

Math 1007: Mathematics and Politics
The George Washington University Fall 2025

Jay Daigle

Contents

1	Voting	2
1.1	Two-Candidate Elections	2
1.1.1	Introductory Scenario	2
1.1.2	Defining our terms	3
1.1.3	Some methods of voting	5
1.1.4	Voting method criteria	9
1.1.5	May's Theorem	12
1.2	Social Choice Functions	14
1.2.1	Ballots and Simple Functions	14
1.2.2	Voting System Criteria	22
1.3	Evaluating Social Choice Functions	26
1.3.1	Plurality method	26
1.3.2	Antiplurality method	27
1.3.3	Borda count	28
1.3.4	Hare's method	29
1.3.5	Coombs's method	31
1.3.6	Copeland's method	32
1.3.7	Black's method	33
1.4	Arrow's Theorem	36
2	Apportionment	38
2.1	Congressional Apportionment and Hamilton's Method	38
2.1.1	Defining apportionment	41

2.1.2	Quotas	43
2.1.3	Hamilton’s Method	43
2.1.4	”Paradoxes” in Hamilton’s method	45
2.2	Divisor Methods	48
2.2.1	Jefferson’s Method	48
2.2.2	Critical Divisors	51
2.2.3	Results on Jefferson’s Method	53
2.2.4	Other divisor methods	55
2.2.5	Other rounding methods	58
2.3	Evaluating Apportionment Methods	63
2.3.1	House Monotonicity	65
2.3.2	Population Monotonicity	65
2.3.3	Relative population monotonicity	67
2.3.4	The New States Paradox	70
2.3.5	The Impossibility Theorem	71
2.4	Balinski and Young Apportionment	73
2.4.1	The Quota Rule for Balinski and Young	76
2.5	Why choose different rounding functions?	77
2.5.1	Degree of Representation and Webster’s Method	78
2.5.2	District Size and Dean’s Method	78
2.5.3	Average district size and Hill’s Method	78
3	Conflict	80
3.1	Zero-Sum Games	80
3.1.1	Naive and Prudent Strategies	81
3.1.2	Best Response and Saddle Points	83
3.1.3	Dominant Strategies	86
3.2	Expectation and Probability	90
3.2.1	Random Variables and Expected Value	91
3.2.2	Mixed Strategies in Games	94
3.2.3	Independent processes and independent probabilities	96
3.3	Solving Zero-Sum Games	98
3.3.1	Prudent mixed strategies	101
3.3.2	Solving 2 by 2 games	103
3.4	Conflict and Cooperation	106

3.4.1	Guarantees and Saddle Points	107
3.4.2	Cooperation and Competition	109
3.4.3	Some important games	109
3.5	Nash Equilibria	111

3 Conflict

3.1 Zero-Sum Games

Example 3.1 (Roshambo/Rock Paper Scissors). Two players face each other and simultaneously choose from three options: rock, paper, or scissors. Rock beats scissors, scissors beats paper, and paper beats rock. Two of the same selection tie.

We can summarize this game in a matrix like this:

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

Figure 3.1: The payoff matrix for Rock Paper Scissors

The row reflects the choice of player 1, and the column reflects the choice of player 2. (We will sometimes just call these players “row” and “column”.) The number in the cell represents the result *for player 1*. So in the (Rock, Paper) cell, we have a -1 because player 1 loses.

Definition 3.2. A *two-person zero-sum game* is a game featuring two players, in which each player adopts a *strategy*, and then the combination of strategies determines a number called the *payoff*. We can think of the payoff as the amount of money player 2 has to pay player 1; if it is negative, this means player 1 has to pay player 2.

Remark 3.3. The phrase “zero-sum” means that player 1 wins the same amount that player 2 loses. In a competitive one-on-one game like Rock Paper Scissors, that has to be true. But it is often possible for both players to win, or lose, simultaneously. We will discuss this idea further in section 3.4.

We can represent a two-player zero-sum game with a *matrix*, a grid of possible payoffs. If Row has m possible strategies and Column has n strategies, we have a $m \times n$ matrix, with m rows and n columns. (I always have trouble remembering whether the rows or the columns come first. But we will always use rows for player 1, and the number of rows will be written first.) The entry in the i th row and the j th column will be written $u_{i,j}$ or sometimes u_{ij} . We call this matrix the *payoff matrix*.

Rock Paper Scissors is obviously a game for children. But the same logic applies in many scenarios with much higher stakes.

Example 3.4. In World War II, the Japanese navy needed to re-supply their base on New Guinea. In order to do this, they needed to sail a convoy around the island of New Britain, passing either to the North or to the South of the island. The United States controlled New Britain and wanted to stop the Japanese convoy by bombing it. The resulting battle is known as the Battle of the Bismarck Sea.

The US was certain to locate the Japanese convoy, but depending on how quickly they found them, they would be able to achieve either more or fewer days of bombing. General George Kenny gave estimates of the number of days of bombing under four possible scenarios. First he considered what would happen if the Japanese sailed north of New Britain. In that case, if the United States searched to the north, the United States would have two days of bombing, but if the United States searched to the south, the United States would have only one day. On the other hand, if the Japanese sailed south and the United States searched south the United States would have three days of bombing, but if the Japanese sailed south and the United States searched north, the United States would have two days of bombing.

We can interpret this as a two-player game between the US and the Japanese. It's zero-sum because the US wanted to maximize the bombing, and the Japanese wanted to minimize it. Thus we can represent the payoffs in the following matrix:

		Japan	
		North	South
United States	North	2	2
	South	1	3

Figure 3.2: The payoff matrix for the Battle of the Bismarck Sea

This matrix allows us to abstract away the details of the situation and just think about the most general version of the structure. That can lose important detail, but can also let us focus in on the fundamental logic of the conflict, which can be extremely useful!

