses of each other. Two of them are in pure-strategies: the degenerate mixed strategy profiles $\langle 1,1\rangle$ and $\langle 0,0\rangle$. In addition, there is one mixed-strategy equilibrium,

$$\langle ({}^2/_3[F], {}^1/_3[B]), ({}^1/_3[F], {}^2/_3[B])\rangle .$$

In the mixed strategy equilibrium, each outcome occurs with positive probability. To calculate the corresponding probability, multiply the equilibrium probabilities of each player

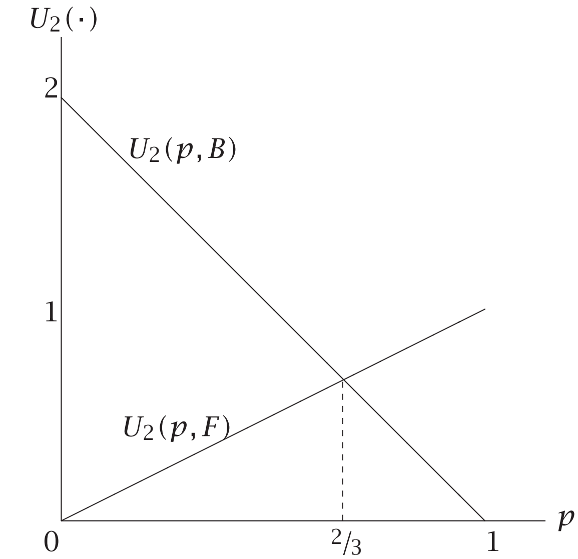

Figure 22: Player 2's Expected Payoffs as a Function of Player 1's Mixed Strategy.

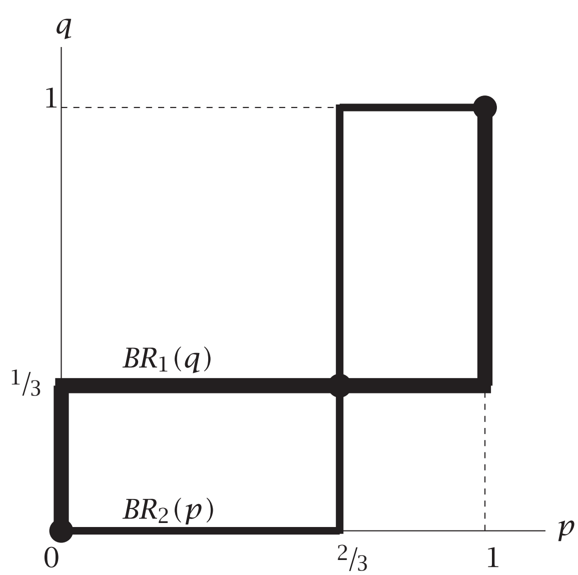

Figure 23: Best Responses in Battle of the Sexes.

choosing the relevant action. This yields $\Pr(F, F) = {}^2/_3 \times {}^1/_3 = {}^2/_9$, $\Pr(B, B) = {}^1/_3 \times {}^2/_3 = {}^2/_9$, $\Pr(F, B) = {}^2/_3 \times {}^2/_3 = {}^4/_9$, and $\Pr(B, F) = {}^1/_3 \times {}^1/_3 = {}^1/_9$. Thus, player 1 and player 2 will meet with probability ${}^4/_9$ and fail to coordinate with probability ${}^5/_9$. Obviously, these probabilities have to sum up to 1. Both players' expected payoff from this equilibrium is $(2)\,{}^2/_9 + (1)\,{}^2/_9 = {}^2/_3$.

2.4 Computing Nash Equilibria

Remember that a mixed strategy σ_i is a best response to σ_{-i} if, and only if, every pure strategy in the support of σ_i is itself a best response to σ_{-i}. Otherwise player i would be able to improve his payoff by shifting probability away from any pure strategy that is not a best response to any that is.

This further implies that in a mixed strategy Nash equilibrium, where σ_i^* is a best response to σ_{-i}^* for all players i, all pure strategies in the support of σ_i^* yield the same payoff when played against σ_{-i}^*, and no other strategy yields a strictly higher payoff. We now use these remarks to characterize mixed strategy equilibria.

REMARK 4. In any finite game, for every player i and a mixed strategy profile σ,

$$U_i(\sigma) = \sum_{s_i \in S_i} \sigma_i(s_i) U_i(s_i, \sigma_{-i}).$$

That is, the player's payoff to the mixed strategy profile is the weighted average of his expected payoffs to all mixed strategy profiles where he plays every one of his pure strategies with a probability specified by his mixed strategy σ_i.

For example, returning to the BoS game, consider the strategy profile $(1/4, 1/3)$. Player 1's expected payoff from this strategy profile is:

$$\begin{aligned} U_1(1/4, 1/3) &= (1/4)\, U_1(F, 1/3) + (3/4)\, U_1(B, 1/3) \\ &= (1/4)\,[(2)\,1/3 + (0)\,2/3] + (3/4)\,[(0)\,1/3 + (1)\,2/3] \\ &= 2/3 \end{aligned}$$

To see that this is equivalent to computing U_1 "directly," observe that the outcome probabilities given this strategy profile are shown in Fig. 24 (p. 20).

	F	B
F	$1/12$	$2/12$
B	$3/12$	$6/12$

Figure 24: Outcome Probabilities for $\langle 1/4, 1/3 \rangle$.

Using these makes computing the expected payoff very easy:

$$U_1(1/4, 1/3) = 1/12(2) + 2/12(0) + 3/12(0) + 6/12(1) = 8/12 = 2/3,$$

which just verifies (for our curiosity) that Remark 4 works as advertised.

The property in Remark 4 allows us to check whether a mixed strategy profile is an equilibrium by examining each player's expected payoffs to his pure strategies only. (Recall that the definition of MSNE I gave you is actually stated in precisely these terms.) Observe in the example above that if player 2 uses her equilibrium mixed strategy and chooses F with probability $1/3$, then player 1's expected payoff from either one of his pure strategies is exactly the same: $2/3$. This is what allows him to mix between them optimally. In general, a player will be willing to randomize among pure strategies only if he is indifferent among them.

PROPOSITION 1. *For any finite game, a mixed strategy profile σ^* is a mixed strategy Nash equilibrium if, and only if, for each player i*

1. $U_i(s_i, \sigma^*_{-i}) = U_i(s_j, \sigma^*_{-i})$ *for all* $s_i, s_j \in \text{supp}(\sigma^*_i)$
2. $U_i(s_i, \sigma^*_{-i}) \geq U_i(s_k, \sigma^*_{-i})$ *for all* $s_i \in \text{supp}(\sigma^*_i)$ *and all* $s_k \notin \text{supp}(\sigma^*_i)$.

That is, the strategy profile σ^* is a MSNE if for every player, the payoff from any pure strategy in the support of his mixed strategy is the same, and at least as good as the payoff from any pure strategy not in the support of his mixed strategy when all other players play their MSNE mixed strategies. In other words, **if a player is randomizing in equilibrium, he must be indifferent among all pure strategies in the support of his mixed strategy**. It is easy to see why this must be the case by supposing that it must not. If he player is not indifferent, then there is at least one pure strategy in the support of his mixed strategy that yields a payoff strictly higher than some other pure strategy that is also in the support. If the player deviates to a mixed strategy that puts a higher probability on the pure strategy that yields a higher payoff, he will strictly increase his expected payoff, and thus the original mixed strategy cannot be optimal; i.e. it cannot be a strategy he uses in equilibrium.

Clearly, a Nash equilibrium that involves mixed strategies cannot be strict because if a player is willing to randomize in equilibrium, then he must have more than one best response. In other words, strict Nash equilibria are always in pure strategies.

We also have a very useful result analogous to the one that states that no player uses a strictly dominated strategy in equilibrium. That is, a dominated strategy is never a best response to any combination of mixed strategies of the other players.

PROPOSITION 2. *A strictly dominated strategy is not used with positive probability in any mixed strategy equilibrium.*

Proof. Suppose that $\langle \sigma^*_1, \sigma^*_{-i} \rangle$ is MSNE and $\sigma^*_1(s_1) > 0$ but s_1 is strictly dominated by s'_1. Suppose first that $\sigma^*_1(s'_1) > 0$ as well. Since both s_1 and s'_1 are used with positive probability in MSNE, it follows that $U_1(s_1, \sigma^*_{-i}) = U_1(s'_1, \sigma^*_{-i})$, which contradicts the fact that s'_1 strictly dominates s_1. Suppose now that $\sigma^*_1(s'_1) = 0$ but then MSNE implies that $U_1(s_1, \sigma^*_{-i}) \geq U_1(s'_1, \sigma^*_{-i})$, which also contradicts the fact that s'_1 strictly dominates s_1. □

This means that when we are looking for mixed strategy equilibria, we can eliminate from consideration all strictly dominated strategies. It is important to note that, as in the case of pure strategies, *we cannot eliminate weakly dominated strategies from consideration* when finding mixed strategy equilibria (because a weakly dominated strategy *can* be used with positive probability in a MSNE).

