Analytic Theory II: Strategic-Form Games of Complete Information (Review)

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

January 24, 2020

Contents

---

1 The Strategic (Normal) Form

Every strategy profile s induces an outcome of the game: a sequence of moves actually taken as specified by the strategies and a probability distribution over the terminal nodes of the game. If the game is one of certainty (no moves by Nature), then s specifies one outcome with certainty. Otherwise, more than one outcome may occur with positive probability. The point is that we can calculate the expected payoffs of all players. Sometimes, it is useful to analyze the game in its **strategic form**, which includes only the players, their actions, and the payoffs in its description.

Putting things a little more formally, let n be the number of players. For each player i, denote the strategy space by S_i. (We shall sometimes write $s_j \in S_i$ to reflect that strategy s_j is a member of the set of strategies S_i.) Let $(s_1, s_2, \ldots, s_n)$ denote a strategy profile, where s_1 is the action of player 1, s_2 is the action of player 2, and so on. Let $S = S_1 \times S_2 \times \ldots \times S_n$ denote the set of strategy profiles.

For each player i, define the vNM expected utility function $U_i : S \to \mathbb{R}$ so that for each $s \in S$ that players choose, $U_i(s)$ is player i's expected payoff from outcome s.

DEFINITION 1. For a game with $\mathcal{I} = \{1, \ldots, n\}$ players, the **strategic (normal) form representation** $G = \{\mathcal{I}, S, U\}$ specifies for each player i a set of strategies S_i and a payoff function $U_i : S \to \mathbb{R}$, where $S = \times S_i$, and $U = (U_1, \ldots, U_n)$.

When we analyze these games, we often assume that players choose their strategies simultaneously, and hence we call them **simultaneous-move games**. However, this does not require that players strictly act at the same time. All that is necessary is that each player acts without knowledge of what others have done. That is, players cannot condition their strategies on observable actions of the other players.

Of course, this ignores the information about timing of moves explicitly specified by the extensive form. The question boils down to whether we think such questions are essential to the situation we are trying to analyze. If they are not, then it should not matter greatly if we simplify our description to exclude such information. In an important sense, the strategic form is a **static model** because it dispenses with the dynamics of timing of moves completely.

This may not be as controversial (or useless) as it sounds. First, as we shall see, there are great many situations that we might profitably analyze without reference to the timing of moves. Second, the simplified representation is actually considerably easier to analyze, so we can benefit from dispensing with information that is not essential. We shall, of course, also see that there are many, many situations where ignoring timing has crucial consequences and our solutions based on the normal form will be quite suspicious precisely because they will discard such information. The question (again) will boil down to the choice of representation, which a researcher has to make based on her skill and experience.

1.1 Reduced Strategic Form

Two pure strategies are *equivalent* if they induce the same probability distribution over the outcomes for all pure strategies for the opponents. Or, putting it a bit more formally:

DEFINITION 2. Given any strategic form game $G = \{\mathcal{I}, S, U\}$, for any player i and any two strategies $s_1, s_2 \in S_i$, the strategies s_1 and s_2 are **payoff-equivalent** if, and only if,

$$U_j(s_1, s_{-i}) = U_j(s_2, s_{-i}), \quad \forall s_{-i} \in S_{-i}, \quad \forall j \in \mathcal{I}.$$

That is, no matter what all other players do, no player cares whether i uses s_1 or s_2. Let's parse this expression. To see whether two strategies for player 1 are payoff-equivalent, we take each strategy of player 2 in turn and compare the payoffs that player 1 obtains from playing s_1 and s_2 against that, then we compare the payoffs that player 2 obtains from player 1 playing s_1 and s_2 against her strategy. If either of these two comparisons produces a difference, stop: the two strategies are not payoff-equivalent. If, on the other hand, they yield the same payoffs in both cases, proceed to the next strategy for player 2 and repeat the process. If you exhaust all strategies for player 2 in this way and the comparisons have not yielded any differences, then the two strategies for player 1 are payoff-equivalent.

Consider the game in Fig. 1 (p. 3), where the strategies AE and AF for player 1 are equivalent.

		Player 2	
		c	d
Player 1	AE	1, 1	1, 1
	AF	1, 1	1, 1
	BE	−1, 1	3, 2
	BF	−1, 1	4, 0

Figure 1: Simple Game with Equivalent Pure Strategies.

In this example, the two pure strategies AE and AF always lead to the same outcome because the game ends when the first action is taken and so the second information set is never reached. This happens regardless of what player 2 does at her information set. That is, fix player 2's strategy to be c, then: (i) player 1's payoff from AE is 1, which is the same as his payoff from AF; (ii) player 2's payoff from player 1 choosing AE is 1, which is the same as her payoff from him choosing AF. So neither player cares if player 1 chooses AE or AF if player 2 chooses c. Next, fix player 2's strategy to be d, then: (i) player 1's payoff from AE is 1, which is the same as his payoff from AF; (ii) player 2's payoff from player 1 choosing AE is 1, which is the same as her payoff from him choosing AF. Hence, neither player cares if player 1 chooses AE or AF if player 2 chooses d. Since there are no more strategies for player 2 to check against, we are done: no player cares what player 1 does regardless of what player 2 chooses. Observe that in these comparisons we had to check whether player 1 himself would care, not just whether his opponent would. We can now simplify the normal form representation by removing all but one strategies from every class of equivalent strategies.

DEFINITION 3. The **purely reduced normal form** of an extensive form game is obtained by eliminating all but one member of each equivalence class of pure strategies.

Therefore, we can remove either AE or AF (but not both) to obtain the reduced normal form shown in Fig. 2 (p. 3). The "new" strategy for player 1 is called A.

		Player 2	
		c	d
Player 1	A	1, 1	1, 1
	BE	−1, 1	3, 2
	BF	−1, 1	4, 0

Figure 2: The Reduced Normal Form of the Game from Fig. 1 (p. 3).

The example we just did may be a bit misleading because the payoffs for the players are always the same in all the outcomes regardless of what player 2 chooses. This need not be the case. To see that, consider the strategic form game in Fig. 3 (p. 4).

To decide whether U and D are payoff-equivalent, we first fix player 2's strategy at L and observe that players get $(3, 1)$ no matter which of the two pure strategies under consideration player 1 chooses. We then fix player 2's strategy at R and observe that players get $(-2, 0)$ regardless of whether player 1

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

Player 1 \ Player 2	L	R
U	3, 1	−2, 0
M	4, 3	4, 3

Figure 3: Reducing a Game with Different Payoffs.

chooses U or D. Hence, the two are payoff-equivalent, and we can eliminate one of them. Observe that a player can get different payoffs depending on whether player 2 chooses L or R from strategies that are payoff-equivalent (i.e., player 1 can get either 3 or −2) but this is not the relevant comparison to make. For example, both players get $(4, 3)$ if player 1 chooses M regardless of player 2's action. However, this does not mean that L and R are payoff-equivalent (because players would get different payoffs against either one of these if player 1 chooses a different strategy.)

Consider now the game in Fig. 4 (p. 4).

Player 1 \ Player 2	x	y
(a, c)	6, 0	6, 0
(a, d)	6, 0	6, 0
(a, e)	6, 0	6, 0
(b, c)	8, 0	0, 8
(b, d)	0, 8	8, 0
(b, e)	3, 4	7, 0

Figure 4: Another Game from Myerson.

It is fairly obvious that the strategies (a, c), (a, d), and (a, e) are payoff equivalent to one another because regardless of what player 2 does, the outcome from all three is the same. In other words, player 1 does not care what player 2 does if he chooses any of these three strategies. We can therefore merge these three strategies into a new one, called A, with the resulting payoff matrix in Fig. 5 (p. 4).

Player 1 \ Player 2	x	y
A	6, 0	6, 0
(b, c)	8, 0	0, 8
(b, d)	0, 8	8, 0
(b, e)	3, 4	7, 0

Figure 5: The Purely Reduced Strategic Form of the Game from Fig. 4 (p. 4).

We can reduce this game further, but to do this, we need to introduce the concept of mixed strategies.

2 Mixed Strategies in Strategic Form Games

So far, we have considered only strategies that involve playing a selected action with probability 1. We called these **pure strategies** to emphasize this. We now consider randomized choices.

DEFINITION 4. A **mixed strategy** for player i, denoted by σ_i, is a probability distribution over i's set of pure strategies S_i. Denote the mixed strategy space for player i by Σ_i, where $\sigma_i(s_i)$ is the probability that σ_i assigns to the pure strategy $s_i \in S_i$. The space of mixed strategy profiles is denoted by $\Sigma = \triangle \Sigma_i$.

Thus, if player i has K pure strategies: $S_i = \{s_{i1}, s_{i2}, \ldots, s_{iK}\}$, then a mixed strategy for player i is a probability distribution $\sigma_i = \{\sigma_i(s_{i1}), \sigma_i(s_{i2}), \ldots, \sigma_i(s_{iK})\}$, where $\sigma_i(s_{ik})$ is the probability that player i will choose strategy s_{ik} for $k = 1, 2, \ldots, K$. Since σ_i is a probability distribution, we require that $\sigma_i(s_{ik}) \in [0, 1]$ for all $k = 1, 2, \ldots, K$ and $\sum_{k=1}^{K} \sigma_i(s_{ik}) = 1$. That is, the probabilities must be non-negative and not larger than 1, and should sum up to 1. You can think of a mixed strategy as a lottery whose "outcomes" are pure strategies.

Each player's randomization is statistically independent of those of his opponents,[1] and the payoffs to the mixed strategy profile are the expected values of the corresponding pure strategy payoffs.[2] You should now see why we needed Expected Utility Theory. Player i's payoff from a mixed strategy profile $\sigma \in \Sigma$ in an n-player game is

$$U_i(\sigma) = \sum_{s \in S} \left(\prod_{j=1}^{n} \sigma_j(s_j) \right) u_i(s)$$

Let's parse this expression. The mixed strategy profile σ is a list of mixed strategies, one for each player: $\sigma = \{\sigma_1, \sigma_2, \ldots, \sigma_n\}$. Each of these mixed strategies, e.g. σ_i, is a list of probabilities associated with player i's set of pure strategies. To find the probability of an outcome, we need to calculate the probability that all players choose the pure strategies that produce this outcome. Thus, if the pure strategy profile $s \in S$ produces the outcome we are interested in, the probability of this outcome is the product of probabilities that each player chooses the pure strategy in this profile (because of independence).

Consider first an example from a game without chance moves, like Matching Pennies. To make things specific, let's use the mixed strategy profile $\sigma = \langle(^1\!/_3 H, ^2\!/_3 T), (^1\!/_4 H, ^3\!/_4 T)\rangle$. In this profile, player 1's mixed strategy specifies playing H with probability $^1\!/_3$ and T with probability $^2\!/_3$, and player 2's mixed strategy strategy specifies playing H with probability $^1\!/_4$, and T with probability $^3\!/_4$. There are four pure strategy profiles: $S = \{(H, H), (H, T), (T, H), (T, T)\}$ that produce the four outcomes of the game.