3.1.1 Naive and Prudent Strategies

Definition 3.5. The outcome that gives a player their best possible payoff is the *primary outcome*. For Row this is the largest entry in the payoff matrix; for Column it is the smallest,

or most negative, entry.

Definition 3.6. In the *naive method*, a player chooses the strategy corresponding to their primary outcome. This is called the player’s *naive strategy*, or sometimes the *greedy strategy* or *optimistic strategy*. If both players play their naive strategies, we get the *doubly naive outcome*.

Example 3.7. In Rock Paper Scissors, Row’s primary outcome is a win, represented by 1; Column’s primary outcome is a win, or a loss to Row, represented by -1 . Any strategy is a naive strategy in this case because it is possible to win with any strategy.

In this case it’s not clear what the doubly naive outcome is, exactly, because any outcome can correspond to a doubly naive strategy.

Example 3.8. In example 3.4, the US’s primary outcome is three days of bombing, achieved by searching to the south if the navy goes to the south. Japan’s primary outcome is one day of bombing, if the US searches to the south and Japan goes to the north. So the doubly naive outcome is when the US searches to the south and Japan goes to the north, resulting in one day of bombing.

Example 3.9. Novice chess players often make chess moves based on what could maybe happen if their opponent blunders or walks into their trap. When I studied chess as a kid this was called “hope chess”. (As we’ll see, it was not a recommended approach.)

We can also be less optimistic and more careful.

Definition 3.10. The worst payoff a player can get from a given strategy is that strategy’s *guarantee*. (The player is guaranteed to get at least that good a payoff.)

In the *prudent method*, a player chooses the strategy with the best guarantee, called the *prudent strategy*. We might also call this the *pessimistic strategy*, because it wants to minimize the damage from the worst-case scenario.

If both players play their prudent strategies, we get the *doubly prudent outcome*.

Remark 3.11. Sometimes people will call the naive strategy the “maximax”, because it maximizes the maximum payoff. The prudent strategy is “minimax”, because it minimizes the maximum loss. (More rarely, it’s called “maximin” because it maximizes the minimum payoff.)

Example 3.12. Again, in Rock Paper Scissors, every strategy is a prudent strategy, because every strategy has the same maximum loss.

Example 3.13. In the Battle of the Bismarck Sea, we need to think about each player separately. The US has a guaranteed of 2 for North, and a 1 for South, so their prudent strategy is to search North. Japan has a guarantee of 2 for North, and a 3 for South. (Remember that for Japan, lower numbers are better!) So their prudent strategy is also to pick North.

We can represent this logic in the following *min-max diagram*:

		Japan		
		North	South	
United States	North	2	2	2 ←
	South	1	3	1
		2	3	
		↑		

Figure 3.3: The min-max diagram for the Battle of the Bismarck Sea

Proposition 3.14. *Let r be the guarantee of Row's prudent strategy, and c be the guarantee of Column's prudent strategy. If Row plays a prudent strategy, their payoff will be at least r ; if they play a non-prudent strategy, it is possible their payoff will be lower. Similarly, if Column plays a prudent strategy, their payoff will be at most c , and if they do not, there is a possibility their payoff will be greater.*

Corollary 3.15. *If r is the guarantee of Row's prudent strategy and c is the guarantee of Column's prudent strategy, then $r \leq c$.*

Proof. Suppose row i is a prudent strategy, and column j is a prudent strategy. Then the payoff for the cell corresponding to this doubly prudent strategy is $u_{i,j}$. We must have $u_{i,j} \geq r$, and we must have $u_{i,j} \leq c$, and combining these two statements tells us that $r \leq u_{i,j} \leq c$. □

3.1.2 Best Response and Saddle Points

An obvious move to make while playing two-player games is to try to guess what your opponent is going to do, and respond to that. (Indeed this is about the only way to play Rock Paper Scissors.)

Definition 3.16. A strategy choice by one player is called the *best response* to an opponent strategy if it gives the best payoff against that strategy.

The best response to a naive strategy is called the *counter-naive strategy*. The best response to a prudent strategy is called the *counter-prudent strategy*.

We can summarize information about best responses in a diagram called a *flow diagram*. We lay out a grid in the same pattern as the payoff matrix, but notated with arrows: in each column, we have a vertical arrow pointing to the largest entry, and in each row, we have a horizontal arrow pointing to the smallest entry. These arrows indicate each player’s best response. (You can think of them as giving a map for how to respond if you know what row or column your opponent will select.)

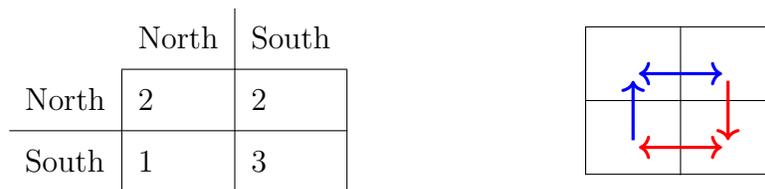


Figure 3.4: The flow diagram for the Battle of the Bismarck Sea

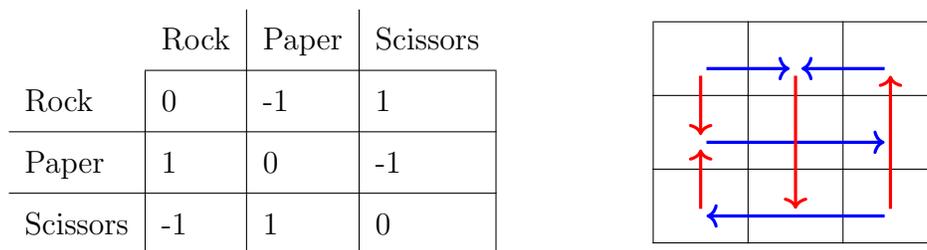


Figure 3.5: The flow diagram for Rock Paper Scissors

Example 3.17 (Battle of the Bismarck Sea).