2.4.1 Myerson's Card Game

The strategic form of the game is given in Fig. 8 (p. 7). It is easy to verify that there are no equilibria in pure strategies. Further, as we have shown, the strategy Ff is strictly dominated, so we can eliminate it from the analysis. The resulting game is shown in Fig. 25 (p. 22).

Let q denote the probability with which player 2 chooses m, and $1 - q$ be the probability with which she chooses p. We now show that in equilibrium player 1 would not play Rf with positive probability. Suppose that $\sigma^*_1(Rf) > 0$; that is, player 1 uses Rf in some MSNE. There are now three possible mixtures that could involve this: (i) $\text{supp}(\sigma^*_1) = \{Rr, Rf, Fr\}$, (ii) $\text{supp}(\sigma^*_1) = \{Rr, Rf\}$, or (iii) $\text{supp}(\sigma^*_1) = \{Rf, Fr\}$.

Let's take (i) and (ii), in which $\sigma^*_1(Rr) > 0$ as well. Since player 1 is willing to mix in equilibrium between (at least) these two pure strategies, it follows that his expected payoff

		Player 2	
		m	p
	Rr	$0, 0$	$1, -1$
Player 1	Rf	$-1/2, 1/2$	$1, -1$
	Fr	$1/2, -1/2$	$0, 0$

Figure 25: The Reduced Strategic Form of the Myerson Card Game.

should be the same no matter which one of them he uses. The expected payoff from Rf is $U_1(Rf, q) = (-1/2)q + (1)(1-q) = 1 - \frac{3}{2}q$, and the expected payoff from Rr is $U_1(Rr, q) = (0)q + (1)(1-q) = 1 - q$. In MSNE, these two have to be equal, so $1 - \frac{3}{2}q = 1 - q$, which implies $5/2q = 0$, or $q = 0$. Hence, in any MSNE in which player 1 puts positive probability on both Rf and Rr requires that $q = 0$; that is, that player 2 chooses p with certainty. This makes intuitive sense, which we can verify by looking at the payoff matrix. Observe that both Rr and Rf give player 1 a payoff of 1 against p but that Rf is strictly worse against m. This implies that should player 2 choose m with positive probability, player 1 will strictly prefer to play Rr. Therefore, player 1 would be willing to randomize between these two pure strategies only if player 2 is expected to choose p for sure.

Given that behavior for player 2, player 1 will never put positive probability on Fr because conditional on player 2 choosing p, Rr and Rf strictly dominate it. In other words, case (i) cannot happen in MSNE.

We now know that if player 1 uses Rr and Rf, he can only do so in case (ii). But if player 1 is certain not to choose Fr, then m strictly dominates p for player 2: $U_2(\sigma_1, m) = 1/2\sigma_1(Rf) > -1 = U_2(\sigma_1, p)$ for any strategy in (ii). This now implies that $q = 1$ because player 2 is certain to choose m. But this contradicts $q = 0$ which we found has to hold for any equilibrium mixed strategy that puts positive weight on both Rr and Rf. Hence, it cannot be the case that player 1 plays (ii) in MSNE either.

This leaves one last possibility to consider, so suppose he puts positive probability on Rf and Fr. Since he is willing to mix, it has to be the case that $U_1(Rf, \sigma_2^*) = U_1(Fr, \sigma_2^*)$. We know that the expected payoffs are $U_1(Rf, \sigma_2^*) = -1/2q + (1-q) = 1/2q = U_1(Fr, \sigma_2^*)$, which implies $q = 1/2$. That is, if player 1's equilibrium mixed strategy is of type (iii), then player 2 must mix herself, and she must do so precisely with probability $1/2$. However, this now implies that $U_1(Rf, 1/2) = 1/4 < 1/2 = U_1(Rr, 1/2)$. That is, player 1's expected payoff from the strategy Rr, which he is not supposed to be using, is strictly higher than the payoff from the pure strategies in the support of the mixed strategy. This means that player 1 will switch to Rr, which implies that case (iii) cannot occur in MSNE either. We conclude that there exists no MSNE in which player 1 puts positive probability on Rf.

In this particular case, you can also observe that Rr strictly dominates Rf for any mixed strategy for player 2 that assigns positive probability to m. Since we know that player 2 must mix in equilibrium, it follows that player 1 will never play Rf with positive probability in any equilibrium. Thus, we can eliminate that strategy. Note that although Rr weakly dominates Rf, this is not why we eliminate Rf. Instead, we are making an equilibrium argument and proving that Rf will never be chosen in any equilibrium with positive probability.

So, any Nash equilibrium must involve player 1 mixing between Rr and Fr. Since he will never play Rf in equilibrium, we can eliminate this strategy from consideration altogether, leaving us with the simple 2×2 game shown in Fig. 26 (p. 23). Let s be the probability of choosing Rr, and $1 - s$ be the probability of choosing Fr.

		Player 2	
		m	p
Player 1	Rr	$0,0$	$1,-1$
	Fr	$^1/_2,-^1/_2$	$0,0$

Figure 26: The Simplified Myerson Game.

Let's find the MSNE for this one. Observe that now we do not have to worry about partially mixed strategies: since each player has only two pure strategies each, any mixture must be complete. Hence, we only need equate the payoffs to find the equilibrium mixing probabilities. Because player 1 is willing to mix, the expected payoffs from the two pure strategies must be equal. Thus, $(0)q + (1)(1-q) = {}^1/_2 q + (0)(1-q)$, which implies that $q = {}^2/_3$. Since player 2 must be willing to randomize as well, her expected payoffs from the pure strategies must also be equal. Thus, $(0)s + -{}^1/_2(1-s) = (-1)s + (0)(1-s)$, which implies that $s = {}^1/_3$. We conclude that the unique mixed strategy Nash equilibrium of the card game: is

$$\left\langle \left(\sigma_1^*(Rr) = {}^1/_3, \sigma_1^*(Fr) = {}^2/_3\right), \left(\sigma_2^*(m) = {}^2/_3, \sigma_2^*(p) = {}^1/_3\right)\right\rangle.$$

That is, player 1 raises for sure if he has a red (winning) card, and raises with probability $^1/_3$ if he has a black (losing) card. Player 2 meets with probability $^2/_3$ when she sees player 1 raise in equilibrium. The expected payoff in this unique equilibrium for player 1 is:

$$({}^1/_2)\left[{}^2/_3(2) + {}^1/_3(1)\right] + ({}^1/_2)\left[{}^1/_3\left({}^2/_3(-2) + {}^1/_3(1)\right) + {}^2/_3(-1)\right] = {}^1/_3,$$

and the expected payoff for player 2, computed analogously, is $-{}^1/_3$. If you are risk-neutral, you should only agree to take player 2's role if offered a pre-play bribe of at least \$0.34 because you expect to lose \$0.33.

Let's think a bit about the intuition behind this MSNE. First, note that player 2 cannot meet or pass with certainty in any equilibrium. If she passed whenever player 1 raised, then player 1 would raise even when he has a losing card. But if that's true, then raising would not tell player 2 anything about the color of the card, and so she expects a 50-50 chance to win if she meets. With these odds, she is better off meeting: her expected payoff would be 0 if she meets (50% chance of winning \$2 and 50% of losing the same amount). Passing, on the other hand, guarantees her a payoff of -1. Of course, if she met with certainty, then player 1 would never raise if he has the losing card. This now means that whenever player 1 raises, player 2 would be certain that he has the winning card, but in this case she surely should not meet: passing is much better with a payoff of -1 versus a truly bad loss of -2. So it has got to be the case that player 2 mixes.

Second, we have seen that player 1 cannot raise without regard for the color of the card in any equilibrium: if he did that, player 2 would meet with certainty, but in that case it is better to fold with a losing card. Conversely, player 1 cannot fold regardless of the color because no matter what player 2 does, raising with a winning card is always better. Hence, we conclude that player 1 must raise for sure if he has the winning card. But to figure out the probability with which he must bluff, we need to calculate the probability with which player 2 will meet a raise. It is these two probabilities that the MSNE pins down.