As usual, the strategy profile σ induces a probability distribution over the outcomes. The probability of each outcome is the product of the probabilities that each player chooses the relevant strategy. For example, the probability of the pure strategy profile (H, H) being played is $(^1\!/_3)(^1\!/_4) = {}^1\!/_{12}$. Analogously, the probabilities of the other pure strategy profiles being played are $\Pr(H, T) = {}^1\!/_4$, $\Pr(T, H) = {}^1\!/_6$, and $\Pr(T, T) = {}^1\!/_2$. (You should verify that these sum to 1, which they must because they are probabilities of exhaustive and mutually exclusive events.) Fig. 6 (p. 5) shows the probability distribution over the four possible outcomes induced by the two mixed strategies.

	H	T
H	$^1\!/_{12}$	$^1\!/_4$
T	$^1\!/_6$	$^1\!/_2$

Figure 6: The probability distribution over outcomes induced by σ.

Player 1's payoffs from these outcomes are $u_1(H, H) = u_1(T, T) = 1$ and $u_1(H, T) = u_1(T, H) = -1$. Multiplying the payoffs by the probability of obtaining them and summing over (the expected utility calculation we have done before) yields an expected payoff of $^1\!/_{12}(1) + {}^1\!/_2(1) + {}^1\!/_4(-1) + {}^1\!/_6(-1) =$

[1] That is, the joint probability equals the product of individual probabilities.

[2] In all cases where we shall calculate mixed strategies, the space of pure strategies will be finite so we do not run into measure-theoretic problems.

$^1/_6$. Thus, player 1's expected payoff from the mixed strategy profile σ as specified above is $^1/_6$. Note how we first did the multiplication term and then summed over all available pure strategy profiles, while multiplying by the utility of each. This is exactly what the expression above does. Recalling that $S = \{(H,H),(H,T),(T,H),(T,T)\}$, we can write:

$$\begin{aligned}
U_1(\sigma) &= \sum_{s \in S} \left(\prod_{j=1}^{2} \sigma_j(s_j) \right) u_1(s) \\
&= \sigma_1(H)\sigma_2(H)u_1(H,H) + \sigma_1(H)\sigma_2(T)u_1(H,T) \\
&\quad + \sigma_1(T)\sigma_2(H)u_1(T,H) + \sigma_1(T)\sigma_2(T)u_1(T,T) \\
&= (^1/_3)(^1/_4)(1) + (^1/_3)(^3/_4)(-1) + (^2/_3)(^1/_4)(-1) + (^2/_3)(^3/_4)(1) \\
&= {}^1/_6.
\end{aligned}$$

Consider now an example from a game that does involve chance moves, like the Card Game, whose strategic form is in Fig. 25 (p. 25). Suppose we wanted to know player 2's expected payoff from the mixed strategy profile $\sigma = \langle(^1/_3, {}^1/_4, {}^5/_{12}, 0), (^1/_3, {}^2/_3)\rangle$. That is, for player 1, $\sigma_1(Rr) = {}^1/_3$, $\sigma_1(Rf) = {}^1/_4$, $\sigma_1(Fr) = {}^5/_{12}$, and $\sigma_1(Ff) = 0$, whereas for player 2, $\sigma_2(m) = {}^1/_3$ and $\sigma_2(p) = {}^2/_3$. So,

$$\begin{aligned}
U_2(\sigma) &= \sum_{s \in S} \left(\prod_{j=1}^{2} \sigma_j(s_j) \right) u_2(s) \\
&= \sigma_1(Rr)\sigma_2(m)u_2(Rr,m) + \sigma_1(Rr)\sigma_2(p)u_2(Rr,p) \\
&\quad + \sigma_1(Rf)\sigma_2(m)u_2(Rf,m) + \sigma_1(Rf)\sigma_2(p)u_2(Rf,p) \\
&= \sigma_1(Fr)\sigma_2(m)u_2(Fr,m) + \sigma_1(Fr)\sigma_2(p)u_2(Fr,p) \\
&\quad + \sigma_1(Ff)\sigma_2(m)u_2(Ff,m) + \sigma_1(Ff)\sigma_2(p)u_2(Ff,p) \\
&= (^1/_3)(^1/_3)(0) + (^1/_3)(^2/_3)(-1) + (^1/_4)(^1/_3)(0.5) + (^1/_4)(^2/_3)(-1) \\
&= (^5/_{12})(^1/_3)(-0.5) + (^5/_{12})(^2/_3)(0) + (0)(^1/_3)(0) + (0)(^2/_3)(0) \\
&= -{}^5/_{12}.
\end{aligned}$$

If you wanted to compute the probability distribution over the outcomes induced by σ, you should get the result in Tab. 7 (p. 6).

	m	p
Rr	$^1/_9$	$^2/_9$
Rf	$^1/_{12}$	$^1/_6$
Fr	$^5/_{36}$	$^5/_{18}$
Ff	0	0

Figure 7: The probability distribution over outcomes for Fig. 25 (p. 25) induced by σ.

As the last example showed, there is no requirement that a mixed strategy puts positive probabilities on all available pure strategies. The **support** of a mixed strategy σ_i is the set of strategies to which σ_i assigns positive probability. This means that we can think of a pure strategy s_i as a **degenerate mixed strategy** that assigns probability 1 to s_i and 0 to all remaining pure strategies (i.e. the support of a degenerate mixed strategy consists of a single pure strategy). A **completely mixed strategy** assigns positive probability to every strategy in S_i.[3]

[3]Completely mixed strategies are important because a strategy profile of completely mixed strategies assigns positive prob-

As mentioned in the previous section, we can further reduce some strategic form games. Consider the game in Fig. 5 (p. 4). Although no other pure strategies are payoff-equivalent, the strategy (b, e) is redundant in an important sense. Suppose player 1 were to choose between the strategy A and (b, d) with a flip of a fair coin. The resulting randomized strategy can be denoted with $\sigma = 0.5[A] + 0.5[b, d]$, and would give the expected payoffs:

$$U(\sigma, x) = (0.5)(6, 0) + (0.5)(0, 8) = (3, 4)$$
$$U(\sigma, y) = (0.5)(6, 0) + (0.5)(8, 0) = (7, 0).$$

In other words, we could get the payoffs from (b, e) from randomizing between the strategies A and (b, d). We formalize this notion as follows:

DEFINITION 5. A strategy $\hat{s}_i \in S_i$ is **randomly redundant** if and only if there exists a mixed strategy $\sigma_i \in \Sigma_i$ such that $\sigma_i(\hat{s}_i) = 0$ and

$$U_j(\hat{s}_i, s_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) u_j(s_i, s_{-i}) \quad \forall s_{-i} \in S_{-i}, \quad \forall j \in \mathcal{I}.$$

That is each player's payoffs from the profiles involving $\hat{s}_i$ can be expressed as the expected payoffs from a mixed strategy for player i that does not have $\hat{s}_i$ in its support. In other words, $\hat{s}_i$ is randomly redundant if there is some way for player i to mix his other pure strategies such that no matter what combination of strategies the other players choose, every player would get the same expected payoff whether i uses $\hat{s}_i$ or mixes in this way.

DEFINITION 6. The **fully reduced normal form** of an extensive form game Γ is obtained from the purely reduced representation of Γ by eliminating all randomly redundant strategies.

The fully reduced normal form Fig. 4 (p. 4) (whose purely reduced normal form is in Fig. 5 (p. 4)) is given in Fig. 8 (p. 7).

		Player 2	
		x	y
	A	6, 0	6, 0
Player 1	(b, c)	8, 0	0, 8
	(b, d)	0, 8	8, 0

Figure 8: The Fully Reduced Strategic Form of the Game from Fig. 4 (p. 4).

Consider the example in Fig. 9 (p. 8): how are we to approach something like this to decide whether there are any strategies that are randomly redundant? Obviously, the only possibilities must involve strategies for player 1, but which one(s)? We can begin by simple elimination by asking whether any two strategies can be mixed to eliminate a third one. We cannot eliminate A by any mixture of two or more of the remaining three pure strategies because player 1's payoff against L is negative if he plays A and non-negative otherwise. Since any mixture of B, C, and D must yield a non-negative payoff against A as well, there is no way to match the payoff from A. It is also impossible to eliminate B with any combination of the other three strategies: player 1's payoff against L is 3, which is strictly greater than any of the other payoffs he could get against L. This means that any mixture of A, C, and D must yield player 1

ability to every possible outcome in the game. As we shall see later, the fundamental solution concept (Nash equilibrium) will not produce any odd results in that situation. Problems with Nash equilibrium (in the sense of unreasonable predictions about optimal behavior) might only occur when the strategy profile induces zero probability for one or more of the possible outcomes.

an expected payoff strictly less than 3, so they cannot match B. It is also impossible to eliminate C; this time, note that player 2's payoff against C when she plays L is -1, which is strictly less than any of her payoffs against the other three strategies for player 1. This means that any mixture of A, B, and D must give player 2 a payoff strictly better than -1 when she chooses L, so it will not be possible to match C.

		Player 2 L	Player 2 R
Player 1	A	$-1, 0$	$-1/3, 1/2$
	B	$3, 1/2$	$-1, 9/8$
	C	$0, -1$	$1, 0$
	D	$3/4, -1/4$	$0, 1/2$

Figure 9: Less Obvious Example.

All of this means that if there is any randomly redundant strategy for player 1, it would have to be D. What mixture of some combination of A, B, and C can work? First, note that it cannot be a mixture between A and B by themselves: player 2's payoff from L would be non-negative and she must get $-1/4$ to match her payoff against D. Can it be a mixture between B and C by themselves? Looking at player 1's payoffs against R, we can see that he gets -1 from B and 1 from C. There is only one way to match the payoff of 0 he obtains from D: mix B and C with equal probabilities. But then player 2's payoff against the mixture would be $9/16$ when she chooses R, which does not match her payoff of $1/2$ against D. Hence, it is not possible to eliminate D with a mixture of B and C alone.