“backwards induction”

Definition 3.18. A *saddle point* is an outcome such that the strategy for each player is the best response to the strategy of the opponent, simultaneously .

A *saddle point strategy* is a strategy that corresponds to a saddle point outcome.

There are various ways of conceptualizing this. An outcome is a saddle point if and only if all the arrows in its row and column point to it. We can also say that a point in row k

and column ℓ is a saddle point if

$$\begin{aligned} u_{k,\ell} &\leq u_{k,j} && \text{for any column } j \\ u_{k,\ell} &\geq u_{i,\ell} && \text{for any row } i. \end{aligned}$$

Saddle points are the simplest example of what game theorists call a *Nash equilibrium*; we'll discuss the more general category in section 3.5.

Theorem 3.19. *A two-person zero-sum game has a saddle point if and only if $r = c$. In that case, the saddle point is a doubly prudent outcome, and the payoff is $r = c$.*

Proof. Since this theorem says “if and only if”, we will need to prove two things.

First suppose $r = c$. Let row k and column ℓ be strategies with guarantee r , so that (k, ℓ) is a doubly prudent outcome. We can conclude that $r \leq u_{k,j}$ for all j , and thus column ℓ is a best response to row k . And $r \geq u_{i,\ell}$ for all i , which means that row k is a best response to column ℓ . Thus (k, ℓ) is a saddle point.

Conversely, suppose (k, ℓ) is a saddle point, and set $v = u_{k,\ell}$. Since we have a saddle point, we know that $u_{k,\ell}$ must be the smallest entry in its row and the largest entry in its column; this implies that the guarantees for both row k and column ℓ must be v .

But we defined r to be the largest guarantee across all rows, so $r \geq v$. And c is the smallest guarantee among the columns, so $c \leq v$. But we know from corollary 3.15 that $r \leq c$, so we have $c \leq v \leq r \leq c$ implying all three numbers are the same. This shows that $r = c$, and also that row k and column ℓ are prudent strategies since they both give our best guarantee. \square

Corollary 3.20. *An outcome (k, ℓ) in a zero-sum game is a saddle point if and only if row k is a saddle point strategy and column ℓ is a saddle point strategy.*

Remark 3.21. This might seem obvious; one direction is in fact as obvious as it seems. The trick is that, at least in theory, there might be more than one saddle point. (In fact this is possible, though mostly in slightly silly ways.) What we need to show is that if (i, j) is a saddle point, so that row i is a saddle point strategy; and also (k, ℓ) is a saddle point, so that column ℓ is a saddle point strategy; then (i, ℓ) must also be a saddle point.

Proof. If (k, ℓ) is a saddle point, then row k and column ℓ are saddle point strategies by definition.

Now suppose row k and column ℓ are saddle point strategies. That means there's a saddle point, say (k, j) , in row k ; and a saddle point, say (i, ℓ) in column ℓ . We want to prove that (k, ℓ) is also a saddle point.

Since the game has a saddle point, we know from theorem 3.19 that $r = c$, and row k and column ℓ are prudent strategies. That implies that (k, ℓ) is doubly prudent, which means that (k, ℓ) must be a saddle point. \square

This means we can find saddle points using a min-max diagram. Specifically, we have saddle points whenever $r = c$. We can see this at work in this 5-by-5 diagram:

2	3	2	2	5	2 ←
0	10	0	3	-9	-9
-2	2	-1	2	7	-2
2	10	2	2	2	2 ←
0	4	1	0	-4	-4
	2	10	2	3	7
	↑		↑		

Figure 3.6: Using a min-max diagram to find saddle points

We see that each player has two prudent strategies. Row 1 and row 4 are prudent, as are column 1 and column 3. That means there are four saddle points: $(1,1)$, $(1,3)$, $(4,1)$, $(4,3)$ are all saddle points. Any one of these guarantees a payoff of 2.

Corollary 3.22. *A strategy in a two-person zero-sum game is a saddle point strategy if and only if it is both prudent and counter-prudent.*

Proof. We know a saddle point strategy is prudent from Theorem 3.19. But in a saddle point, each player is making the optimal response to the other player, by definition. Thus it is the best response to the other player's prudent strategy, and is counter-prudent. \square

3.1.3 Dominant Strategies

Definition 3.23. We say that one row of a matrix *dominates* another if each entry of the first row is at least as large as the corresponding entry of the second row, and at least one entry is strictly larger.

Algebraically: row k dominates row i if $u_{k,j} \geq u_{i,j}$ for each column j , and there is at least one j such that $u_{k,j} > u_{i,j}$.

We say row k *strictly dominates* row i if $u_{k,j} > u_{i,j}$ for each column j .

Example 3.24. Consider the game

1	2	4
5	3	6

Row 2 dominates row 1 because every entry in row 2 is larger than the corresponding entry in row 1. ($5 > 1$, $3 > 2$, and $6 > 4$). In fact, row 2 strictly dominates row 1.

In the game

1	2	3
5	2	3

row 2 dominates row 1, but not strictly.

Example 3.25. In the games

1	2	3
5	1	6

1	2	3
1	2	3

neither row dominates the other.

We can give similar definitions for column dominance. But in this case, it's important to remember that for the column player, a smaller number is better.

Definition 3.26. We say that one column of a matrix *dominates* another if each entry of the first column is at least as small as the corresponding entry of the second column, and at least one entry is strictly smaller.

Algebraically: column ℓ dominates column j if $u_{i,\ell} \leq u_{i,j}$ for each row i , and there is at least one i such that $u_{i,\ell} < u_{i,j}$.