Intuitively, upon seeing player 1 raise, player 2 would still be unsure about the color of the card, although she would have an updated estimate of that probability of winning. She should become more pessimistic if player 1 raises with a strictly higher probability on a winning card. Hence, she would use this new probability of victory to decide her mixture. Bayes Rule will give you precisely this updated probability:

$$\begin{aligned}
\Pr[\text{black}|\text{1 raises}] &= \frac{\Pr[\text{1 raises}|\text{black}] \times \Pr[\text{black}]}{\Pr[\text{1 raises}|\text{black}] \times \Pr[\text{black}] + \Pr[\text{1 raises}|\text{red}] \times \Pr[\text{red}]} \\
&= \frac{\sigma_1(Rr)\,(1/2)}{\sigma_1(Rr)\,(1/2) + (1)\,(1/2)} = \frac{(1/3)\,(1/2)}{(1/3)\,(1/2) + (1)\,(1/2)} \\
&= 1/4.
\end{aligned}$$

In other words, upon seeing player 1 raise, player 2 revises her probability of winning (the card being black) from $1/2$ down to $1/4$. Given this probability, what should her best response be? The expected payoff from meeting under these new odds is $1/4(2) + 3/4(-2) = -1$, which is the same as her payoff from passing. This should not be surprising: player 1's mixing probability must be making her indifferent if she is willing to mix. For her part, she must choose the mixture that makes player 1 willing to mix between his two pure strategies, and this mixture is to meet with probability $2/3$. That is, *player 1's mixed strategy makes player 2 indifferent, which is required if she is to mix in equilibrium. Conversely, her strategy must be making player 1 indifferent between his pure strategies, so he is willing to mix too.*

It is important to note that player 1 is not mixing *in order* to make player 2 indifferent between meeting and passing: instead this is a *feature* (or requirement) of optimal play. To see that, suppose that his strategy did not make her indifferent, then she would either meet or pass for sure, depending on which one is better for her. But as we have just seen, playing a pure-strategy cannot be optimal because of the effect it will have on player 1's behavior. Therefore, optimality itself requires that player 1's behavior will make her indifferent. In other words, players are not looking to enure that their opponents are indifferent so that they would play the appropriate mixed strategy. Rather, their own efforts to find an optimal strategy render their opponents indifferent.[4]

By the way, you have just solved an incomplete information signaling game! Recall that in the original description, player 1 sees the color of the card (so he is privately informed about it) and can "signal" this to player 2 through his behavior. Observe that his action does reveal some, but not all, information: after seeing him raise, player 2 updates to believe that her probability of winning is worse than random chance. We shall see this game again when we solve more games of incomplete information and we shall find this MSNE is also the perfect Bayesian equilibrium. For now, aren't you glad that on the first day you learn what a Nash equilibrium is, you get to solve a signaling game which most introductory classes wouldn't even teach?

[4] This does not mean that there isn't a philosophical problem here: if a player is indifferent among several pure strategies, then there appears to be no compelling reason to expect him to choose the "right" (equilibrium) mixture that would rationalize his opponent's strategy. Clearly, any deviation from the equilibrium mixture cannot be supported if the other player guesses it—she will simply best-respond by playing the strategy that becomes better for her. That's why any other non-equilibrium mixture cannot be supported as a part of equilibrium: if it were a part of equilibrium, then the opponent will know it and expect it, but if this were true, she will readjust her play accordingly. The question is: if a player is indifferent among his pure strategies, then how would his opponent guess which "deviating" mixture he may choose? This is obviously a problem in a single-shot encounter when the indifferent player may simply pick a mixture at random (or even choose a pure strategy directly); after all, he is indifferent. In that case, there may be no compelling reason to expect behavior that resembles Nash equilibrium. Pure-strategy Nash equilibria, especially the strict ones, are more compelling in that respect. However, Harsanyi's purification argument (which I mentioned in class but which we shall see in action soon) gets neatly around this problem because in that interpretation, there is no actual randomization.

2.4.2 Another Simple Game

To illustrate the algorithm for solving strategic form games, we now go through a detailed example using the game from Myerson, p. 101, reproduced in Fig. 27 (p. 25). The algorithm for finding all Nash equilibria involves (a) checking for solutions in pure strategies, and (b) checking for solutions in mixed strategies. Step (b) is usually the more complicated one, especially when there are many pure strategies to consider. You will need to make various guesses, use insights from dominance arguments, and utilize the remarks about optimal mixed strategies here.

		Player 2		
		L	M	R
Player 1	U	7, 2	2, 7	3, 6
	D	2, 7	7, 2	4, 5

Figure 27: A Strategic Form Game.

We begin by looking for pure-strategy equilibria. U is only a best response to L, but the best response to U is M. There is no pure-strategy equilibrium involving player 1 choosing U. On the other hand, D is a best response to both M and R. However, only L is a best response to D. Therefore, there is no pure-strategy equilibrium with player 1 choosing D for sure. This means that any equilibrium must involve a mixed strategy for player 1 with $\text{supp}(\sigma_1) = \{U, D\}$. In other words, player 1 must mix in any equilibrium. Turning now to player 2's strategy, we note that there can be no equilibrium with player 2 choosing a pure strategy either. This is because player 1 has a unique best response to each of her three strategies, but we have just seen that player 1 must be randomizing in equilibrium.

We now have to make various guesses about the support of player 2's strategy. We know that it must include at least two of her pure strategies, and perhaps all three. There are four possibilities to try.

- $\text{supp}(\sigma_2) = \{L, M, R\}$. Since player 2 is willing to mix, she must be indifferent between her pure strategies, and therefore:

$$2\sigma_1(U) + 7\sigma_1(D) = 7\sigma_1(U) + 2\sigma_1(D) = 6\sigma_1(U) + 5\sigma_1(D).$$

  We require that the mixture is a valid probability distribution, or $\sigma_1(U) + \sigma_1(D) = 1$. Note now that $2\sigma_1(U) + 7\sigma_1(D) = 7\sigma_1(U) + 2\sigma_1(D) \Rightarrow \sigma_1(U) = \sigma_1(D) = 1/2$. However, $7\sigma_1(U) + 2\sigma_1(D) = 6\sigma_1(U) + 5\sigma_1(D) \Rightarrow \sigma_1(U) = 3\sigma_1(D)$, a contradiction. Therefore, there can be no equilibrium that includes all three of player 2's strategies in the support of her mixed strategy.

- $\text{supp}(\sigma_2) = \{M, R\}$. Since player 1 is willing to mix, it must be the case that $2\sigma_2(M) + 3\sigma_2(R) = 7\sigma_2(M) + 4\sigma_2(R) \Rightarrow 0 = 5\sigma_2(M) + \sigma_2(R)$, which is clearly impossible because both $\sigma_2(M) > 0$ and $\sigma_2(R) > 0$. Hence, there can be no equilibrium where player 2's support consists of M and R. (You can also see this by inspecting the payoff matrix: if player 2 is choosing only between M and R, then D strictly dominates U for player 1. This means that player 1's best response will be D but we already know that he must be mixing, a contradiction.)[5]

[5] Alternatively, you could simply observe that if player 2 never chooses L, then D strictly dominates U for player 1. But if he is certain to choose D, then player 2 strictly prefers to play L, a contradiction.

- $\text{supp}(\sigma_2) = \{L, M\}$. Because player 1 is willing to mix, it follows that $7\sigma_2(L) + 2\sigma_2(M) = 2\sigma_2(L) + 7\sigma_2(M) \Rightarrow \sigma_2(L) = \sigma_2(M) = 1/2$. Further, because player 2 is willing to mix, it follows that $2\sigma_1(U) + 7\sigma_1(D) = 7\sigma_1(U) + 2\sigma_1(D) \Rightarrow \sigma_1(U) = \sigma_1(D) = 1/2$.

  So far so good. We now check for profitable deviations. If player 1 is choosing each strategy with positive probability, then choosing R would yield player 2 an expected payoff of $(1/2)(6) + (1/2)(5) = 11/2$. Thus must be worse than any of the strategies in the support of her mixed strategy, so let's check M. Her expected payoff from M is $(1/2)(7) + (1/2)(2) = 9/2$. That is, the strategy which she is sure not to play yields an expected payoff strictly higher than any of the strategies in the support of her mixed strategy. Therefore, this cannot be an equilibrium either.