This leaves us with just one more possibility: mix A, B, and C to eliminate D. If D is randomly-redundant, then the following system of equations must have a unique solution:

$$
\begin{aligned}
-\sigma_1(A) + 3\sigma_1(B) &= 3/4 \\
-1/3\sigma_1(A) - \sigma_1(B) + \sigma_1(C) &= 0 \\
1/2\sigma_1(B) - \sigma_1(C) &= -1/4 \\
1/2\sigma_1(A) + 9/8\sigma_1(B) &= 1/2,
\end{aligned}
$$

such that $\sigma_1(A) + \sigma_1(B) + \sigma_1(C) = 1$ and $\sigma_1(a) \in (0, 1)$ for all $a \in \{A, B, C\}$. From the last equation, we obtain $\sigma_1(A) = 1 - 9/4\sigma_1(B)$. Plugging this into the first equation and multiplying both sides by 4 then gives us $-4 + 9\sigma_1(B) + 12\sigma_1(B) = 3$, which then yields the solution $\sigma_1(B) = 7/21 = 1/3$. Plugging this into the third equation yields $1/6 - \sigma_1(C) = -1/4$, so $\sigma_1(C) = 5/12$. Finally, plugging these two into the second equation reduces it to $-1/3\sigma_1(A) - 1/3 + 5/12 = 0$, which implies $\sigma_1(A) = 1/4$. Of course, since we know the probabilities must sum up to 1, we could have just computed $\sigma_1(A) = 1 - \sigma_1(B) - \sigma_1(C)$ to obtain the same result. This way, however, we can verify that the sum is unity, so we have not messed up any of our calculations. We now have the mixed strategy $\sigma_1 = (1/4, 1/3, 5/12, 0)$ which yields the same expected payoffs to either player as D does against L, and the same expected payoffs to either player as D does against R. Hence, D is randomly redundant and we can safely eliminate it without losing anything in the process.

One question you may have at this point is what happens if there are more than one randomly-redundant strategies: would it matter which one gets eliminated first? What if we use some pure strategy to eliminate another and then eliminate that pure strategy itself: does that mean we have to restore the one we originally eliminated or is it possible to eliminate it without using that pure strategy? As it turns out, it does not matter which order you do the elimination in: if you can eliminate a pure strategy d by a mixed strategy that has s, s', and s'' in its support and then s itself gets eliminated by another mixed strategy with only s' and s'' in its support, then it is possible to eliminate d with a mixed strategy that only has s' and s'' in its support. Let's see an example that illustrates this, so consider Fig. 10 (p. 9).

	L	R
A	$1, 2$	$-2, 0$
B	$0, 3$	$-1/2, 2$
C	$-1, 4$	$1, 4$
D	$-1/4,\ 13/4$	$-1/8,\ 5/2$

	L	R
A	$1, 2$	$-2, 0$
B	$0, 3$	$-1/2, 2$
C	$-1, 4$	$1, 4$

	L	R
A	$1, 2$	$-2, 0$
C	$-1, 4$	$1, 4$

Figure 10: Order of Elimination Does Not Matter.

The mixed strategy $\sigma = (1/4, 1/4, 1/2, 0)$ makes D randomly-redundant in the original game on the left, producing the reduced normal form in the middle. But then $\sigma' = (1/2, 0, 1/2)$ makes B randomly-redundant in that intermediate form, producing the fully reduced form on the right. The question then is: since we used B to eliminate D in the first step, would we still be able to eliminate D now that we B itself is gone? That is, do we need B to keep D out? The claim is that since B can be eliminated by A and C, then it should be possible to eliminate D with only these two strategies as well. What is the appropriate mixture then? Since mixing A and C with equal weights eliminates B, let's distribute the weight on B in the original σ evenly to A and C and check if the result can eliminate D. That is, add $1/8$ to the probabilities σ assigns to A and C to consider $\sigma'' = (3/8, 0, 5/8, 0)$ in the original game. It is straightforward to verify that this strategy makes D randomly redundant: against L player's expected payoff is $3/8 - 5/8 = -2/8 = -1/4$ and player 2's expected payoff is $3/8(2) + 5/8(4) = 26/8 = 13/4$; analogously, against R, player 1's expected payoff is $3/8(-2) + 5/8 = -1/8$, and player 2's expected payoff is $5/8(4) = 5/2$. This means that we can use σ'' to eliminate D and then σ' to eliminate B, yielding the same fully reduced form.

It is sometimes quite tricky to identify randomly redundant strategies. It may be worth your while to try anyway because by reducing the number of strategies to consider for the analysis, you will greatly simplify your task (you will see what I mean when we begin solving the games next time). Unless we explicitly state otherwise, we shall take the *reduced strategic form representation* to mean the fully reduced form.

You might wonder why we are eliminating redundant strategies: after all, the ones we remove from considerations do, in fact, specify ways to play the game and reach possibly different outcomes. For instance, in the reduced strategic form in Fig. 10 (p. 9), there are no outcomes $\langle D, L\rangle$ or $\langle D, R\rangle$, which were both available in the original specification. Aren't we losing something when we do not consider them? If there are several redundant strategies, does it not matter which ones we eliminate? The answer is that for the *analysis* of the game, it will not matter. When we find solutions that involve a strategy that has other payoff-equivalent ones in the original game, then we will immediately know that the original game has more solutions: we would obtain those by replacing the strategy with the payoff-equivalent ones we eliminated. Thus, suppose for instance that in the reduced form we found solutions in which A and C are played with probability $1/2$ each. Because we know that this mixed strategy is payoff equivalent to the pure strategy B, we immediately know that there are solutions in which player 2's strategy is the same but player 1 plays B instead of that particular mixed strategy. If, however, the solution involved A and C with some other probabilities, then there will be no solutions that involve B. Thus, when we want to provide a substantive interpretation for the solution, we have to remember the payoff-equivalent strategies.

3 Nash Equilibrium

The most common definition of rationality in game theory is based on the idea that players would choose strategies that yield the highest expected payoff given what they think the other players are doing; that is, players would choose the **best response** to their expectations about the behavior of others. Since all players are "rational" in that sense, they must expect the others to be choosing their best responses as well.

In other words, all players must be best-responding to each other. When this happens, no player would have an incentive to change their strategy because, by definition, it cannot be improved upon by any other strategy. This is why a profile of strategies that are mutual best responses is called an **equilibrium**, and it was named in honor of John Nash who proved that most games must have at least one strategy profile with that property.

Scholars often refer to particular definitions of rationality as *solution concepts*, which is presumably meant to emphasize the fact that any conceptual definition of rationality is one among many. Nash equilibrium is foundational not only because it is the most commonly used one but because it underlies many stronger definitions of rationality like subgame perfection, perfect Bayesian equilibrium, and so on.

For the remainder of this course, we shall define **rational behavior** as *choosing the best response to one's expectations about the behaviors of others*.

3.1 Nash Equilibrium in Pure Strategies

Rational players think about actions that the other players might take, and then choose strategies that yield the highest expected payoff given their expectations about the others. Such strategies are called **best responses** (or *best replies*).

DEFINITION 7. 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 $BR_i(s_{-i})$. Consider, for example, the game in Fig. 11 (p. 10): In this game, $BR_1(L) = \{M\}$, $BR_1(C) = \{U, M\}$, and $BR_1(R) = \{U\}$. Also,

		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 11: The Best Response Game.

$BR_2(U) = \{C, R\}$, $BR_2(M) = \{R\}$, 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.)

The best response correspondence for a player is a function of their beliefs about what the other players are doing. When the other players are rational in the same sense, these beliefs are not arbitrary: the player must expect the others to be choosing best responses as well. This allows us to create a specific solution concept based on that definition of rationality: a Nash equilibrium is a strategy profile such that each player's strategy is a best response to the other players' strategies:

DEFINITION 8 (NASH EQUILIBRIUM). The strategy profile $(s_i^*, s_{-i}^*) \in S$ is a **pure-strategy Nash equilibrium** (PSNE) 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 9. The strategy profile (s_i^*, s_{-i}^*) is a PSNE 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, the payoff from the profile $\langle s_i^*, s_{-i}^* \rangle$ is at least as good as the payoff from the profile $\langle s_i, s_{-i}^* \rangle$ 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.[4]

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 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, 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. 11 (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

[4]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.

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. 12 (p. 12).

		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 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.

3.1.1 Diving Money

(Osborne, 38.2) Two players have \$10 to divide. Each names an integer $0 \leq k \leq 10$. If $k_1 + k_2 \leq 10$, 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.

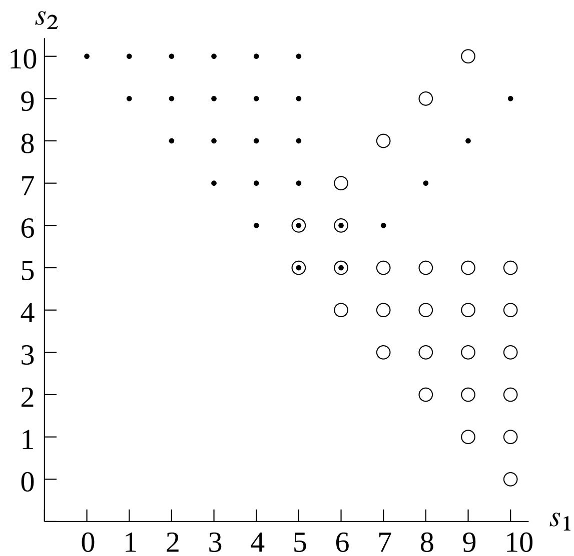

Figure 13: Best Responses in the Dividing Money Game.

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 $0 \le k_1 \le 4$. There are several cases to consider:

- if $k_2 \le k_1$, 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 $k_1 + k_2 \le 10$, then he gets k_1; if he plays 5 instead, then:
  - if $5 + k_2 \le 10$, 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 $0 \le k_2 \le 4$ 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. 14 (p. 13).

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

Figure 14: 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.[5] In this case, you could re-do Fig. 13 (p. 12) with only $s_i \geq 5$ to get all four solutions.

3.1.2 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 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:

$$\begin{aligned} x^* &= 1 + cy^* \\ y^* &= 1 + cx^*. \end{aligned}$$

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.

3.1.3 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. 15 (p. 15). The three discs at the points where the best response functions intersect represent the three pure-strategy Nash equilibria we found above.

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

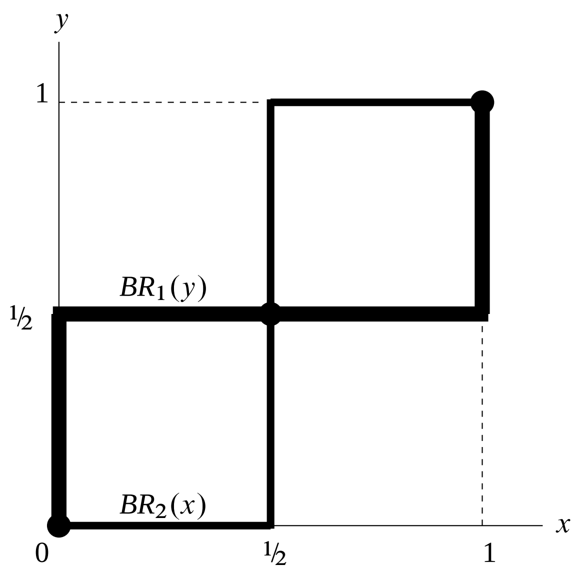

Figure 15: Best Responses in the Modified Partnership Game.

3.2 Strict Nash Equilibrium

Consider the game in Fig. 16 (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.)

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

Figure 16: The Battle of Bismarck Sea.

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 10. 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.