We say column ℓ *strictly dominates* column j if $u_{i,\ell} > u_{i,j}$ for each row i .

Example 3.27. In the game

1	2	4
5	4	6

columns 1 and 2 strictly dominate column 3, but neither 1 nor 2 dominates the other.

The basic idea is that dominated strategies are obviously bad, and no rational/reasonable player would ever select one. Therefore you can remove them from the game without affecting the analysis. Removing dominated strategies is called *reduction*. Sometimes removing some columns causes one row to now be dominated, and you can do further reduction; if a matrix has been reduced and cannot be reduced further, that is a *complete reduction*.

Example 3.28. Let's follow this process for the Battle of the Bismarck Sea:

		Japan	
		North	South
United States	North	2	2
	South	1	3

Neither row dominates the other, but column 1 dominates column 2. So we can reduce this game to

		Japan	
		North	
United States	North	2	
	South	1	

Now the US has a dominated strategy; in this reduced game, South is dominated by North. Removing South as an option leaves a completely reduced game:

		Japan	
		North	
United States	North	2	

This is essentially the same conclusion we reached using flow diagrams in example 3.17: North is prudent for both players.

Example 3.29. Consider the game

0	-1	-2	5	4
-3	1	2	3	6
-4	-5	-6	-7	7

Can ask all the questions we already asked. Row's naive strategy is row 3, aiming for the 7. Column's naive strategy is column 4. (How does the naive strategy work out for them?)

Row's prudent strategy is row 1 with -2, and Column's prudent strategy is column 1, with the 0. How does this work out for them?

Column's counter-naive strategy is column 4, and Row's counter-naive strategy is row 1. Column's counter-prudent strategy is column 3, and Row's counter-prudent strategy is row 1.

We can ask things like, what is the counter-counter-prudent strategy? If Row is playing the counter-prudent strategy they pick row 1, so Column should pick column 3. If Column is playing the counter-prudent strategy they pick column 3, so Row should pick column 2.

We could draw out a full flow diagram, but easier to look for dominated strategies. Probably pick column first, since it's easier for things to be dominated there. The last column is quite bad, and is dominated by all the first three columns, so we can remove that.

0	-1	-2	5
-3	1	2	3
-4	-5	-6	-7

Note this removes Row's original naive strategy goal; this is why the naive strategy is naive!

Given the removal of the last column, we now have row 1 dominating row 3, so we can get rid of row 3 entirely. Then column 4 is dominated by any other column and we lose that as well, getting this game:

0	-1	-2
-3	1	2

Now it makes sense to build a flow diagram for this completely reduced game:

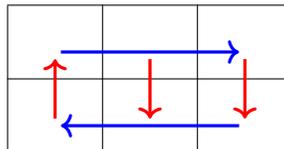


Figure 3.7: The flow diagram for example 3.29

We can see this game has no saddle point, because every cell has an arrow leaving it. That means the original game can't have a saddle point either. We can assume that reasonable players will only make plays corresponding to plays in the reduced form of the game: Row will pick between rows 1 and 2, while Column will pick columns 1, 2, and 3. But there's not a stable strategy here.

To talk about this situation better we need to understand a little about randomness and probability.

3.2 Expectation and Probability

We need to talk about randomness, which mathematicians think about as a “theory of probability”. This is unfortunately going to require us to develop a lot of terminology.

We want to study random processes. The possible results of the random processes are *outcomes* and the set of all possible outcomes is the *sample space*. (This term comes from statistics, where we think about taking a random sample of all the people we could interview, or something like that.)

Example 3.30. Tossing a coin has possible outcomes “heads” and “tails”. We might write the sample space as $\{h,t\}$.

Rolling a six-sided die has six possible outcomes. The sample space is $\{1, 2, 3, 4, 5, 6\}$.

To describe how likely each outcome is, we assign it a number between 0 and 1, called the *probability*. If we have n possible outcomes, we call the probabilities p_1, p_2, \dots, p_n . Each probability satisfies $0 \leq p_k \leq 1$; we think of $p_k = 0$ as “impossible” and $p_k = 1$ as “guaranteed”. Further, in this formulation exactly one outcome will happen, so $p_1 + p_2 + \dots + p_n = 1$; this represents the fact that it is guaranteed that one and only one outcome will occur.

We often group these probabilities together into a list $P = (p_1, p_2, \dots, p_n)$, which we call a *probability distribution*.

In ordinary English we often refer to probabilities as percentages, in which case we often use the word *chance*. “The chance of rain tomorrow is 50%” is equivalent to saying the probability of rain tomorrow is $p = 1/2$.

Example 3.31. In a (fair) coin toss, the probability of heads and tails are each 50%. We can write $P = (1/2, 1/2)$.

When rolling a die, we get the probability distribution $P = (1/6, 1/6, 1/6, 1/6, 1/6, 1/6)$.

Remark 3.32. In this formulation, we assume we have listed every possible outcome. This is always a simplification; for instance, we didn’t include the possibility that a coin will land on edge. But it’s also a useful assumption, and it pairs well with our game-theory assumption that we’ve listed every possible strategy.

It is possible to do probability theory when there are infinitely many possible outcomes. But this involves calculus, so we won’t be talking about it any further in this course. (Again, our games will only have finitely many possible strategies, so this doesn’t introduce any limitations that will be relevant to us.)

It's very important to note that sometimes, one outcome is more likely than another. That might seem obvious, but it's also a point that is frequently quite subtle. But if I pick a day at random, the probability of getting a weekend is $2/7$. If I pick a shirt at random from my closet, it's much more likely to be blue than yellow.

In common language, we often say "choose at random" to mean every choice is equally likely. We want a broader sense of randomness, so we will say "equally likely" when we mean that. If we want to be especially fancy, we can say the *uniform distribution* on n outcomes is $P = (1/n, 1/n, \dots, 1/n)$.