- $\text{supp}(\sigma_2) = \{L, R\}$. Since player 1 is willing to mix, it follows that $7\sigma_2(L) + 3\sigma_2(R) = 2\sigma_2(L) + 4\sigma_2(R) \Rightarrow 5\sigma_2(L) = \sigma_2(R)$, which in turn implies $\sigma_2(L) = 1/6$, and $\sigma_2(R) = 5/6$. Further, since player 2 is willing to mix, it follows that $2\sigma_1(U) + 7\sigma_1(D) = 6\sigma_1(U) + 5\sigma_1(D) \Rightarrow \sigma_1(D) = 2\sigma_1(U)$, which in turn implies $\sigma_1(U) = 1/3$, and $\sigma_1(D) = 2/3$.

  Can player 2 do better by choosing M? Her expected payoff would be $(1/3)(7) + (2/3)(2) = 11/3$. Any of the pure strategies in the support of her mixed strategy yields an expected payoff of $(1/3)(2) + (2/3)(7) = (1/3)(6) + (2/3)(5) = 16/3$, which is strictly better. Therefore, the mixed strategy profile:

  $$\langle (\sigma_1(U) = 1/3, \sigma_1(D) = 2/3), (\sigma_2(L) = 1/6, \sigma_2(R) = 5/6) \rangle$$

  is the unique Nash equilibrium of this game. The expected equilibrium payoffs are $11/3$ for player 1 and $16/3$ for player 2.

This exhaustive search for equilibria may become impractical when the games become larger (either more players or more strategies per player). There are programs, like the late Richard McKelvey's *Gambit*, that can search for solutions to many games.

2.4.3 Choosing Numbers

Players 1 and 2 each choose a positive integer up to K. Thus, the strategy spaces are both $\{1, 2, \ldots, K\}$. If the players choose the same number then player 2 pays \$1 to player 1, otherwise no payment is made. Each player's preferences are represented by his expected monetary payoff. The claim is that the game has a mixed strategy Nash equilibrium in which each player chooses each positive integer with equal probability.[6]

It is easy to see that this game has no equilibrium in pure strategies: If the strategy profile specifies the same numbers, then player 2 can profitably deviate to any other number; if the strategy profile specifies different numbers, then player 1 can profitably deviate to the number that player 2 is naming. However, this is a finite game, so Nash's Theorem tells us there must be an equilibrium. Thus, we know we should be looking for one in mixed strategies.

The problem here is that there is an infinite number of potential mixtures we have to consider. We attack this problem methodically by looking at types of mixtures instead of individual ones.

[6]It is not clear how you get to this claim. This is the part of game theory that often requires some inspired guesswork and is usually the hardest part. Once you have an idea about an equilibrium, you can check whether the profile is one. There is usually no mechanical way of finding an equilibrium.

STRAIGHT LOGIC. Let's prove that players must put positive probability on each possible number in equilibrium, and that they must play each number with exactly the same probability. Suppose, to the contrary, that player 1 does not play some number, say z, with positive probability. Then player 2's best response is to play z for sure, so she will not mix. However, given that she will choose z for sure, player 1 is certain to deviate and play z for sure himself. Therefore, player 1 must put positive probability on all numbers. But if player 1 mixes over all numbers, then so must player 2. To see this, suppose to the contrary that she does not and instead plays some number, say y, with probability zero. But then player 1 can do strictly better by redistributing the positive weight he attaches to y to the numbers which player 2 chooses with positive probability, a contradiction to the fact that player 1 must mix over all numbers in equilibrium. Therefore, both players must mix over all numbers. But this now implies that they must be choosing each number with the same probability. To see this, recall that since they are mixing, they must be indifferent among their pure strategies. The only way player 1 will be indifferent among his choices is when player 2 chooses each number in the support of her mixed strategy with the same probability. If that were not true and she chose some numbers with high probability, then playing these numbers would give player 1 an expected payoff higher than playing any of the other numbers, a contradiction of the equilibrium supposition. If player 1 himself chose some numbers with higher probability, then choosing any numbers other than these would give player 2 a strictly better payoff, a contradiction too. Hence, both players must randomize over all numbers and they must assign equal probabilities to them. There is only one way to do this: they pick each number with probability $^1/_K$.

Let's verify that this is MSNE by applying Proposition 1. Since all strategies are in the support of this mixed strategy, it is sufficient to show that each strategy of each player results in the same expected payoff. (That is, we only use the first part of the proposition.) Player 1's expected payoff from each pure strategy is $^1/_K(1) + (1 - {}^1/_K)(0) = {}^1/_K$ because player 2 chooses the same number with probability $^1/_K$ and a different number with the complementary probability. Similarly, player 2's expected payoff is $^1/_K(-1) + (1 - {}^1/_K)(0) = -{}^1/_K$. Thus, this strategy profile is a mixed strategy Nash equilibrium.

MORE INVOLVED PROOF. This essentially replicates the above argument but in a slightly more roundabout way and uses more notation. We want to show that the MSNE with each player choosing any number with probability $^1/_K$ is the unique MSNE.

Let (σ_1^*, σ_2^*) be a MSNE where $\sigma_i^*(k)$ is the probability that player i's mixed strategy assigns to the integer k. Given that player 2 uses σ_2^*, player 1's expected payoff to choosing the number k is $\sigma_2^*(k)$. From Proposition 1, if $\sigma_1^*(k) > 0$, then $\sigma_2^*(k) \geq \sigma_2^*(j)$ for all numbers j. That is, if player 1 assigns positive probability to choosing some number k, then the equilibrium probability with which player 2 is choosing this number must be at least as great as the probability of any other number. If this were not the case, it would mean that there exists some number $m \neq k$ which player 2 chooses with a higher probability in equilibrium. But in that case, player 1's equilibrium strategy would be strictly dominated by the strategy that chooses m with a higher probability (because it would yield player 1 a higher expected payoff). Therefore, the mixed strategy σ_1^* could not be optimal, a contradiction.

Since $\sigma_2^*(j) > 0$ for at least some j, it follows that $\sigma_2^*(k) > 0$. That is, because player 2 must choose some number, she must assign a strictly positive probability to at least one number from the set. But because the equilibrium probability of choosing k must be at least as high as the probability of any other number, the probability of k must be strictly positive.

We conclude that in equilibrium, if player 1 assigns positive probability to some arbitrary number k, then player 2 must do so as well.

Now, player 2's expected payoff if she chooses k is $-\sigma_1^*(k)$, and since $\sigma_2^*(k) > 0$, it must be the case that $\sigma_1^*(k) \leq \sigma_1^*(j)$ for all j. This follows from Proposition 1. To see this, note that if this did not hold, it would mean that player 1 is choosing some number m with a strictly lower probability in equilibrium. However, in this case player 2 could do strictly better by switching to a strategy that picks m because the expected payoff would improve (the numbers are less likely to match). But this contradicts the optimality of player 2's equilibrium strategy σ_2^*.

What is the largest equilibrium probability with which k is chosen by player 1? We know that it cannot exceed the probability assigned to any other number. Because there are K numbers, this means that it cannot exceed $1/K$. To see this, note that if there was some number to which the assigned probability was strictly greater than $1/K$, then there must be some other number with a probability strictly smaller than $1/K$, and then $\sigma_1^*(k)$ would have to be no greater than that smaller probability. We conclude that $\sigma_1^*(k) \leq 1/K$.

We have now shown that if in equilibrium player 1 assigns positive probability to some arbitrary number k, it follows that this probability cannot exceed $1/K$. Hence, the equilibrium probability of choosing any number to which player 1 assigns positive probability cannot exceed $1/K$.

But this now implies that player 1 must assign $1/K$ to each number and mix over all available numbers. Suppose not, which would mean that player 1 is mixing over $n < K$ numbers. From the proof above, we know that he cannot assign more than $1/K$ probability to each of these n numbers. But because his strategy must be a valid probability distribution, the individual probabilities must sum up to 1. In this case, the sum up to $n/K < 1$ because $n < K$. The only way to meet the requirement would be to assign at least one of the numbers a strictly larger probability, a contradiction. Therefore, $\sigma_1^*(k) = 1/K$ for all k.

A symmetric argument establishes the result for player 2. We conclude that there are no other mixed strategy Nash equilibria in this game.

2.4.4 Defending Territory

General A is defending territory accessible by 2 mountain passes against General B. General A has 3 divisions at his disposal and B has 2. Each must allocate divisions between the two passes. A wins the pass if he allocates at least as many divisions to it as B does. A successfully defends his territory if he wins at both passes.

General A has four strategies at his disposal, depending on the number of divisions he allocates to each pass: $S_A = \{(3,0), (2,1), (1,2), (0,3)\}$. General B has three strategies he can use: $S_B = \{(2,0), (1,1), (0,2)\}$. We construct the payoff matrix as shown in Fig. 28 (p. 28).