3.3 Nash Equilibrium in Mixed Strategies

The most common example of a game with no Nash equilibrium in pure strategies is MATCHING PENNIES, which is given in Fig. 17 (p. 16).

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

Figure 17: Matching Pennies.

This is a strictly competitive (zero-sum) situation, in which the gain for one player is the loss of the other.[6] 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 11. A mixed strategy profile σ^* is a **mixed-strategy Nash equilibrium** (MSNE) if, and 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:

$$U_1(H, \sigma_2) = (1)q + (-1)(1-q) = 2q - 1$$
$$U_1(T, \sigma_2) = (-1)q + (1)(1-q) = 1 - 2q.$$

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

$$U_1(H, \sigma_2) \geq U_1(T, \sigma_2)$$
$$2q - 1 \geq 1 - 2q$$
$$q \geq 1/2.$$

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.

[6] 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.

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. 18 (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 $p = q = 1/2$.

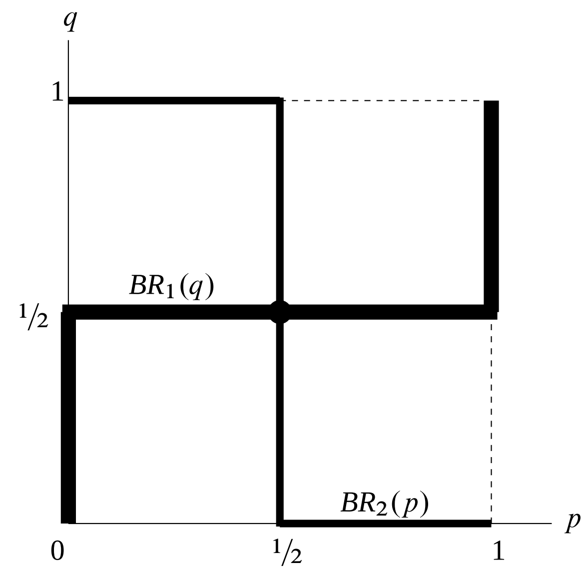

Figure 18: Best Responses in Matching Pennies.

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 12. A mixed strategy profile σ^* is a MSNE if, and only if, for all players i,

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

Since expected utilities are linear in the probabilities, if a player uses a non-degenerate mixed strategy in a Nash equilibrium, then they must be indifferent between all pure strategies to which they assign 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.)

The fact that a player who is willing to mix in MSNE must be indifferent among the pure strategies used with positive probability raises several questions about this definition of rationality. First, many scholars (myself included) are uncomfortable with the idea that players randomize their actions. Second, even when players are indifferent among several pure strategies, it is unclear why they should be randomizing with the distribution required by MSNE. Third, it appears that the other players are choosing their strategies in order to make the player uncertain, and so able to mix, which seems like an odd thing to do.

We shall have an occasion to discuss these at length later in the course. For now, I will note that it is possible to interpret mixed strategies in a way that alleviates both concerns. The idea behind MSNE is not that a player is randomizing but that the other players cannot predict with certainty what that player will do. The player could be choosing a pure strategy – based, perhaps, on factors known only to that player – as long as the other players do not know what these factors are, that player's behavior will appear random to them. Their expectations about the distribution of these random acts anchor their own best responses — if it were different, then their own behavior would change. In other words, nobody is picking strategies at random in a way to make others indifferent. The distribution reflects the type of uncertainty necessary to rationalize the best responses of the other players; that is, it reflects what beliefs they must have about the strategy of the player that they are uncertain about. It is now easy to understand why a player whose behavior appears unpredictable to the others must be expected by them to be indifferent among the pure strategies used with positive probability: if this were not the case, then there would exist a strategy that is a best response, and the others must expect that player to choose it; i.e., they cannot be uncertain.

3.3.1 Battle of the Sexes

Let's model a situation where two players, $i \in \{1, 2\}$, want to decide between two types of entertainment to which they want to go together but the decision must be made without knowledge of what the other will do (say they are in their offices and the phones are down so they cannot communicate beforehand). The two available pieces of entertainment for the night are a boxing match (fight) and a ballet. For each player then, the set of actions consists of (1) go to the fight, or (2) go to the ballet. Note that the actions are exhaustive and mutually exclusive. This means that each player has two pure strategies, so the set is called the *strategy space* for the player.

Continuing with the example, the strategy profile then consists of one strategy for each of the two players. This gives us four different strategy profiles: (1) player 1 goes to the fight, player 2 goes to the fight; (2) player 1 goes to the fight, player 2 goes to the ballet; (3) player 1 goes to the ballet, player 2 goes to the fight; and (4) player 1 goes to the ballet, player 2 goes to the ballet. We shall specify an outcome (strategy profile) by listing first the strategy for player 1 and then the strategy for player 2. Thus, the four outcomes above can be written as (1) *(Fight, Fight)*; (2) *(Fight, Ballet)*; (3) *(Ballet, Fight)*; and (4) *(Ballet,Ballet)*.

Since each strategy profile produces a different outcome in this game, the game has 4 possible outcomes, in 2 of which the players go together to the same place, and 2 in which they fail to coordinate. Each player has (ordinal) preferences over these four outcomes. In other words, each player ranks these outcomes according to their desirability using some criterion. As we know, if preferences are rational, we can represent them numerically. Hence, we use appropriate numbers whose ordinal ranking represents the preferences as payoffs. Each outcome then consists of two elements which specify the payoff for each player for this outcome. This is often called the *payoff vector*.

Player 1 prefers going to the fight whereas player 2 prefers going to the ballet.[7] However, both prefer to go together regardless of the type of entertainment. Their worst outcome is when they end up alone at any of the places and it does not matter which place they happen to be at. Thus, player 1's preference ordering

[7] In the original game, player 1 was a man and player 2 was a woman. There is now a more politically-correct version of the BoS game, called *Bach or Stravinsky,* which involves two sexless players deciding between concerts of music by the two composers. Since any rational person with taste would clearly choose Stravinsky, I find that version uninteresting.

is:

$$(F, F) \succ (B, B) \succ (F, B) \sim (B, F)$$

and player 2's preference ordering is:

$$(B, B) \succ (F, F) \succ (F, B) \sim (B, F)$$

Now that we have specified the ordinal rankings, we need to choose a payoff function to represent the orderings. Denote player 1's utility function by u_1, and player 2's utility function by u_2. We need two functions such that:

$$u_1(F, F) > u_1(B, B) > u_1(F, B) = u_1(B, F)$$
$$u_2(B, B) > u_2(F, F) > u_2(F, B) = u_2(B, F).$$

One possible and simple specification is

$$u_1(F, F) = u_2(B, B) = 2$$
$$u_1(B, B) = u_2(F, F) = 1$$
$$u_1(F, B) = u_1(B, F) = u_2(F, B) = u_2(B, F) = 0.$$

A convenient way of describing the (finite) strategy spaces of the players and their payoff functions for two-player games is to use a bi-matrix,[8] as illustrated in Fig. 19 (p. 19).

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

Figure 19: Battle of the Sexes.

Note: the Battle of the Sexes game represents a situation where players must coordinate their actions but where they have opposed preferences over the coordinated outcomes. We shall see two other types of coordination games: pure coordination (where players only care about coordinating) and Pareto coordination (where both strictly prefer one of the coordinated outcomes to the other).

Recall that although we call this a *simultaneous-moves* game, it is not necessary for players to actually act at the same time. All that is required is that each player acts with no knowledge about how the other player acts. In our BoS game, this can be achieved by requiring the players to make their choices without having access to a communication device.

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. 20 (p. 20) 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

[8]This is just like a regular matrix except each entry consists of two numbers instead of one.

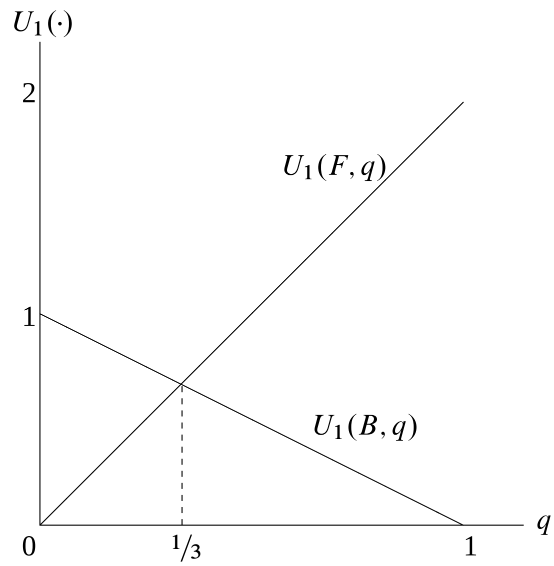

Figure 20: Player 1's Expected Payoffs as a Function of Player 2's Mixed Strategy.

be indifferent between any mixing of them. Thus, we 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. 21 (p. 21) 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. 22 (p. 21). 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 choosing the relevant action.

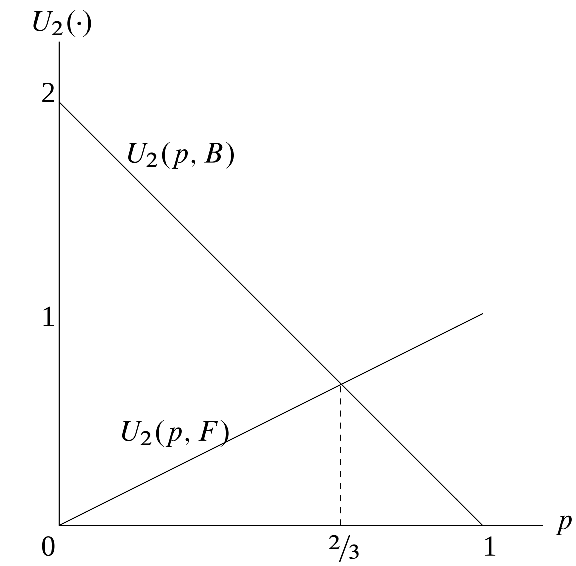

Figure 21: Player 2's Expected Payoffs as a Function of Player 1's Mixed Strategy.

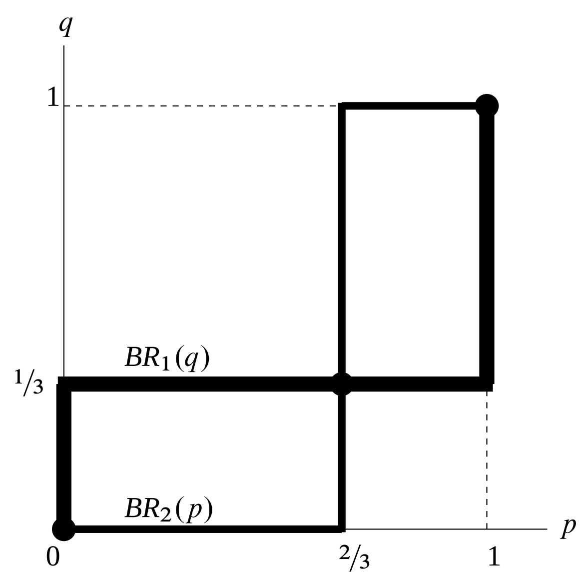

Figure 22: Best Responses in Battle of the Sexes.