In many cases, we have non-uniform distributions "build out of" uniform distributions. For instance, suppose I want a probability distribution $P = (1/6, 1/4, 1/3, 1/4)$. We can check these probabilities sum to 1. More usefully, we can write these as $(2/12, 3/12, 4/12, 3/12)$. If we want to produce one of these numbers at random, we can pick twelve objects, and label them say 1, 2, 3, or 4. If we label two objects with a 1, three with a 2, four with a 3, and three with a 4, we get the desired probability distribution.

We can do the same thing with a spinner; we can chop a wheel into segments of appropriate size and spin something around on that wheel. People who regularly need to generate random numbers like this, such as poker players, will often use the seconds hand of a clock, or the ones digit of the seconds on a digital clock.

3.2.1 Random Variables and Expected Value

In this section we're going to have to introduce a couple of terms that were really badly chosen; neither of them means anything like what it sounds like they mean. Unfortunately, they were badly chosen in the early 1900s and now we're stuck with them.

Definition 3.33. Suppose we have a random process on a sample space with n outcomes and a probability distribution $P = (p_1, p_2, \dots, p_n)$. A *random variable* X on this sample space is a function that assigns a real number to each of the n possible outcomes.

Mathematically, a random variable X is a function that takes in an outcome from the sample space, and outputs some real number. Since there are n possible outcomes, we can write X as x_1, x_2, \dots, x_n .

Conceptually, a random variable is trying to measure the value, or benefit, of each outcome. If we're going to roll a six-sided die, it makes sense to say a 1 is worth 1 point, a 2 is worth 2 points, and so on, so we get $x_i = i$. If we're going to randomly get a penny, a dime, or a quarter, we might say the random variable has $x_1 = 1, x_2 = 10, x_3 = 25$. In this

context, the random variable is representing the *payoff* a gambler or gamer gets from each possible outcome.

This leads us to our second badly-chosen definition.

Definition 3.34. Consider a random process on a sample space with n outcomes, with a probability distribution $P = (p_1, p_2, \dots, p_n)$. Let X be a random variable that assigns the payoff x_k to outcome k . Then we define the *expected value* of X to be

$$E = E(X) = p_1x_1 + p_2x_2 + \cdots + p_nx_n.$$

Conceptually, the expected value is the average payout you get from this random process. If you play the game a hundred times, you will probably get a total payout of about $100 \cdot E(X)$.

Example 3.35. What is the expected value of rolling a six-sided die? We get

$$\begin{aligned} E &= \frac{1}{6} \cdot 1 + \frac{1}{6} \cdot 2 + \frac{1}{6} \cdot 3 + \frac{1}{6} \cdot 4 + \frac{1}{6} \cdot 5 + \frac{1}{6} \cdot 6 \\ &= \frac{1}{6} (1 + 2 + 3 + 4 + 5 + 6) = \frac{1}{6} \cdot 21 = 3.5. \end{aligned}$$

Example 3.36. Suppose a lottery pays \$100,000,000 with a probability $1/150,000,000$. It also has lower prizes: \$200,000 with probability $1/3,000,000$, or \$10,000 with probability $1/150,000$, or \$10 with probability $1/300$. Let's find the expected value of a ticket.

There are five outcomes to this lottery. We can build the following table:

Prize:	Grand	Second	Third	Fourth	Lose
Payoff:	\$100,000,000	\$200,000	\$10,000	\$10	\$0
Probability:	$1/150,000,000$	$1/3,000,000$	$1/150,000$	$1/300$	0.997

Then the expected value is

$$\begin{aligned} E &= \frac{100,000,000}{150,000,000} + \frac{200,000}{3,000,000} + \frac{10,000}{150,000} + \frac{10}{300} + 0 \cdot 0.997 \\ &\approx 0.667 + 0.067 + 0.067 + 0.033 + 0 \\ &\approx 0.833. \end{aligned}$$

Thus the expected payoff is about 83 cents.

In this case, you'd absolutely take a free ticket, right? On average it's worth 83 cents. But suppose the ticket costs a dollar. Then instead you'd get the following table

Prize:	Grand	Second	Third	Fourth	Lose
Payoff:	\$99,999,999	\$199,999	\$9,999	\$9	−\$1
Probability:	1/150,000,000	1/3,000,000	1/150,000	1/300	0.997

Then the expected value is

$$\begin{aligned}
 E &= \frac{99,999,999}{150,000,000} + \frac{199,999}{3,000,000} + \frac{9,999}{150,000} + \frac{9}{300} + (-1) \cdot 0.997 \\
 &\approx 0.667 + 0.067 + 0.067 + 0.033 + -1 \\
 &\approx -0.167.
 \end{aligned}$$

On average you lose about 17 cents per ticket.

Poll Question 3.2.1. In example 3.35, do you expect to get a 3.5?

In example 3.36, do you expect to lose 17 cents?

I said at the start that the term “expected value” isn’t really well chosen. It’s the average you get, but often you don’t expect to get the expected value. In a lottery I *expect* to get zero, but because that’s by far the most likely result. But because there are rare very valuable outcomes, the expected value of the ticket is 83 cents.

However, expected value is a useful framework for analyzing decisions, especially in contexts where similar decisions will be made repeatedly. (We’ll try to return to this idea of “repeatedness” at the end of the course.) Therefore, we will use it as a measure of how good an outcome is.

Definition 3.37. Given a choice among two random variables on two random processes, the *expected value principle* says a rational player will choose the one with the largest expected value.

We wrote this as a definition of “expected value principle”, but you can also think of it as a definition of “rational” in the context of game theory and other sorts of decision theory and economic analysis.

One important caveat is that we don’t always want to compute expected value in *dollars*. Not all dollars are created equal, in the sense that for most people, losing fifty thousand dollars would be much more than fifty times as bad as losing one thousand dollars.