		General B		
		$(2,0)$	$(1,1)$	$(0,2)$
General A	$(3,0)$	$1,-1$	$-1,1$	$-1,1$
	$(2,1)$	$1,-1$	$1,-1$	$-1,1$
	$(1,2)$	$-1,1$	$1,-1$	$1,-1$
	$(0,3)$	$-1,1$	$-1,1$	$1,-1$

Figure 28: Defending Territory.

This is a strictly competitive game, which (not surprisingly) has no pure strategy Nash equilibrium. Thus, we shall be looking for MSNE. Denote a mixed strategy of General A by (p_1, p_2, p_3, p_4), and a mixed strategy of General B by (q_1, q_2, q_3).

First, suppose that in equilibrium $q_2 > 0$. Since General A's expected payoff from his strategies $(3, 0)$ and $(0, 3)$ are both less than any of the other two strategies, it follows that in such an equilibrium $p_1 = p_4 = 0$. In this case, General B's expected payoff to his strategy $(1, 1)$ is then -1. However, either one of the other two available strategies would yield a higher expected payoff. Therefore, $q_2 > 0$ cannot occur in equilibrium.

What is the intuition behind this result? Observe that the strategy $(1, 1)$ involves General B dividing his forces and sending one division to each pass. However, this would enable General A to defeat both of them for sure: he would send 2 divisions to one pass, and 1 division to the other. That is, he would play either $(2, 1)$ or $(1, 2)$ but in either case, General B would lose for sure. Given that at least one of the passes will be defended by 1 division, General B would do strictly better by attacking a pass in full force: he would lose if he happens to attack the pass defended by 2 divisions but would win if he happens to attack the pass defended by a single division. Thus, he would deviate from the strategy $(1, 1)$, so it cannot occur in equilibrium.

We conclude that in equilibrium General B must attack in full force one of the passes. Note now that he must not allow General A to guess which pass will be attack in that way. The only way to do so is to attack each with the same probability. If this were not the case and General B attacked one of the passes with a higher probability, then General A's best response would be to defend that pass with at least 2 divisions for sure. But then General B would strictly prefer to attack the other pass. Hence, in equilibrium it has to be the case that General B attacks both passes with probability $q_1 = q_2 = 1/2$.

Continuing with this logic, since General A now expects a full-scale attack on each pass with equal probability, he knows for sure that he will lose the war with probability $1/2$. This is so because there is no way to defend both passes simultaneously against a full-scale attack. The allocations $(3, 0)$ and $(2, 1)$ leave the second pass vulnerable if General B happens to choose it, and the allocations $(1, 2)$ and $(0, 3)$ leave the first pass vulnerable. Hence, General A's best bet is to choose between these two combinations with equal probability. That is, he can defend successfully the first pass and lose the second with the allocations $(3, 0)$ and $(2, 1)$, and defend successfully the second pass and lose the first with the allocations $(1, 2)$ and $(0, 3)$. Using our notation, his strategy would be to play $p_1 + p_2 = p_3 + p_4 = 1/2$.

This, however, is not enough to pin down equilibrium strategies. Observe that if General A plays $(0, 3)$ and $(3, 0)$ with high probabilities, then General B can attempt to split his forces: doing so would give him an opportunity to sneak 1 division through an undefended pass. But we already know that $(1, 1)$ cannot be an equilibrium strategy. This implies that in equilibrium General A must not be too likely to leave a pass undefended. Since, as we have seen, General B will launch a full-scale attack on each of the passes with equal probability, his expected payoff is 0: given General A's strategy, he will win with probability $1/2$ and lose with probability $1/2$. Dividing his forces should not improve upon that expectation. This will be so if the overall probability of General A leaving a pass undefended is no greater than $1/2$. That is, $p_1 + p_4 \leq 1/2$. If that were not so, then General B would divide his forces and win with probability greater than $1/2$, a contradiction to the equilibrium supposition that he is equally likely to win and lose. Thus, we conclude that the game has infinitely many MSNE. In all of these General B attacks each of the passes in full strength with equal probability: $q_1 = q_3 = 1/2$. General A, on the other hand is equally likely to prevail at either pass: $p_1 + p_2 = p_3 + p_4 = 1/2$, and not too likely to leave a pass undefended: $p_1 + p_4 \leq 1/2$.

More formally, given that in any equilibrium $q_2 = 0$, what probabilities would B assign to the other two strategies in equilibrium? Since $q_2 = 0$, it follows that $q_3 = 1 - q_1$. General A's expected payoff to $(3,0)$ and $(2,1)$ is $2q_1 - 1$, and the payoff to $(1,2)$ and $(0,3)$ is $1 - 2q_1$. If $q_1 < 1/2$, then in any equilibrium $p_1 = p_2 = 0$. In this case, B has a unique best response, which is $(2,0)$, which implies that in equilibrium $q_1 = 1$. But if this is the case, then either of A's strategies $(3,0)$ or $(2,1)$ yields a higher payoff than any of the other two, contradicting $p_1 = p_2 = 0$. Thus, $q < 1/2$ cannot occur in equilibrium. Similarly, $q_1 > 1/2$ cannot occur in equilibrium. This leaves $q_1 = q_3 = 1/2$ to consider.

If $q_1 = q_3 = 1/2$, then General A's expected payoffs to all his strategies are equal. We now have to check whether General B's payoffs from this profile meet the requirements of Proposition 1. That is, we have to check whether the payoffs from $(2,0)$ and $(0,2)$ are the same, and whether this payoff is at least as good as the one to $(1,1)$. The first condition is:

$$\begin{aligned} -p_1 - p_2 + p_3 + p_4 &= p_1 + p_2 - p_3 - p_4 \\ p_1 + p_2 &= p_3 + p_4 = 1/2 \end{aligned}$$

General B's expected payoff to $(2,0)$ and $(0,2)$ is then 0, so the first condition is met. Note now that since $p_1 + p_2 + p_3 + p_4 = 1$, we have $1 - (p_1 + p_4) = p_2 + p_3$. The second condition is:

$$\begin{aligned} p_1 - p_2 - p_3 + p_4 &\leq 0 \\ p_1 + p_4 &\leq p_2 + p_3 \\ p_1 + p_4 &\leq 1 - (p_1 + p_4) \\ p_1 + p_4 &\leq 1/2 \end{aligned}$$

Thus, we conclude that the set of mixed strategy Nash equilibria in this game is the set of strategy profiles:

$$\big((p_1, 1/2 - p_1, 1/2 - p_4, p_4), (1/2, 0, 1/2)\big) \text{ where } p_1 + p_4 \leq 1/2.$$

This, of course, is precisely what we found with less algebra above. (But the algebra does make it very easy.)

2.4.5 Choosing Two-Thirds of the Average

(Osborne, 34.1) Each of 3 players announces an integer from 1 to K. If the three integers are different, the one whose integer is closest to $2/3$ of the average of the three wins \$1. If two or more integers are the same, \$1 is split equally between the people whose integers are closest to $2/3$ of the average.

Formally, $N = \{1, 2, 3\}$, $S_i = \{1, 2, \ldots, K\}$, and $\triangle S = S_1 \times S_2 \times S_3$. There are K^3 different strategy profiles to examine, so instead we analyze *types* of profiles.

Suppose all three players announce the same number $k \geq 2$. Then $2/3$ of the average is $2/3k$, and each gets $\$1/3$. Suppose now one of the players deviates to $k - 1$. Now $2/3$ of the average is $2/3k - 2/9$. We now wish to show that the player with $k - 1$ is closer to the new $2/3$ of the average than the two whose integers where k:

$$\begin{aligned} 2/3k - 2/9 - (k - 1) &< k - (2/3k - 2/9) \\ k &> 5/6 \end{aligned}$$

Since $k \geq 2$, the inequality is always true. Therefore, the player with $k - 1$ is closer, and thus he can get the entire \$1. We conclude that for any $k \geq 2$, the profile (k, k, k) cannot be a Nash equilibrium.

The strategy profile $(1, 1, 1)$, on the other hand, is NE. (Note that the above inequality works just fine for $k = 1$. However, since we cannot choose 0 as the integer, it is not possible to undercut the other two players with a smaller number.)

We now consider an strategy profile where not all three integers are the same. First consider a profile, in which one player names a highest integer. Denote an arbitrary such profile by (k^*, k_1, k_2), where k^* is the highest integer and $k_1 \geq k_2$. Two thirds of the average for this profile is $a = {}^2/_9(k^* + k_1 + k_2)$. If $k_1 > a$, then k^* is further from a than k_1, and therefore k^* does not win anything. If $k_1 < a$, then the difference between k^* and a is $k^* - a = {}^7/_9k^* - {}^2/_9k_1 - {}^2/_9k_2$. The difference between k_1 and a is $a - k_1 = {}^2/_9k^* - {}^7/_9k_1 + {}^2/_9k_2$. The difference between the two is then ${}^5/_9k^* + {}^5/_9k_1 - {}^4/_9k_2 > 0$, so k_1 is closer to a. Thus k^* does not win and the player who offers it is better off by deviating to k_1 and sharing the prize. Thus, no profile in which one player names a highest integer can be Nash equilibrium.