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 $(2)\,2/9 + (1)\,2/9 = 2/3$.

3.4 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 1. 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. 23 (p. 22).

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

Figure 23: 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 1 works as advertised.

The property in Remark 1 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 MSNE 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 MSNE.* □

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).

3.4.1 Myerson's Card Game

Consider the following game from Roger Myerson.

EXAMPLE 1. (MYERSON'S CARD GAME.) There are two players, labeled "player 1" and "player 2."[9] At the beginning of this game, each player puts a dollar in a pot. Next, player 1 draws a card from a shuffled deck of cards in which half the cards are red and half are black. Player 1 looks at his card privately and decides whether to raise or fold. If player 1 folds, then he shows his card to player 2 and the game ends; player 1 takes the money in the pot if the card is red, but player 2 takes the money if the card is black. If player 1 raises, then he adds another dollar to the pot and player 2 must decide whether meet or pass. If she passes, the game ends and player 1 takes all the money in the pot. If she meets, she puts another dollar in the pot, and then player 1 shows his card to player 2 and the game ends; if the card is red, player 1 takes all the money in the pot, but if it is black, player 2 takes all the money.

The extensive form of this game is in Fig. 24 (p. 24).

Let us convert this to strategic form. von Neumann and Morgenstern suggested a procedure for simplifying games in extensive form by constructing the strategic form G of any Γ. This is done in an algorithmic way. First, we find all pure strategies for the players. Next, we construct the expected outcomes for all

[9] We establish the following convention: odd-numbered players are male, and even-numbered players are female. For a generic player, we shall always use the generic male pronoun.

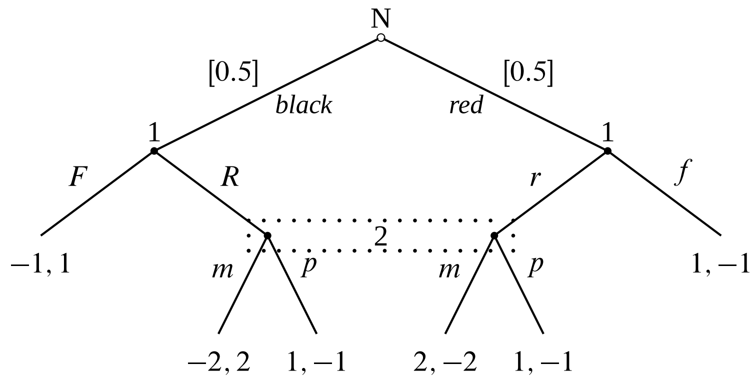

Figure 24: Myerson's Card Game: Extensive Form.

strategy profiles. Finally, we redefine the utility functions on the outcomes to be utility functions on the profiles with expected outcomes.

Consider the following scenario. The two players are going to play this tomorrow and today they have to plan their moves in advance. Player 1 does not know the color that he will draw but he can condition his strategy on the card color because he knows that he will see it before choosing whether to raise or fold. As we have seen, he has four pure strategies, $S_1 = \{Rr, Rf, Fr, Ff\}$. Player 2, on the other hand, will only ever get to move if player 1 raises, so her pure strategies are $S_2 = \{m, p\}$. The strategy profiles are:

$$S = S_1 \times S_2 = \left\{ \langle Rr, m\rangle, \langle Rr, p\rangle, \langle Rf, m\rangle, \langle Rf, p\rangle, \langle Fr, m\rangle, \langle Fr, p\rangle, \langle Ff, m\rangle, \langle Ff, p\rangle \right\}.$$

We now have to define the expected utility functions for the player. Recall that originally, we defined the utility functions directly in terms of the outcome. However, even if we knew here which strategy profile will be realized (that is, what strategy each player has chosen), we cannot predict the actual outcome of the game because it will depend on the color of the card, which is a chance move. For example, suppose player 1 has chosen the strategy Fr and player 2 has chosen m, and so the strategy profile is $\langle Fr, m\rangle$. The outcome will be folding by player 1 if the card is black, and raising by player 1 and meeting by player 2 if the card is red. Player 1's payoff will be -1 if the card is black, and 2 if the card is red.

So what payoff should player 1 expect from the profile $\langle Fr, m\rangle$? Its expected payoff, of course. Choosing the strategy Fr given that player 2 will be choosing m is equivalent to choosing a lottery, in which player 1 would get -1 with probability 0.5, and 2 with probability 0.5. We know how to compute the expected utility in this case:

$$U_1(Fr, m) = {}^1\!/_2 \times u_1(\text{black}, F) + {}^1\!/_2 \times u_1(\text{red}, r, m) = {}^1\!/_2 \times (-1) + {}^1\!/_2 \times (2) = 0.5.$$

In analogous manner, we would compute player 2's expected payoff:

$$U_2(Fr, m) = {}^1\!/_2 \times u_2(\text{black}, F) + {}^1\!/_2 \times u_2(\text{red}, r, m) = {}^1\!/_2 \times (1) + {}^1\!/_2 \times (-2) = -0.5.$$

Continuing in this way, we define the expected utility functions for the two players on all strategy profiles, and arrive the the normal form representation of this game of uncertainty shown in Fig. 25 (p. 25).

The strategic game in Fig. 25 (p. 25) describes how the utilities of the players depend on the strategies they choose *at the beginning of the game*. We know from our expected utility theorem that a player would choose the strategy that yields the highest expected payoff because this would be consistent with his preferences. In other words, players will make choices that maximize their expected payoff.

In general, given any extensive form game Γ, its normal form representation G can be constructed as follows. The set of players remains the same. For any player $i \in \mathcal{I}$, let the set of strategies S_i in the

		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 25: The Strategic Form of the Game from Fig. 24 (p. 24).

normal form game be the same as the set of strategies in the extensive form. For any strategy profile $s \in S$ and any node x in the tree of Γ, define $P(x|s)$ to be the probability that the path of play will go through node x, when the path of play starts at the initial node, and at any decision node in the path, the next node is determined by the relevant player's strategy in s, and, at any node where nature moves, the next node is determined by the probability distribution given in Γ. At any terminal node $z \in \mathcal{Z}$, let $u_i(z)$ be player i's payoff from outcome z. Then, for any strategy profile $s \in S$ and any $i \in \mathcal{I}$, let $U_i(s)$ be:

$$U_i(s) = \sum_{z \in \mathcal{Z}} P(z|s)u_i(z).$$

That is, $U_i(s)$ is player i's expected utility if all players implement the strategies according to s. If G is derived from Γ in this way, it is called the **strategic (normal) form representation** of Γ.

The first step in solving for Nash equilibrium (either in pure or mixed strategies) should always be the elimination of any strictly dominated strategies. In this game, no pure strategy is strictly dominated by another pure strategy. However, the strategy $s_1 = Ff$ is strictly dominated by the mixed strategy $\sigma_1 = (0.5)[Rr] + (0.5)[Fr]$:

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

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.[10]

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.

After eliminating Ff, we end up with the reduced strategic form of the game:

It is clear by inspection that this game has no PSNE, so let's look for one in mixed strategies. 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.[11]

[10] How did we know to try this mixed strategy? Notice that Ff is weakly dominated by Fr and Rr, and strictly dominated by Fr against m and by Rr against p. This means that mixing (in any way, actually) between Fr and Rr would yield a strictly higher payoff against either m or p.

[11] How do we know to show that? Even though one cannot safely eliminate weakly dominated strategies from consideration for inclusion in Nash equilibrium, they often can be eliminated with equilibrium reasoning. That is, by supposing that they are being used with positive probability and finding a contradiction in the assumption that the strategy is a best response. Sometimes this exercise allows us to eliminate weakly dominated strategies, and sometimes it does not. Here, it does.

		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 26: Myerson's Card Game: Reduced Strategic Form.

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) supp$(\sigma_1^*) = \{Rr, Rf, Fr\}$, (ii) supp$(\sigma_1^*) = \{Rr, Rf\}$, or (iii) 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 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. 27 (p. 27). 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 27: Myerson's Card Game: Further Reduction after Equilibrium Reasoning.

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 $-{}^1/_3$. Of course, we could have simply exploited the fact that in MSNE all pure strategies yield the same expected payoff to obtain:

$$U_1^* = U_1(Rr) = \sigma_2^*(p) = {}^1/_3 \quad \text{and} \quad U_2^* = U_2(p) = -\sigma_1^*(Rr) = -{}^1/_3.$$

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.[12]

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?

3.4.2 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. 28 (p. 29). The algorithm for finding all Nash equilibria

[12]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.

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 28: 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.)[13]

- $\text{supp}(\sigma_2) = \{L, M\}$. 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

[13] 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.

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.

3.4.3 Choosing Numbers

Players 1 and 2 each choose a positive integer up to K. Thus, the strategy spaces are both $\{1, 2, \dots, K\}$. 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.[14]

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.

Let us conjecture that players must put positive probability on each possible number in equilibrium. 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.

OK, so the probability distribution has full support. But what is the equilibrium distribution? Since players are mixing, they must be indifferent among their pure strategies. The only way player 1 will

[14]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.

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.

3.4.4 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. 29 (p. 31).

		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 29: 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:

$$\begin{aligned} p_1 - p_2 - p_3 + p_4 &\leq 0 \\ p_1 + p_4 &\leq p_2 + p_3 \end{aligned}$$

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

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.)

3.4.5 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.

3.4.6 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. 30 (p. 34).

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

Figure 30: 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 $n_A \leq k$ denote the number of citizens who vote for A and let $n_B \leq m$ denote the number of citizens who vote for B. We restrict our attention to the case where $n_A \geq n_B$ (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) $n_A \geq n_B + 2$ or $n_B \geq n_A + 2$: 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) $n_A = n_B \leq k$: 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) $n_A \geq n_B + 2$ or $n_B \geq n_A + 2$: 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.[15]

4 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 13. 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. 31 (p. 35) might help. You can probably already see that Prisoners' Dilemma and Stag Hunt are

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

Figure 31: The Symmetric Game.

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

DEFINITION 14. 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. 32 (p. 36). It has three Nash equilibria in pure strategies: (A, A), (C, A), and (A, C). Only (A, A) is symmetric.

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

[15] Finding the MSNE is quite involved.

	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 32: Another Symmetric Game.

4.1 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. 33 (p. 36).

		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 33: 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.

4.2 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 result is that as the number of bystanders goes up, the probability that any particular person will call for help goes down. 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.[16]

4.3 All-Pay Auction

There are $n > 1$ bidders for an object, each of whom values it at $v > 0$. All players simultaneously submit bids, $s_i \geq 0$, and the winner is the bidder who submits the highest bid (if there are multiple highest bidders, the winner is chosen randomly among them). Everyone pays their bid to the auctioneer regardless of whether they win or not. The payoffs are thus $v - s_i$ if bidder i is the winner, and $-s_i$ if she is not.