This is the principle insurance works on. The expected value of an insurance policy is usually negative: the insurance company takes in more in premiums than it pays out in claims. (This is how they make money.) The reason to buy insurance is that you are turning a small probability of a catastrophe into a predictable, manageable expense.

When we want to do an analysis that deals with issues like this, we sometimes talk about *utility*, the value to a person of having something. The idea I expressed above is sometimes describes as the diminishing marginal utility of money: your first dollar is more useful than your millionth dollar. In a lot of economic analysis, this is implemented by assuming utility is the logarithm of income; we won't be doing that because we don't want to talk about logarithms.

Instead, we'll assume games are already expressed in terms of utility, or whatever payoff we actually care about. For the purpose of this formalism, we're generally going to be trying to maximize the expected value of our game-playing strategy.

3.2.2 Mixed Strategies in Games

We now have the tools to properly analyze Rock Paper Scissors. Recall the payoff matrix:

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

This game has no saddle points; every strategy is simultaneously naive and prudent. But for any strategy you pick, the other player has a counter-strategy that will cause you to lose.

Any predictable strategy will predictably lose in this game, because if an opponent can predict your moves, they will play the counterplays. We need to introduce unpredictability, which means randomness.

Definition 3.38. Consider an $m \times n$ zero-sum two-player matrix game.

The original strategy choices in a game, corresponding to single rows or columns of the matrix, are called *pure strategies*.

A *mixed strategy* for Row is a probability distribution $P = (p_1, \dots, p_m)$ on their set of m pure strategies. In principle, Row will choose one row, or one pure strategy, at random, choosing row k with probability p_k .

A mixed strategy for Column is a probability distribution $Q = (q_1, \dots, q_n)$ on their set of n pure strategies.

Example 3.39. Suppose Column plays Rock Paper Scissors with a mixed strategy $Q = (1/4, 1/2, 1/4)$. This means they will randomly play Rock 1/4 of the time, Paper 1/2 of the

time, and Scissors $1/4$ of the time. (This is why gamblers will generate random numbers from their watches.)

If Column commits to this mixed strategy, the game for Row becomes, in effect, a lottery.

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0
	$1/4$	$1/2$	$1/4$

For instance, suppose Row plays Rock. Now we see they have a $1/4$ chance of getting 0, a $1/2$ chance of getting -1 , and a $1/4$ chance of getting 1. The expected value of Rock is

$$E(\text{Rock}) = \frac{1}{4} \cdot 0 + \frac{1}{2} \cdot (-1) + \frac{1}{4} \cdot 1 = -1/4.$$

On average, Row will lose $1/4$ of a point for every game they play.

We can also calculate the other expected values. We have

$$E(\text{Paper}) = \frac{1}{4} \cdot 1 + \frac{1}{2} \cdot 0 + \frac{1}{4} \cdot (-1) = 0$$

$$E(\text{Scissors}) = \frac{1}{4} \cdot (-1) + \frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 0 = 1/4$$

Thus with a strategy of Scissors, Row will win a quarter of a point per game on average.

It can be convenient to talk about a pure strategy as just a specific kind of mixed strategy.

Definition 3.40. A *basic mixed strategy* is a probability distribution with every probability except one equal to 0. We write P_i for Row's basic mixed strategy that sets $p_i = 1$ and thus plays row i ; we write Q_j for Column's basic mixed strategy that sets $q_j = 1$ and thus plays column j .

We want to think about what happens when both players are playing mixed strategies. This makes all the payoffs uncertain; since each strategy has a random component, the payoff is also partially random. So we're mostly going to be concerned with the expected value of a given mixed strategy.

We can start by talking about what happens when one player plays a pure strategy against another mixed strategy; this will be an important tool for what comes later.

Lemma 3.41. Consider a $m \times n$ matrix game. If Row plays a pure strategy of row i against Column's mixed strategy $Q = (q_1, \dots, q_n)$, then the expected value of the payoff is

$$E(P_i, Q) = q_1 u_{i,1} + q_2 u_{i,2} + \dots + q_n u_{i,n}.$$

Similarly, if Column plays row j against Row's mixed strategy $P = (p_1, \dots, p_m)$, then the expected value of the payoff is

$$E(P, Q_j) = p_1 u_{1,j} + p_2 u_{2,j} + \dots + p_m u_{m,j}.$$

Corollary 3.42. If P_i and Q_j are basic mixed strategies, then $E(P_i, Q_j) = u_{i,j}$.

3.2.3 Independent processes and independent probabilities

Often we want to think about two separate random processes. In our specific context, we may have Row choosing a strategy randomly, and Column also choosing a strategy randomly.

Definition 3.43. Suppose we have two random processes, one given by a probability distribution $P = (p_1, \dots, p_m)$ and the other by a probability distribution $Q = (q_1, \dots, q_n)$. We say the processes are *independent* if the probability of the compound outcome (i, j) is $p_i q_j$.

Often it's important in life to figure out if two processes are actually independent.

Example 3.44. If we flip two (fair) coins, the first coin has outcomes h, t with probabilities 0.5, 0.5, and the second coin has outcomes h, t with probabilities 0.5, 0.5.

If we flip both coins, there are four outcomes hh, ht, th, tt , each with probability $0.5 \cdot 0.5 = 0.25$. These processes are independent.

Example 3.45. If we flip two coins that are taped together head-to-tail, then either both will land heads up or both will land tails up. Then coin 1 has a probability distribution of $(0.5, 0.5)$ and coin 2 has a probability distribution of $(0.5, 0.5)$, but the odds that both coins will land heads up is *also* (0.5) , and we get the distribution $(0.5, 0, 0, 0.5)$. This distribution is not independent.