Consider now a profile in which two players name highest integers. Denote this profile by (k^*, k^*, k) with $k^* > k$. Then $a = {}^4/_9k^* + {}^2/_9k$. The midpoint of the difference between k^* and k is ${}^1/_2(k^* + k) > a$. Therefore, k is closer to a and wins the entire \$1. Either of the two other players can deviate by switching to k and thus share the prize. Thus, no such profile can be Nash equilibrium.

This exhausts all possible strategy profiles. We conclude that this game has a unique Nash equilibrium, in which all three players announce the integer 1.

2.4.6 Voting for Candidates

(Osborne, 34.2) There are n voters, of which k support candidate A and $m = n - k$ support candidate B. Each voter can either vote for his preferred candidate or abstain. Each voter gets a payoff of 2 if his preferred candidate wins, 1 if the candidates tie, and 0 if his candidate loses. If the citizen votes, he pays a cost $c \in (0, 1)$.

(a) What is the game with $m = k = 1$?

(b) Find the pure-strategy Nash equilibria for $k = m$.

(c) Find the pure-strategy Nash equilibria for $k < m$.

We tackle each part in turn:

(a) Let's draw the bi-matrix for the two voters who can either (V)ote or (A)bstain. This is depicted in Fig. 29 (p. 31).

		B Supporter	
		V	A
A Supporter	V	$1 - c, 1 - c$	$2 - c, 0$
	A	$0, 2 - c$	$1, 1$

Figure 29: The Election Game with Two Voters.

Since $0 < c < 1$, this game is exactly like the Prisoners' Dilemma: both citizens vote and the candidates tie.

(b) Here, we need to consider several cases. (Keep in mind that each candidate has an equal number of supporters.) Let $n_A \le k$ denote the number of citizens who vote for A and let $n_B \le m$ denote the number of citizens who vote for B. We restrict our attention to the case where $n_A \ge n_B$ (the other case is symmetric, so there is no need to analyze it separately). We now have to consider several different outcomes with corresponding classes of strategy profiles: (1) the candidates tie with either (a) all k citizens voting for A or (b) some of them abstaining; (2) some candidate wins either (a) by one vote or (b) by two or more votes. Thus, we have four cases to consider:

   (a) $n_A = n_B = k$: Any voting supporter who deviates by abstaining causes his candidate to lose the election and receives a payoff of $0 < 1 - c$. Thus, no voting supporter wants to deviate. This profile is a Nash equilibrium.

   (b) $n_A = n_B < k$: Any abstaining supporter who deviates by voting causes his candidate to win the election and receives a payoff of $2 - c > 1$. Thus, an abstaining supporter wants to deviate. This profile is not Nash equilibrium.

   (c) $n_A = n_B + 1$ or $n_B = n_A + 1$: Any abstaining supporter of the losing candidate who deviates by voting causes his candidate to tie and increases his payoff from 0 to $1 - c$. These profiles are not Nash equilibria.

   (d) $n_A \ge n_B + 2$ or $n_B \ge n_A + 2$: Any supporter of the winning candidate who switches from voting to abstaining can increase his payoff from $2 - c$ to 2. Thus, these profiles cannot be Nash equilibria.

   Therefore, this game has a unique Nash equilibrium, in which everybody votes and the candidates tie.

(c) Let's apply very similar logic to this part as well:

   (a) $n_A = n_B \le k$: Any supporter of B who switches from abstaining to voting causes B to win and improves his payoff from 1 to $2 - c$. Such a profile cannot be a Nash equilibrium.

   (b) $n_A = n_B + 1$ or $n_B = n_A + 1$, with $n_A < k$: Any supporter of the losing candidate can switch from abstaining to voting and cause his candidate to tie, increasing his payoff from 0 to $1 - c$. Such a profile cannot be a Nash equilibrium.

   (c) $n_A = k$ or $n_B = k + 1$: Any supporter of A can switch from voting to abstaining and save the cost of voting for a losing candidate, improving his payoff from $-c$ to 0. Such a profile cannot be a Nash equilibrium.

   (d) $n_A \ge n_B + 2$ or $n_B \ge n_A + 2$: Any supporter of the winning candidate can switch from voting to abstaining and improve his payoff from $2 - c$ to 2. Such a profile cannot be a Nash equilibrium.

   Thus, when $k < m$, the game has no Nash equilibrium (in pure strategies).

3 Symmetric Games

A useful class of normal form games can be applied in the study of interactions which involve anonymous players. Since the analyst cannot distinguish among the players, it follows that they have the same strategy sets (otherwise the analyst could tell them apart from the different strategies they have available).

DEFINITION 10. A two-player normal form game is **symmetric** if the players' sets of strategies are the same and their payoff functions are such that

$$u_1(s_1, s_2) = u_2(s_2, s_1) \text{ for every } (s_1, s_2) \in S.$$

That is, player 1's payoff from a profile in which he chooses strategy s_1 and his opponent chooses s_2 is the same as player 2's payoff from a profile, in which she chooses s_1 and player 1 chooses s_2. Note that these do not really have to be equal, it just has to be the case that the outcomes are ordered the same way for each player. (Thus, we're not doing interpersonal comparisons.) Once we have the same ordinal ranking, we can always rescale the appropriate utility function to give the same numbers as the other. Therefore, we continue using the equality while keeping in mind what it is supposed to represent. A generic example, as in Fig. 30 (p. 33) might help. You can probably already see that Prisoners' Dilemma and Stag

	A	B
A	w, w	x, y
B	y, x	z, z

Figure 30: The Symmetric Game.

Hunt are symmetric while BoS is not. We now define a special solution concept:

DEFINITION 11. A strategy profile (s_1^*, s_2^*) is a **symmetric Nash equilibrium** if it is a Nash equilibrium and $s_1^* = s_2^*$.

Thus, in a symmetric Nash equilibrium, all players choose the same strategy in equilibrium. For example, consider the game in Fig. 31 (p. 33). It has three Nash equilibria in pure strategies: (A, A), (C, A), and (A, C). Only (A, A) is symmetric.

	A	B	C
A	1,1	2,1	4,1
B	1,2	5,5	3,6
C	1,4	6,3	0,0

Figure 31: Another Symmetric Game.

Let's analyze several games where looking for symmetric Nash equilibria make sense.

3.1 Heartless New Yorkers

A pedestrian is hit by a taxi (happens quite a bit in NYC). There are n people in the vicinity of the accident, and each of them has a cell phone. The injured pedestrian is unconscious and requires immediate medical attention, which will be forthcoming if at least one of the n people calls for help. Simultaneously and independently each of the n bystanders decides whether to call for help or not. Each bystander obtains v units of utility if the injured person receives help. Those who call pay a personal cost of $c < v$. If no one calls, each bystander receives a utility of 0. Find the symmetric Nash equilibrium of this game. What is the probability no one calls for help in equilibrium?

We begin by noting that there is no symmetric Nash equilibrium in pure strategies: If no bystander calls for help, then one of them can do so and receive a strictly higher payoff of

$v - c > 0$. If all call for help, then any one can deviate by not calling and receive a strictly higher payoff $v > v - c$. (Note that there are n asymmetric Nash equilibria in pure strategies: the profiles, where exactly one bystander calls for help and none of the others do, are all Nash equilibria. However, the point of the game is that these bystanders are anonymous and do not know each other. Thus, it makes sense to look for a symmetric equilibrium.)

Thus, the symmetric equilibrium, if one exists, should be in mixed strategies. Let p be the probability that a person does not call for help. Consider bystander i's payoff of this mixed strategy profile. If each of the other $n - 1$ bystanders does not call for help, help will *not* arrive with probability p^{n-1}, which means that it will be called (by at least one of these bystanders) with probability $1 - p^{n-1}$.

What is i to do? His payoff is $[p^{n-1}(0) + (1 - p^{n-1})v] = (1 - p^{n-1})v$ if he does not call, and $v - c$ if he does. From Proposition 1, we must find p such that the payoffs from his two pure strategies are the same:

$$\begin{aligned}(1 - p^{n-1})v &= v - c \\ p^{n-1} &= {}^{c}/_{v} \\ p^{*} &= ({}^{c}/_{v})^{\frac{1}{n-1}}\end{aligned}$$

Thus, when all other bystanders play $p = p^{*}$, i is indifferent between calling and not calling. This means he can choose any mixture of the two, and in particular, he can choose p^{*} as well. Thus, the symmetric mixed strategy Nash equilibrium is the profile where each bystander calls with probability $1 - p^{*}$.