Let us first check if this game has equilibria in pure strategies. For fun, let us deal with asymmetric strategy profiles first. Consider any strategy profile where not all positive bids are the same. Then anyone who has submitted a losing bid could improve their payoff by submitting a zero bid. No such profile can be an equilibrium. Consider now a profile where there are $k \geq 2$ positive bids that are all the same, and everyone else bids nothing. Since the object is allocated randomly among the highest bidders, each of them expects $v/k - s_i$. Any one of these bidders could deviate to a bid $s_i + \varepsilon$ and ensure a win with a payoff $v - (s_i + \varepsilon)$. But then

$$v - s_i - \varepsilon > \frac{v}{k} - s_i \quad \Leftrightarrow \quad \varepsilon < \frac{v(k-1)}{k},$$

which means that such a profitable deviation always exists. Consider now a profile where there is exactly one highest bidder with s^* and everyone else bids zero. If $s^* < v$, then any of the players who bids nothing could deviate to some $s^* + \varepsilon$ and win with a strictly positive payoff. This is not possible only when $s^* = v$. However, the strategy profile where only one player bids cannot be an equilibrium because there is always a lower bid that is also winning (that player can deviate to, say, half of the supposed optimal bid). There are no asymmetric PSNE in this game.

There are also no symmetric PSNE. If nobody bids, any player could profit by making a tiny positive bid and win. If everyone bids the same and the bid is $s^* < v$, then it is profitable to increase the bid slightly to break the tie and win for sure. And if everyone bids $s^* = v$, each player is better off not bidding at all. What about symmetric MSNE?

Let $F(x)$ denote the cdf induced by the mixed strategy (assume that it is atomless). That is, if a player bids x, then the probability it will exceed all other bids is $F(x)^{n-1}$. Since no player would ever bid more than v with positive probability (even winning in that case is worse than not bidding), it follows that $F(v) = 1$ in any MSNE. Since the player is willing to mix in equilibrium, they must be indifferent among all bids in the support of F, or:

$$F(x)^{n-1} v - x = \overline{u},$$

where $\overline{u}$ is some, as yet unknown, payoff that is constant in the bid. Rearranging terms yields:

$$F(x) = \left(\frac{\overline{u} + x}{v} \right)^{\frac{1}{n-1}}.$$

[16] This is why emergency training often says that when there are several bystanders one should not just shout "Someone call 911!" but should instead point to a specific person and shout "You call 911!" This has the effect of immediately coordinating expectations on one of the asymmetric Nash equilibria. By the way, this result is often used to "explain" the story of Kitty Genovese, who was stabbed to death in 1964. The NYT claimed at the time that there had been 38 witnesses who saw or heard the murder without a single one of them calling the police or rushing to help her. The story has been largely debunked: it is not clear that anyone actually saw the murder, those who heard anything were apparently very few in number, many of them not recognizing what the sounds meant, and two people did call the police.

Using $F(v) = 1$, this means that

$$\left(\frac{\overline{u} + v}{v}\right)^{\frac{1}{n-1}} = 1 \quad \Rightarrow \quad \overline{u} = 0,$$

which pins down the expected payoff and yields the solution:

$$F(x) = \left(\frac{x}{v}\right)^{\frac{1}{n-1}}.$$

In other words, the strategy mixes over all bids $x \leq v$ such that the cdf satisfies the condition above, and yields an expected payoff of zero. No bids $x > v$, which are not in the support of the mixed strategy, can improve upon that payoff because they yield $v - x < 0$.

Competition among the bidders has left them with no expected surplus from the auction. Since the mixed strategy (the pdf) that induces $F(x)$ derived above is

$$f(x) = \frac{\left(\frac{x}{v}\right)^{\frac{1}{n-1}}}{(n-1)x},$$

the expected payment for each bidder is:

$$\int_0^v x f(x) \mathrm{d}x = \left.\frac{x\left(\frac{x}{v}\right)^{\frac{1}{n-1}}}{n}\right|_0^v = \frac{v}{n},$$

which means that the expected revenue for the auctioneer is v.

5 Some Canonical Games

Several simple 2-by-2 games have proven to be especially useful in thinking through the strategic aspect of certain situations. We have already seen the Battle of the Sexes, which is a canonical example of what Schelling calls "mixed-motive" scenarios: situations have both cooperative and conflictual aspects. In the BoS game, players very much want to coordinate (on going together) but disagree which of the cooperative outcomes should obtain (they prefer different entertainment). This is sometimes referred to as a game of ranked coordination as opposed to one where players simply want to coordinate without preference over the way they do it. One example of this would be choosing which side of the road to drive on: both driving on the left or both driving on the right are preferable to the alternatives, with neither being superior to the other. (A similar mixed-motive scenario is the Game of Chicken, where players wish to take different actions, and the risk is that they match instead.)

Another famous coordination game is the Stag Hunt (SH), where the problem is not that players disagree over the preferred cooperative outcome but that they might not be able to trust each other enough to achieve it. This is different from the Prisoner's Dilemma (PD), where each player has a strict individual incentive not to cooperate. Since most of you are quite familiar with the PD game, let us take a closer look at the (much more interesting) SH.

5.1 The Stag Hunt

The Prisoner's Dilemma is one type of social problem which assumes that unilateral defection is preferable to mutual cooperation. There are, however, situations in which mutual cooperation is the most preferred outcome for both players. And yet, as we shall now see, this in no way guarantees their ability to cooperate.

The classic illustration of such a social dilemma is due to Jean-Jacques Rousseau, and the story goes as follows. Two hunters must decide whether to cooperate, C, and hunt a stag together, or defect, D, and chase after a rabbit individually. If the both stalk the stag, they are certain to catch it, and they can feast on it. However, it requires both of them to stalk it, and if even one of them does not, the stag is certain to get away. If, on the other hand, a hunter goes chasing a bunny rabbit, he is certain to catch one regardless of what the other one does. Assume that if the other one is also hunting for rabbits, the noise they both make scares the tastiest rabbits away and they can only catch old and nasty hares with lower nutritional value. In other words, if a hunter go after a rabbit, there is a slight preference that he does so on his own. Even the best rabbit is worse for a hunter than his share of the stag. There is only time to stalk the stag or hunt for rabbits, they cannot do both.

We set up the situation as a simultaneous-move two-player game. Each of the hunters has two strategies: cooperate, C, or defect, D. The possible outcomes are: both cooperate and catch the stag (Stag), one hunter chases a rabbit and the other stalks the stag (Yummy Bunny and Hunger, respectively), and both hunt for rabbits (Stale Hare). The preference orderings are:

	Stag	$\succ$	Yummy Bunny	$\succ$	Stale Hare	$\succ$	Hunger
Hunter 1	(C, C)	$\succ$	(D, C)	$\succ$	(D, D)	$\succ$	(C, D)
Hunter 2	(C, C)	$\succ$	(C, D)	$\succ$	(D, D)	$\succ$	(D, C)

Table 1: Preferences in the Stag Hunt.

Compare the rankings of the strategy profiles to those in the Prisoner's Dilemma. As before, unreciprocated cooperation is the worst possible outcome for each player, and mutual defection is the second worst outcome. Unlike the PD, however, the preferences in a Stag Hunt situation are such that both players prefer mutual cooperation to unilateral defection.

The best responses in pure strategies involve matching what the other player is doing. Thus, if the other hunter is expected to go after the stag, cooperating (Stag) is preferable to defecting (Yummy Bunny). Analogously, if the other hunter is expected to go after a rabbit, defecting (Stale Hare) is preferable to cooperating (Hunger). This means that there are two PSNE: $\langle C, C\rangle$ and $\langle D, D\rangle$.[17]

Unlike the PD, mutual cooperation can be sustained in equilibrium. Unfortunately, like the PD, mutual defection can also be an equilibrium. In that sense, assuming that both players prefer mutual cooperation to every other possible outcome does not actually mean that they will cooperate. This is a fairly startling result and it is worth thinking through why it happens.

Recall that a best response is a strategy that is optimal given what you think the other player is doing. In this sense, cooperation is best if you think the other is cooperating. In a Nash equilibrium, these expectations are self-enforcing in the sense that *your* expectation of the other player choosing to cooperate rationalizes *your* choice to cooperate, which in turn validates *their* expectation that you will cooperate, which then rationalizes *their* choice to cooperate, and this in turn validates *your* expectation that they will cooperate, closing the circle of mutually supporting expectations.

Unfortunately, the exact same logic applies in the case of defection. If you think your partner will defect, you will defect as well, which validates their expectation that you will defect, which rationalizes their defection, which in turn validates your expectation that they will defect. Again, the circle is complete and we have an equilibrium with mutually supporting expectations.

The question then seems to boil down to where we "begin" the circle of expectations. For instance, if we think one of the hunters expects the other to cooperate, we end up with the cooperative equilibrium. If, on the other hand, we think one of the hunters expects the other to defect, we end up with the non-cooperative

[17]This means that there is a MSNE as well. We shall derive it later when we consider this game in generic form.

equilibrium. So which expectation is more likely? Without knowing the hunters and their relationship, it is impossible to say for sure.[18] However, we could ask ourselves: if I were one of these hunters, which is the *least risky* choice to make? That is, which choice gives me an outcome that leaves me least vulnerable to the behavior of the other hunter?

In a sense, we are trying to protect ourselves from a mistaken expectation. Let's say I generally trust the other hunter to cooperate but I also know that sometimes he gets tempted when he sees rabbits, and I am not entirely sure that he will not see a rabbit or that if he sees one while stalking the stag, he won't abandon the stalking in order to chase after the rabbit. Now, if I cooperate, I would get the stag if he does not get distracted but I will end up hungry if he does. If I defect, I would get the juicy rabbit if does not get distracted, and I will end up with a stale hare if he does. When I cooperate, the worst possible thing that can happen to me is to go hungry. When I defect, the worst possible thing that can happen to me is to end up with a stale hare. In that sense, defection is less risky because it leaves me less vulnerable in the case that I have misjudged my partner or he makes a mistake.