Pretty much any time we talk about “flipping two coins” or something, we mean to say that we are flipping them independently.

Example 3.46.

This allows us to compute expected values for mixed strategies played against other mixed strategies.

Example 3.47. Suppose in Rock Paper Scissors Row plays the mixed strategy $P = (1/4, 1/2, 1/4)$, and Column plays $Q = (1/6, 1/3, 1/2)$. We can represent these strategies by writing them along the sides of the payoff matrix.

	Rock	Paper	Scissors	
Rock	0	-1	1	1/4
Paper	1	0	-1	1/2
Scissors	-1	1	0	1/4
	1/6	1/3	1/2	

Since the probabilities are independent, we can work out the probability of landing in each cell:

	Rock	Paper	Scissors	
Rock	1/24	1/12	1/8	1/4
Paper	1/12	1/6	1/4	1/2
Scissors	1/24	1/12	1/8	1/4
	1/6	1/3	1/2	

Now we can compute the expected payoff of this pair of mixed strategies:

$$\begin{aligned}
 E(P,Q) &= \frac{1}{24} \cdot (0) + \frac{1}{12} \cdot (-1) + \frac{1}{8} \cdot (1) \\
 &\quad + \frac{1}{12} \cdot (1) + \frac{1}{6} \cdot (0) + \frac{1}{4} \cdot (-1) \\
 &\quad + \frac{1}{24} \cdot (-1) + \frac{1}{12} \cdot (1) + \frac{1}{8} \cdot (0) \\
 &= -1/12.
 \end{aligned}$$

Lemma 3.48. Suppose in an $m \times n$ game with payoffs $u_{i,j}$ that Row plays $P = (p_1, \dots, p_m)$ and Column plays $Q = (q_1, \dots, q_n)$. Then the probability of the outcome (i,j) is $p_i q_j$, and the expected value of the payoff is computed by adding up the numbers $p_i q_j u_{i,j}$ for all values of i and j .

We can write this sum

$$\begin{aligned} E(P,Q) &= p_1q_1u_{1,1} + p_2q_1u_{2,1} + \cdots + p_mq_1u_{m,1} \\ &\quad + p_1q_2u_{1,2} + p_2q_2u_{2,2} + \cdots + p_mq_2u_{m,2} \\ &\quad \vdots \\ &\quad + p_1q_nu_{1,n} + p_2q_nu_{2,n} + \cdots + p_mq_nu_{m,n} \end{aligned}$$

which gets quite clunky to write out. It's often more convenient to notate things this way:

Lemma 3.49. *Suppose in an $m \times n$ game with payoffs $u_{i,j}$ that Row plays $P = (p_1, \dots, p_m)$ and Column plays $Q = (q_1, \dots, q_n)$. Then the expected value of the payoff can be computed either by*

$$\begin{aligned} E(P,Q) &= p_1E(P_1, Q) + \cdots + p_mE(P_m, Q) \\ E(P,Q) &= q_1E(P, Q_1) + \cdots + q_nE(P, Q_n). \end{aligned}$$

Proof. We can expand out either of these sums, to see that either one adds up all the $p_iq_ju_{i,j}$, which by lemma 3.48 is $E(P,Q)$. \square

3.3 Solving Zero-Sum Games

If a game has a saddle point, both players benefit from playing the saddle point, and the game will converge to that point stably. But if it doesn't, there's no single stable pair of strategies. No matter what cell the pair of players winds up in, at least one player can benefit from moving, which produces an infinite cycle of mind games.

The mathematician John von Neumann showed that each player in a two-person zero-sum game has an optimal strategy—for a specific sense of “optimal”—that they can stably settle on, as long as they're allowed to choose a mixed strategy.

Example 3.50. Suppose in a game of Rock Paper Scissors, Column plays the strategy $Q = (2/7, 4/7, 1/7)$. We want to find an optimal response for Row.

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0
	2/7	4/7	1/7

We can compute the expected value of the payoff for any pure strategy Row can select:

$$E(P_1, Q) = \frac{2}{7} \cdot (0) + \frac{4}{7} \cdot (-1) + \frac{1}{7} \cdot (1) = -3/7$$

$$E(P_2, Q) = \frac{2}{7} \cdot (1) + \frac{4}{7} \cdot (0) + \frac{1}{7} \cdot (-1) = 1/7$$

$$E(P_3, Q) = \frac{2}{7} \cdot (-1) + \frac{4}{7} \cdot (1) + \frac{1}{7} \cdot (0) = 2/7$$

(Note that the sum of the three expected payoffs is 0. Is that a surprise?)

So the best pure strategy response is to play Scissors, with an expected payoff of 2/7 per game. Can Row do better with a mixed strategy?

The answer turns out to be that if we know Column's mixed strategy, Row always has a pure strategy that will do as well as possible.

Lemma 3.51. *There is always a pure strategy among the best responses a player has to any pure or mixed strategy played by their opponent.*

Proof. The basic idea is that a mixed strategy is a combination of pure strategies, and so it can't ever be better than its best component.

Suppose P_k is Row's best pure strategy response to Column playing a mixed strategy Q . That means that

$$E(P_k, Q) \geq E(P_i, Q)$$

for any row i .

But we know from lemma 3.49 that if we have any mixed strategy $P = (p_1, \dots, p_m)$, then

$$\begin{aligned} E(P, Q) &= p_1 E(P_1, Q) + p_2 E(P_2, Q) + \dots + p_m E(P_m, Q) \\ &\leq p_1 E(P_k, Q) + p_2 E(P_k, Q) + \dots + p_m E(P_k, Q) \\ &= (p_1 + p_2 + \dots + p_m) E(P_k, Q) = E(P_k, Q). \end{aligned}$$

Thus no mixed strategy can be better than its best pure strategy component. □

Note that this is something people often find counterintuitive. If your opponent is going to play scissors half the time, it *feels* like you should try to play rock half the time. But in fact you'll do better by playing rock all the time, to exploit your opponent's overuse of scissors.