To answer the second question, we compute the probability which equals:

$$p^{*n} = ({}^{c}/_{v})^{\frac{n}{n-1}}$$

Since $n/(n-1)$ is decreasing in n, and because $c/v < 1$, it follows that the probability that nobody calls is increasing in n. The unfortunate (but rational) result is that as the number of bystanders goes up, the probability that any particular person will call for help goes down.[7] Intuitively, the reason for this is that while person i's payoff to calling remains the same regardless of the number of bystanders, the payoff to not calling increases as that number goes up, so he becomes less likely to call. This is not surprising. What is surprising, however, is that as the size of the group increases, the probability that at least one person will call for help decreases.

3.2 Rock, Paper, Scissors

Two kids play this well-known game. On the count of three, each player simultaneously forms his hand into the shape of either a rock, a piece of paper, or a pair of scissors. If both pick the same shape, the game ends in a tie. Otherwise, one player wins and the other loses according to the following rule: rock beats scissors, scissors beats paper, and paper beats rock. Each obtains a payoff of 1 if he wins, −1 if he loses, and 0 if he ties. Find the Nash equilibria.

We start by the writing down the normal form of this game as shown in Fig. 32 (p. 35).

[7] There is an infamous case of this dynamic in action: the murder of Kitty Genovese in New York City in 1964, with 38 neighbors looking on without calling the police.

		Player 2		
		R	P	S
	R	$0,0$	$-1,1$	$1,-1$
Player 1	P	$1,-1$	$0,0$	$-1,1$
	S	$-1,1$	$1,-1$	$0,0$

Figure 32: Rock, Paper, Scissors.

It is immediately obvious that this game has no Nash equilibrium in pure strategies: The player who loses or ties can always switch to another strategy and win. This game is symmetric and we shall look for symmetric mixed strategy equilibria first.

Let p, q, and $1-p-q$ be the probability that a player chooses R, P, and S respectively. We first argue that we must look only at completely mixed strategies (that is, mixed strategies that put positive probability on every available pure strategy). Suppose not, so $p_1 = 0$ in some (possibly asymmetric) MSNE. If player 1 never chooses R, then playing P is strictly dominated by S for player 2, so she will play either R or S. However, if player 2 never chooses P, then S is strictly dominated by R for player 1, so player 1 will choose either R or P in equilibrium. However, since player 1 never chooses R, it follows that he must choose P with probability 1. But in this case player 2's optimal strategy will be to play S, to which either R or S are better choices than P. Therefore, $p_1 = 0$ cannot occur in equilibrium. Similar arguments establish that in any equilibrium, any strategy must be completely mixed.

We now look for a symmetric equilibrium. Player 1's payoff from R is $p(0) + q(-1) + (1 - p - q)(1) = 1 - p - 2q$. His payoff from P is $2p + q - 1$. His payoff from S is $q - p$. In a MSNE, the payoffs from all three pure strategies must be the same, so:

$$1 - p - 2q = 2p + q - 1 = q - p$$

Solving these equalities yields $p = q = {}^1\!/_3$. Thus, whenever player 2 plays the three pure strategies with equal probability, player 1 is indifferent between his pure strategies, and hence can play any mixture. In particular, he can play the same mixture as player 2, which would leave player 2 indifferent among his pure strategies. This verifies the first condition in Proposition 1. Because these strategies are completely mixed, we are done. Each player's strategy in the symmetric Nash equilibrium is $({}^1\!/_3, {}^1\!/_3, {}^1\!/_3)$. That is, each player chooses among his three actions with equal probabilities.

Is this the only MSNE? We already know that any mixed strategy profile must consist only of completely mixed strategies in equilibrium. Arguing in a way similar to that for the pure strategies, we can show that there can be no equilibrium in which players put different weights on their pure strategies.

Generally, you should check for MSNE in all combinations. That is, you should check whether there are equilibria, in which one player chooses a pure strategy and the other mixes; equilibria, in which both mix; and equilibria in which neither mixes. Note that the mixtures need not be over the entire strategy spaces, which means you should check every possible subset.

Thus, in a 2×2 two-player game, each player has three possible choices: two in pure strategies and one that mixes between them. This yields 9 total combinations to check. Similarly, in a 3×3 two-player game, each player has 7 choices: three pure strategies, one completely mixed, and three partially mixed. This means that we must examine 49 combinations! (You can see how this can quickly get out of hand.) Note that in this case, you must check *both* conditions of Proposition 1.

4 Strictly Competitive Games

This is a special class of games that is not studied any more as much as it used to be. Nevertheless, it is important to know about them because (a) the results are not difficult, and (b) the results will be useful in later parts of the course.

A strictly competitive game is a two-player game where players have strictly opposed rankings over the outcomes. A good example is MATCHING PENNIES. That is, when comparing various strategy profiles, whenever one player's payoff increases, the other player's payoff decreases. Thus, there is no room for coordination or compromise. More formally,

DEFINITION 12. A **two-player strictly competitive game** is a two-player game with the property that, for every two strategy profiles, $s, s' \in S$,

$$u_1(s) > u_1(s') \Leftrightarrow u_2(s) < u_2(s').$$

A special case of strictly competitive games are the zero-sum games where the sum of the two players' payoffs is zero (e.g. MATCHING PENNIES).

		Player 2	
		L	R
Player 1	U	$3,2$	$0,4$
	D	$6,1$	$1,3$

Figure 33: A Strictly Competitive Game.

Consider the (non-zero-sum) strictly competitive game in Fig. 33 (p. 36). What is the worst-case scenario that player 1 could ever face? This is the case where player 2 chooses R, which yields a smaller payoff to player 1 whether he chooses U or D (he gets $0 < 3$ if he chooses U and $1 < 6$ if he chooses D).

More generally, the *worst* payoff that player i can get when he plays the (possibly mixed) strategy σ_i is defined by

$$w_i(\sigma_i) \equiv \min_{s_j \in S_j} u_i(s_i, s_j).$$

Which means that we look at all strategies available to player j to find the one that gives player i the smallest possible payoff if he plays σ_i. In other words, if player i chooses σ_i, he is guaranteed to receive a payoff of *at least* $w_i(\sigma_i)$. This is the smallest payoff that player 2 can hold player 1 to given player 1's strategy. A security strategy gives player i the best of the worst. That is, it solves $\max_{\sigma_i \in \Sigma_i} w_i(\sigma_i)$:

DEFINITION 13. A strategy $\hat{\sigma}_i \in \Sigma_i$ for player i is called a **security strategy** if it solves the expression

$$\max_{\sigma_i \in \Sigma_i} \min_{s_j \in S_j} u_i(\sigma_i, s_j),$$

which also represents player i's **security payoff**.

Returning to our example in Fig. 33 (p. 36), player 1's security strategy is D because given that player 2 is minimizing player 1's payoff by playing R, player 1 can maximize it by choosing D (because $1 > 0$). Similarly, player 1 can hold player 2 to at most 3 by playing D, to which player 2's best response is R.

Although there is no general relationship between security strategies and equilibrium strategies, an important exception exists for strictly competitive games.

Proposition 3. *If a strictly competitive game has a Nash equilibrium,* (σ_1^*, σ_2^*)*, then* σ_1^* *is a security strategy for player 1 and* σ_2^* *is a security strategy for player 2.*

Note that in our example, the unique Nash equilibrium is (D, R), which consists of the exact security strategies we found for both players.

If we look at Matching Pennies, we note that $w_1(H) = -1$ when $s_2 = T$, and $w_1(T) = -1$ when $s_2 = H$. On the other hand $w_1(1/2) = 0$, so player 1's security strategy is to mix between his two actions with equal probability. A symmetric argument establishes the same result for the other player. Again, the unique Nash equilibrium of this game is in security strategies.

5 Five Interpretations of Mixed Strategies

See Osborne and Rubinstein's *A Course in Game Theory*, pp. 37-44 for a more detailed treatment of this subject. Here, I only sketch several substantive justifications for mixed strategies.

5.1 Deliberate Randomization

The notion of mixed strategy might seem somewhat contrived and counter-intuitive. One (naïve) view is that playing a mixed strategy means that the player deliberately introduces randomness into his behavior. That is, a player who uses a mixed strategy commits to a randomization device which yields the various pure strategies with the probabilities specified by the mixed strategy. After all players have committed in this way, their randomization devices are operated, which produces the strategy profile. Each player then consults his randomization device and implements the pure strategy that it tells him to. This produces the outcome for the game.