In case you are wondering, this can be formalized precisely. The notion of *risk-dominance* is due to Harsanyi and Selten, and for this game it can applied as follows. For each equilibrium, we can compute the product of losses if someone deviates from it. Suppose you are Hunter 1, and consider your situation. You are supposed to play the cooperative equilibrium $\langle C, C\rangle$ but instead you deviate it. Since C is a best response to C, this deviation is going to cost you: your payoff from $\langle D, C\rangle$ cannot exceed the equilibrium payoff by the very definition of equilibrium. In this case, you are going to suffer a *deviation loss* of $L_1 = u_1(C, C) - u_1(D, C)$. Consider now the non-cooperative equilibrium $\langle D, D\rangle$ and suppose you deviate from your strategy. This time, you will end up at $\langle C, D\rangle$ with a deviation loss of $L_1' = u_1(D, D) - u_1(C, D)$. Compare now your two deviation losses: if the loss from $\langle D, D\rangle$ is greater than the loss from $\langle C, C\rangle$, $L_1' > L_1$, then you should be *less likely* to deviate from $\langle D, D\rangle$. Intuitively, you stand to lose more if you do so, so you would have less incentive to do it. From the other player's perspective, then, $\langle D, D\rangle$ appears less risky: you are more likely to stick with the equilibrium strategy. We can now apply the same argument to the other player, if her deviation loss from $\langle D, D\rangle$, $L_2' = u_2(D, D) - u_2(D, C)$, exceeds her deviation loss from $\langle C, C\rangle$, $L_2 = u_2(C, C) - u_2(C, D)$, it makes sense that you should consider it more likely that she should stick with her equilibrium strategy under $\langle D, D\rangle$.

Putting these two together, we can compute the *risk-dominance* of one equilibrium profile over another. Take the product of the deviation losses for the players: for $\langle C, C\rangle$ it is $L_1 \times L_2$, whereas for $\langle D, D\rangle$ it is $L_1' \times L_2'$. The profile with the *higher* product of losses is said to be risk-dominant: it is the one that players are more likely to stick with. In this game, the assumptions that having the stag is marginally better than a rabbit whereas that the failure to catch anything leads to starvation boil down to $L_i' > L_i$. This means that the risk-dominant profile is $\langle D, D\rangle$. As a result, we would expect $\langle D, D\rangle$ to be the equilibrium players coordinate on, and mutual defection will be the outcome.

The risk-dominance argument would select the non-cooperative equilibrium even though one might initially believe that rational actors would surely coordinate on the cooperative one: after all, both of them would get better payoffs in $\langle C, C\rangle$ than they do in $\langle D, D\rangle$. In the context of a stag hunt, the advantage of avoiding the worst-case scenario might not be obvious, at least not as obvious as it is when we recast the Stag Hunt as an arms race (which we shall shortly do).[19]

Even small doubts about his trustworthiness may make me think about defection. Now, it gets worse if you consider what this means for my partner. Suppose he is aware that I harbor small doubts about his ability to resist temptation. Suppose he is resolved to resist it too. The problem is that when he is aware

[18] This is where the MSNE would be a natural choice as a the solution as it would involve failure to coordinate with positive probability.

[19] Evolutionary models in which reproduction rates depend on relative success from interactions also select the risk-dominant equilibrium.

of my doubt, he knows that I may be tempted to protect myself to avoid going home hungry. But this then makes him even more tempted to defect in order to protect himself from being left with nothing. And of course, I am aware of all of this, which makes me even more suspicious that he might actually defect, which in turn makes me more likely to select the strategy that leaves me least vulnerable to that defection. In other words, *we are very likely to end up in the non-cooperative equilibrium!*

This is a very pessimistic result: we both prefer the cooperative equilibrium to everything else, and this fact is common knowledge. And yet, even small amounts of doubt about the trustworthiness of the other player along with desire to protect oneself from being wrong about the other is almost certain to produce the second worst outcome for both us. In the Prisoner's Dilemma, players are tempted to defect from the cooperative outcome because doing so gives them unambiguous benefit. In the Stag Hunt, this is not so: each player is certain to lose if he unilaterally defects from the cooperative outcome. In both cases, however, mutual defection is likely to happen.

The advantage of a SH-like situation over a PD-like situation is that the social dilemma is solvable in principle in the first case but not in the latter. For instance, if we manage to coordinate expectations and attain a level of trust between ourselves, we will cooperate in SH but still will not cooperate in PD. The cooperative outcome can be sustained in equilibrium in SH but not in PD, which implies that one possible solution to cooperation failure in SH is to work on expectations.

5.2 The Arms Race: Which Model?

To see the conceptual difference between the PD and the SH, let us model an arms race as either a PD or a SH. Suppose it is determined that a new technology has just emerged and that it allows both us and our enemy to produce a super weapon that can guarantee winning a confrontation against an opponent who does not have it. The confrontation is very important. If both have the weapon, the effects cancel each other out. It takes a year to construct the weapon, but once built, it becomes immediately useful. The weapon is quite costly and each nation must shift resources from consumer goods to the military sector, which is politically unattractive. Should we build the weapon or not?

We have already simplified the situation drastically in this description. Let's now represent it with a game. There are two players, "us" and "they." Each has two options: defect and build the weapon, D, or cooperate and do not build it, C. There are four outcomes: both build the weapon (an arms race), only one builds the weapon (the one that does wins), or neither does (status quo). If only the enemy arms, we don't pay the cost of arming but lose the confrontation, which is really bad: *defeat*. If we arm and the enemy arms as well, then we pay the cost but since nobody can get the upper hand, no confrontation occurs: *arms race*. If we are the one side with the weapon, then we pay the cost but win the confrontation, which is really good: *victory*. If neither side arms, no confrontation occurs: *status quo*. If we arm, we pay the cost of doing so regardless of whether the weapon is used or not. The strategic form is:

		Player 2	
		C	D
Player 1	C	Status Quo	Defeat, Victory
	D	Victory, Defeat	Arms Race

Figure 34: The Arms Race.

5.2.1 As a Prisoner's Dilemma

For this scenario, assume that the confrontation is very important so that the benefits of winning it exceed the costs of producing the weapon. The preference ordering is as follows:

$$\text{Victory} \succ \text{Status Quo} \succ \text{Arms Race} \succ \text{Defeat}.$$

Note that victory is preferred to the status quo because the benefits from winning the confrontation are so high that even when we factor in the costs of building the weapon, it is still better than the status quo life with the enemy. The status quo, however, is preferred to an arms race because with an arms race we pay the costs of building the weapon but we don't get anything out of it except that the enemy can't defeat us, which is what the status quo already is. Finally, the arms race is preferred to defeat because losing is so disastrous that it is worse than avoiding the costs of building the weapon. Since the situation is symmetrical, our opponent has similar preferences.

These preferences give each individual player very strong incentives to build the weapon: each is strictly better better off doing so irrespective of what the other player does. If the other player is expected to cooperate by not building, defection yields Victory, which is preferable to the Status Quo. If the other player is expected to defect by building, defection yields an Arms Race, which is preferable to Defeat. C is strictly dominated by D, and thus the unique Nash equilibrium to this game is $\langle D, D \rangle$. The equilibrium outcome is an arms race: both players lose because they pay the costs of building the weapons but do not get any benefit from having them.

5.2.2 As a Stag Hunt

One possible objection to depicting the Arms Race dilemma as a PD is that it seems to require the actors to be aggressive in the sense that they both prefer to compel the other to capitulate than live with the status quo. Historically, even classic antagonists sometimes become essentially status quo powers over time. We could argue that the Arms Race had ceased to be a PD and had become a SH situation:

	Status Quo	$\succ$	Victory	$\succ$	Arms Race	$\succ$	Defeat
Player 1	(C, C)	$\succ$	(D, C)	$\succ$	(D, D)	$\succ$	(C, D)
Player 2	(C, C)	$\succ$	(C, D)	$\succ$	(D, D)	$\succ$	(D, C)

Table 2: Stag Hunt Preferences for the Arms Race Game.

Mutual disarmament would be the cooperative outcome which preserves the status quo and avoids the expense of building weapons. If the other side is expected to cooperate, then each player prefers to do so as well. On the other hand, if one fails to arm when the other one does, the disarmed player would be saddled with the worst possible outcome: defeat. Prudential reasoning suggests that the less risky choice is to arm: you would get your second-best choice is the opponent is cooperating and you would end up in an arms race if he defects as well. An arms race, while expensive, is much preferable to defeat. Small amounts of suspicion about the opponent's intent would then make $\langle D, D \rangle$ the more likely outcome.

The logic of the arms race in a SH-like scenario is fundamentally one of mistrust, risk-aversion, and prudential reasoning. The logic of an arms race in a PD-like scenario is one of desire to exploit the other side's cooperative effort combine with a desire to avoid being saddled with the worst possible outcome. In this sense, the Stag Hunt is probably captures the dynamics of fear-induced hostility much better than a Prisoner's Dilemma.

In international politics, one cannot know the intent and motivations of one's opponent (or partner). We cannot peek into the heads of decision-makers to verify that they do not intend to attack us, which is (of course) what they usually claim. Intentions are not only unverifiable, they are volatile. Changing governments, the particular mood of the leader, or many other factors may change the evaluation of the desirability of attack on a moment's notice. This is why states normally do not rely on intentions, they are forced to *infer* intent from *observable* capabilities and behavior.

This is where suspicion comes into play. If I cannot be certain that my opponent has no intention to attack me, I must admit the possibility (however small) that he might do so. Since being defeated is the worst possible scenario for me, prudential reasoning might lead me risk losing the cooperative outcome in favor of securing, at the very least, a costly preservation of the status quo. So I build some weapons to guarantee my security. Unfortunately, my act of increasing my security immediately decreases the security of my opponent. He would reason as follows: "I was almost sure that he did not have hostile intent but now I see him arming. I know he claims it is purely for defense but is that so? Perhaps he intends to catch me unprepared and defeat me? And even if that is not so, he clearly does not trust me enough or else he would not have started arming. I would like to reassure him that I can be trusted but the only way to do so is to remain unarmed, which unfortunately is very risky if he does happen to have aggressive intent. So I better arm just to make sure I will not have to surrender in that eventuality."

My opponent then arms as well, which makes me even less secure. We both have matched each other in armaments, the status quo survives, but we also learned that we cannot trust each other not to arm. Because we cannot observe intent, we can only see the arming decision which could be because the other side is afraid or it could be because the other side is aggressive. Reassurance being too risky, we opt for the prudential choice and continue arming, further increasing the suspicion and hostility. The process feeds on itself and rationalizes the non-cooperative outcome, just as in the original Stag Hunt story. The process, in which small doubts lead to defensive measures which increase the insecurity of the opponent, who reacts with defensive measures of his own, which increases my insecurity and as well as my doubts leading to further defensive measures on my part, is called the *Security Dilemma*, and it is very similar to the Stag Hunt scenario.

Notice that once the suspicion starts, it is in the interest of the players to restore trust and get the cooperative equilibrium. Unfortunately, trust can only be restored if one of the players decides to take the risk and plunge into unilateral disarmament. If his opponent turns out to have a SH preference structure (prefers the status quo without arms to victory), then this gesture would be reciprocated and the players could potentially go to a stable cooperative solution. If, on the other hand, one's opponent turns out to have a PD preference structure, then one risks defeat. If one suspects that the opponent has PD preferences or if one's opponent is so suspicious that he would ignore the gesture, no player would make the necessary first step to achieving cooperation.

What model you think represents the Arms Race problem best depends on what you think the structure of the preferences is. If you think of the Arms Race as a Prisoner's Dilemma, you would not recommend trust-building and risky unilateral actions: the opponent is sure to ignore anything you say and would not reciprocate restraint because exploiting your weakness is preferable to cooperation. If you think of the Arms Race as a Stag Hunt, on the other hand, you would recommend trust-building, and you might even recommend a dramatic unilateral gesture that runs serious risks but that can persuade the opponent of your peaceful intent.

5.3 Generic Conflict Games

We have now seen several canonical games like Chicken, the Stag Hunt, and the Prisoner's Dilemma. When the games involve no uncertainty either because of chance moves outside the players' control or because of mixed strategies, the precise values of the payoffs do not matter, only their ordinal ranking does. However, when the game does involve chance — as it must whenever some player uses a mixed strategy — then the cardinal values become important. Why it is so is a bit technical, but essentially it is because risky choices involve attitudes toward risks and the sizes of the payoffs loom large in those calculations. When I am running a 20% risk of disaster for an 80% chance of the other player capitulating, it certainly matters not merely that disaster is worse than him capitulating but also just how much worse it is. The worse it is, the less willing I become to take my chances. Von Neumann and Morgenstern's Expected Utility Theory in fact specifies the assumptions about preferences over risky choices we need to make in order to ensure that we can represent these preferences with numbers and calculate the resulting expected utilities.

Consider a generic two-player simultaneous-move game where each player has only two pure strategies: escalate (E) or not $\sim E$. We can represent it in a 2-by-2 payoff matrix, as in Fig. 35 (p. 45). The mnemonics for the variables are W for "war", V for "victory", D for "defeat", and S for "status quo".

		Player 2	
		$\sim E$	E
Player 1	$\sim E$	S, S	D, V
	E	V, D	W, W

Figure 35: The Generic Game.

We shall now see how varying the ordinal rankings among these variables yields all the games we have seen so far, and how we can glean some additional insights from representing them in this form. First, however, we shall make a crucial assumption that we shall maintain more or less throughout all models that we are going to analyze: we shall assume that our players are not war-loving and do not like defeat: they always prefer both the status quo and victory to either war or defeat. In our notation, we are going to assume that

ASSUMPTION 1. Players are not war-loving, $S \succ W$ and $V \succ W$, and want to avoid defeat, $S \succ D$ and $V \succ D$.

The only variation we shall allow is between the rankings of S and V — which we can think of as the strength of the incentive players have to take advantage of the cooperative behavior of the opponent (do they reward cooperation with restraint and obtain S or do they exploit it and obtain V), and the rankings of W and D — which we can think of as their fear of being exploited (do they prefer to let it happen and obtain D, or would they rather avoid it and obtain W).[20]

What can we say about this game? We know that $\langle E, e\rangle$ will be an equilibrium whenever $W \succ D$. Moreover, it will be the unique equilibrium if $V \succ S$ too. In other words, if the complete ordering is

$$V \succ S \succ W \succ D,$$

then the game is a **Prisoner's Dilemma** and its unique equilibrium yields the payoffs that are second-worst for the players. When the fear of being exploited ($W \succ D$) combines with a desire to take advantage of the other player ($V \succ S$), then players will be unable to coordinate on a cooperative outcome regardless of the amount of communication they are allowed to engage in.

If, on the other hand, $S \succ V$, then $\langle \sim E, \sim e\rangle$ will be an equilibrium as well. When the ordering is

$$S \succ V \succ W \succ D,$$

[20]We are making these assumptions because otherwise our insights will be superficial: it is not going to be very helpful if we found out that players go to war in equilibrium when they both value war the most. This is not to say that this cannot happen but that the analysis is trivial. It would be much more interesting if we found that players go to war in equilibrium even though war is among their least-preferred outcomes.

then the game is a **Stag Hunt**, and it has two-pure strategy equilibria, with $\langle \sim E, \sim e \rangle$ being the one both players prefer (it, in fact, yields the highest possible payoff for each player), but where $\langle E, e \rangle$ is risk-dominant, making it more likely for the players to coordinate on that profile and obtain their next-to-worst payoffs. Thus, making the status quo more attractive — which eliminates the desire to take advantage of the other player — can help, but the resulting situation (which still has the fear of being exploited looming as the worst possible outcome) still presents players with a difficult dilemma where the outcome can be very dependent on the amount of trust they have for each other. In most circumstances, this trust will not be enough to overcome to fear, and players will again end up with their next-to-worst outcome.[21]

You might be tempted to conclude that perhaps it is the fear of being exploited that is causing the problem here, so let's suppose players do not have it ($D \succ W$) but that they still want to take advantage of each other $V \succ S$. The resulting preference ordering will be

$$V \succ S \succ D \succ W,$$

and you can verify that this makes this a **Game of Chicken**. The two pure-strategy Nash equilibria are $\langle E, \sim e \rangle$ and $\langle \sim E, e \rangle$ but we know that there is going to be another one in mixed strategies as well. To find it, let p and q be probabilities with which player 1 and player 2 escalate, respectively. The expected payoff for player 1 can be computed as follows:

$$\begin{aligned} U_1(E, q) &= q u_1(W) + (1-q) u_1(V) = u_1(V) - q\big[u_1(V) - u_1(W)\big] \\ U_1(\sim E, q) &= q u_1(D) + (1-q) u_1(S) = u_1(S) - q\big[u_1(S) - u_1(D)\big]. \end{aligned}$$

We know that player 1 will only be willing to mix when indifferent between his pure strategies, so in the MSNE it must be the case that $U_1(E, q) = U_1(\sim E, q)$. Solving this tells us that player 1 will mix only when he thinks that player 2 is going to escalate with probability

$$q = \frac{u_1(V) - u_1(S)}{u_1(V) - u_1(S) + u_1(D) - u_1(W)}.$$

The preference ordering ensures that this is a valid probability. We further conclude that whenever player 1 is mixing, player 2 must be mixing as well, which in turn pins down the precise probability with which she must expect player 1 to escalate, which we derive by setting $U_2(p, E) = U_2(p, \sim E)$, or:

$$p = \frac{u_2(V) - u_2(S)}{u_2(V) - u_2(S) + u_2(D) - u_2(W)}.$$

We already know that in the MSNE the probability of war is positive, but we can say something more about the crisis. For example, we can ask questions like: "What happens to the probability that player 1 escalates if player 2's payoff from victory, $u_2(V)$, increases?" Try answering this first without analyzing the model. You might reason as follows: well, since player 2's payoff from victory is now larger than before and she can only get this outcome by escalating, she should be more willing to escalate. In other words, increasing the payoff for victory should make her more willing to take risks to achieve that outcome, so q should go up. But since this makes escalation more dangerous for player 1 and his payoffs have not changed, he should be less willing to escalate. Thus, the increase in the victory payoff for player 2 must mean that she is more likely to secure the prize without a fight, and that the overall likelihood of war is smaller.

The first surprise is that player 2 will not, in fact, escalate with a higher probability in equilibrium. As you can see from the expression above, q is entirely independent of $u_2(V)$. This is because in equilibrium

[21] In fact, the Stag Hunt, like the Chicken game, also has an equilibrium in mixed strategies. It is specified exactly in the same way as we shall do for the Chicken game, so there is no need to do it here.

her escalation probability reflects player 1's expectations about her behavior that make *him* indifferent, and this calculation naturally only involves player 1's payoffs. Since these have not changed, q will not change either.

But how can that be? Our intuition seems to demand that an increase in $u_2(V)$ must have some effect on behavior… and it does, just not where you would first expect it. Consider player 1's strategy. You can see that p is a function of $u_2(V)$, and you can easily verify that it is, in fact, strictly *increasing* in that value:

$$\frac{\mathrm{d}\, p}{\mathrm{d}\, u_2(V)} = \frac{u_2(D) - u_2(W)}{\left[u_2(V) - u_2(S) + u_2(D) - u_2(W)\right]^2} > 0.$$

In other words, increasing player 2's payoff from victory must make player 1 *more* likely to escalate in equilibrium! What?!?! This just made matters even more confusing!

This, however, what being "in equilibrium" really means. It means that players must be willing to stick to their strategies. Initially, player 2 is indifferent and so willing to play the mixed strategy. When her payoff from victory increases and nothing else changes, however, she will no longer be willing to mix: the expected payoff from escalation given the probability that player 1 escalates will now be strictly greater than the expected payoff from not escalating, and as a result she would actually strictly prefer to escalate. But if she is going to escalate, then player 1 will no longer be willing to mix either. In other words, the strategies would no longer constitute an equilibrium. If player 1 cannot predict what his opponent is going to do in equilibrium (i.e., player 2 is mixing), then it must be that player 2 is expected to continue to be indifferent after $u_2(V)$ increases. Since none of the other payoffs have changed, the only way this can happen is through an increase in player 1's probability of escalation (which makes her bad outcome more likely). Since this puts more weight on the war outcome, it decreases the expected payoff from escalation for player 2 even when $u_2(V)$ goes up. Thus, if the mixed strategies are going to remain optimal, an increase in $u_2(V)$ will be met with an increase in p.

In other words, our intuitive logic has some parts right (e.g., that increasing $u_2(V)$ will make player 2 prefer escalation) but fails to consider the entire effect (e.g., what happens when you put this fact together with the requirement that players choose best responses). This is why simple intuition might sometimes prove quite misleading.

Finally, observe that since p goes up and q remains constant, an increase in $u_2(V)$ also leads to an increase in the equilibrium probability of war, which is Pr(War) $= pq$. Thus, an increase in the value for victory for one of the players makes the other one more aggressive, and it makes it more likely that they will end up fighting.

Analogous arguments establish that when a player's value for war increases, then the probability with which his opponent escalates in equilibrium must increase as well (p is increasing in $u_2(W)$ just like q is increasing in $u_1(W)$). This also seems counter-intuitive: a player's dislike of fighting decreases but as a result his opponent becomes more likely to escalate. The overall effect might be less surprising: the equilibrium probability of war increases.

Conversely, when a player's value for the status quo increases, then his opponent's probability of escalation must go down (p is decreasing in $u_2(S)$). This is surprising when you recall that the opponent prefers to take advantage of such failures to escalate. The overall effect, however, might be what you expect: the equilibrium probability of war decreases. At least we obtain an unambiguous prediction: if one is interested in preserving peace, then making the status quo more valuable (or war more costly) is the way to go.

6 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.

6.1 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.

6.2 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.

6.3 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.)

6.4 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.

6.5 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).

7 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,[22] convex,[23] nonempty subset of a finite-dimensional Euclidean space;[24]
2. $r(\sigma)$ is nonempty for all σ;
3. $r(\sigma)$ is convex for all σ;
4. r is upper hemi-continuous.[25]

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 $\alpha \in (0, 1)$,

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

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. ■

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.

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

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

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

[25] 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*.