This interpretation makes sense for games where players try to outguess each other (e.g. strictly competitive games, poker, and tax audits). However, it has two problems.

First, the notion of mixed strategy equilibrium does not capture the players' motivation to introduce randomness into their behavior. This is usually done in order to influence the behavior of other players. We shall rectify some of this once we start working with extensive form games, in which players move can sequentially.

Second, and perhaps more troubling, in equilibrium a player is indifferent between his mixed strategy and any other mixture of the strategies in the support of his equilibrium mixed strategies. His equilibrium mixed strategy is only one of many strategies that yield the same expected payoff given the other players' equilibrium behavior.

5.2 Equilibrium as a Steady State

Osborne (and others) introduce Nash equilibrium as a steady state in an environment in which players act repeatedly and ignore any strategic link that may exist between successive interactions. In this sense, a mixed strategy represents information that players have about past interactions. For example, if 80% of past play by player 1 involved choosing strategy A and 20% involved choosing strategy B, then these frequencies form the beliefs each player can form about the future behavior of other players when they are in the role of player 1. Thus, the corresponding belief will be that player 1 plays A with probability .8 and B with probability .2. In equilibrium, the frequencies will remain constant over time, and each player's strategy is optimal given the steady state beliefs.

5.3 Pure Strategies in an Extended Game

Before a player selects an action, he may receive a private signal on which he can base his action. Most importantly, the player may not consciously link the signal with his action (e.g. a player may be in a particular mood which made him choose one strategy over another). This sort of thing will appear random to the other players if they (a) perceive the factors affecting the choice as irrelevant, or (b) find it too difficult or costly to determine the relationship.

The problem with this interpretation is that it is hard to accept the notion that players deliberately make choices depending on factors that do not affect the payoffs. However, since in a mixed strategy equilibrium a player is indifferent among his pure strategies in the support of the mixed strategy, it may make sense to pick one because of mood. (There are other criticisms of this interpretation, see O&R.)

5.4 Pure Strategies in a Perturbed Game

Harsanyi introduced another interpretation of mixed strategies, according to which a game is a frequently occurring situation, in which players' preferences are subject to small random perturbations. Like in the previous section, random factors are introduced, but here they affect the payoffs. Each player observes his own preferences but not that of other players. The mixed strategy equilibrium is a summary of the frequencies with which the players choose their actions over time.

Establishing this result requires knowledge of Bayesian Games, which we shall obtain later in the course. Harsanyi's result is so elegant because even if no player makes any effort to use his pure strategies with the required probabilities, the random variations in the payoff functions induce each player to choose the pure strategies with the right frequencies. The equilibrium behavior of other players is such that a player who chooses the uniquely optimal pure strategy for each realization of his payoff function chooses his actions with the frequencies required by his equilibrium mixed strategy.

5.5 Beliefs

Other authors prefer to interpret mixed strategies as beliefs. That is, the mixed strategy profile is a profile of beliefs, in which each player's mixed strategy is the common belief of all other players about this player's strategies. Here, each player chooses a single strategy, not a mixed one. An equilibrium is a steady state of beliefs, not actions. This interpretation is the one we used when we defined MSNE in terms of best responses. The problem here is that each player chooses an action that is a best response to equilibrium beliefs. The set of these best responses includes every strategy in the support of the equilibrium mixed strategy (a problem similar to the one in the first interpretation).

6 The Fundamental Theorem (Nash, 1950)

Since this theorem is such a central result in game theory, we shall present a somewhat more formal version of it, along with a sketch of the proof. A *finite* game is a game with finite number of players and a finite strategy space. The following theorem due to John Nash (1950) establishes a very useful result which guarantees that the Nash equilibrium concept provides a solution for every finite game.

THEOREM 1. *Every finite game has at least one mixed strategy equilibrium.*

Recall that a pure strategy is a degenerate mixed strategy. This theorem does not assert the existence of an equilibrium with non-degenerate mixing. In other words, every finite game will have at least one equilibrium, in pure or mixed strategies.

The proof requires the idea of best response correspondences we discussed. However, it is moderately technical in the sense that it requires the knowledge of continuity properties of correspondences and some set theory. I will give the outline of the proof here but you should read Gibbons pp. 45-48 for some additional insight.

Proof. Recall that player i's best response correspondence $BR_i(\sigma_{-i})$ maps each strategy profile σ to a set of mixed strategies that maximize player i's payoff when the other players play σ_{-i}. Let $r_i = BR_i(\sigma)$ for all $\sigma \in \Sigma$ denote player i's best reaction correspondence. That is, it is the set of best responses for all possible mixed strategy profiles. Define $r : \Sigma \rightrightarrows \Sigma$ to be the Cartesian product of the r_i. (That is, r is the set of all possible combinations of the players best responses.) A *fixed point* of r is a strategy profile $\sigma^* \in r(\sigma^*)$ such that, for each player, $\sigma_i^* \in r_i(\sigma^*)$. In other words, a fixed point of r is a Nash equilibrium.

The second step involves showing that r actually has a fixed point. Kakutani's fixed point theorem establishes four conditions that together are sufficient for r to have a fixed point:

1. Σ is compact,[8] convex,[9] nonempty subset of a finite-dimensional Euclidean space;[10]

2. $r(\sigma)$ is nonempty for all σ;

3. $r(\sigma)$ is convex for all σ;

4. r is upper hemi-continuous.[11]

We must now show that Σ and r meet the requirements of Kakutani's theorem. Since Σ_i is a simplex of dimension $\#S_i - 1$ (that is, the number of pure strategies player i has less 1), it is compact, convex, and nonempty. Since the payoff functions are continuous and defined on compact sets, they attain maxima, which means $r(\sigma)$ is nonempty for all σ. To see the third case, note that if $\sigma' \in r(\sigma)$ and $\sigma'' \in r(\sigma)$ are both best response profiles, then for each player i and $\alpha \in (0, 1)$,

$$u_i(\alpha\sigma_i' + (1-\alpha)\sigma_i'', \sigma_{-i}) = \alpha u_i(\sigma_i', \sigma_{-i}) + (1-\alpha)u_i(\sigma_i'', \sigma_{-i}),$$

[8] Any sequence in Σ has a subsequence that converges to a point in Σ. Alternatively, a compact set is closed and bounded.

[9] Σ is convex if every convex combination of any two points in the set is also in the set.

[10] For our purposes, the Euclidean space is the same as $\mathbb{R}^n$, i.e. the set of n-tuples of real numbers.

[11] A correspondence is upper-hemicontinuous at x_0 if every sequence in which $r(x) \to x_0$ has a limit which lies in the image set of x_0. That is, if $(\sigma^n, \hat{\sigma}^n) \to (\sigma, \hat{\sigma})$ with $\hat{\sigma}^n \in r(\sigma^n)$, then $\hat{\sigma} \in r(\sigma)$. This condition is also sometimes referred to as $r(\cdot)$ having a *closed graph*.

that is, if both σ_i' and σ_i'' are best responses for player i to σ_{-i}, then so is their weighted average. Thus, the third condition is satisfied. The fourth condition requires sequences but the intuition is that if it were violated, then at least one player will have a mixed strategy that yields a payoff that is strictly better than the one in the best response correspondence, a contradiction.

Thus, all conditions of Kakutani's fixed point theorem are satisfied, and the best reaction correspondence has a fixed point. Hence, every finite game has at least one Nash equilibrium. $\square$

Somewhat stronger results have been obtained for other types of games (e.g. games with uncountable number of actions). Generally, if the strategy spaces and payoff functions are well-behaved (that is, strategy sets are nonempty compact subset of a metric space, and payoff functions are continuous), then Nash equilibrium exists. Most often, some games may not have a Nash equilibrium because the payoff functions are discontinuous (and so the best reply correspondences may actually be empty).

Note that some of the games we have analyzed so far do not meet the requirements of the proof (e.g. games with continuous strategy spaces), yet they have Nash equilibria. This means that Nash's Theorem provides sufficient, but not necessary, conditions for the existence of equilibrium. There are many games that do not satisfy the conditions of the Theorem but that have Nash equilibrium solutions.

Now that existence has been established, we want to be able to characterize the equilibrium set. Ideally, we want to have a unique solution, but as we shall see, this is a rare occurrence which happens only under very strong and special conditions. Most games we consider will have more than one equilibrium. In addition, in many games the set of equilibria itself is hard to characterize.