However, that doesn't mean we should only ever play a pure strategy. If Row adopts their optimal pure strategy, then Column can switch to an optimal pure strategy that beats them; we have returned to mind games. We instead want to find an equilibrium, where our opponent doesn't have an effective counter-play.

Definition 3.52. A mixed strategy outcome (P, Q) in a zero-sum game is called an *equilibrium* or *Nash equilibrium* if P is a best response to Q , and Q is a best response to P . We call these strategies P and Q *equilibrium strategies*.

We can see that a pair of mixed strategies (P, Q) is an equilibrium if and only if the following two statements are true:

$$\begin{aligned} E(P, Q) &\geq E(R, Q) && \text{for any Row mixed strategy } R \\ E(P, Q) &\leq E(P, S) && \text{for any Column mixed strategy } S \end{aligned}$$

A Nash equilibrium is a generalization of a saddle point.

Lemma 3.53. A pure strategy outcome (k, ℓ) is a saddle point if and only if the corresponding basic mixed strategy outcome (P_k, Q_ℓ) is a Nash equilibrium.

Example 3.54. Let's look back at Rock Paper Scissors one more time. We've seen that $Q = (2/7, 4/7, 1/7)$ isn't an equilibrium strategy, or really a good one; Row can exploit it by playing Scissors all the time.

But what if Column plays $Q' = (1/3, 1/3, 1/3)$? We can calculate Row's expected payoffs again:

$$\begin{aligned} E(P_1, Q') &= \frac{1}{3} \cdot (0) + \frac{1}{3} \cdot (-1) + \frac{1}{3} \cdot (1) = 0 \\ E(P_2, Q') &= \frac{1}{3} \cdot (1) + \frac{1}{3} \cdot (0) + \frac{1}{3} \cdot (-1) = 0 \\ E(P_3, Q') &= \frac{1}{3} \cdot (-1) + \frac{1}{3} \cdot (1) + \frac{1}{3} \cdot (0) = 0 \end{aligned}$$

Thus *regardless* of what strategy Row plays, their expected payoff is zero; Row no longer has an ability to affect the result of the game.

Definition 3.55. A mixed strategy is called a *neutralizing strategy* if the expected payoff is the same for every possible response by the opponent.

An outcome (P, Q) is a *neutralizing outcome* if both P and Q are neutralizing strategies.

Lemma 3.56. *A neutralizing outcome in a zero-sum game is a Nash equilibrium.*

Proof. Let (P, Q) be a neutralizing outcome. Since P is a neutralizing strategy for Row, every Column strategy is a best response to P , and thus in particular the strategy Q is a best response to P . Similarly, P is a best response to Q . \square

By looking for neutralizing outcomes, we can ultimately prove von Neumann's equilibrium theorem.

Theorem 3.57 (von Neumann's equilibrium theorem). *Every two-person zero-sum game has a Nash equilibrium.*

We'll prove a limited version of this soon, and we'll discuss the general version of it in section 3.5.

Definition 3.58. The *equilibrium method* for a zero-sum game is the method in which players choose one of their equilibrium strategies.

Example 3.59. In Rock Paper Scissors, the neutralizing strategy for Row is $P = (1/3, 1/3, 1/3)$, and the neutralizing strategy for Column is $Q = (1/3, 1/3, 1/3)$. That is, a neutralizing strategy involves choosing the three options randomly with equal probability.

By lemma 3.56, this is a Nash equilibrium.

3.3.1 Prudent mixed strategies

Definition 3.60. The *guarantee* of a mixed strategy is the expected value of the payoff when the opponent plays their best response to the mixed strategy.

By lemma 3.51 this is pretty straightforward to compute. We know that there is a pure strategy among the opponent's best responses, so we can just compute the expected value of each *pure* strategy against a given mixed strategy, and the worst result is the guarantee of that mixed strategy.

Example 3.61. In example 3.50, we examined Column's mixed strategy $Q = (2/7, 4/7, 1/7)$; we saw it has expected values of $E(\text{Rock}) = -3/7$, $E(\text{Paper}) = 1/7$, and $E(\text{Scissors}) = 2/7$. The worst of these for Column is $2/7$, so that's the guarantee.

If we consider Column's strategy $Q' = (1/3, 1/3, 1/3)$, we just saw that the expected value of any response is 0. Thus the guarantee of this strategy is 0.

Definition 3.62. The *prudent mixed strategy* for a player is the mixed strategy with the best guarantee. (This is the highest guarantee for Row, and the lowest for Column).

The *prudent mixed strategy method* is the method in which each player plays their prudent mixed strategy.

We'll show that every game has a prudent mixed strategy.

Example 3.63. The mixed strategies $P' = (1/3, 1/3, 1/3)$ and $Q' = (1/3, 1/3, 1/3)$ are Row's and Column's unique prudent mixed strategies in Rock Paper Scissors, and their guarantees are $\bar{r} = 0$ and $\bar{c} = 0$.

Why? consider any Row mixed strategy $P = (p_1, p_2, p_3)$ that isn't the same as P' . Then at least one probability p_k must be bigger than $1/3$. If Column plays the counterplay to the strategy with biggest probability, then they will win on average, and thus Row will lose on average; their guarantee, then, is negative. That's worse than the guarantee of zero.

Theorem 3.64 (von Neumann's Min-Max Theorem). *In a two-person zero-sum game, every Nash equilibrium (P, Q) is doubly prudent, with $\bar{r} = \bar{c}$. In particular, every such game has prudent mixed strategies for both players. Conversely, every doubly prudent mixed strategy outcome (P, Q) is a Nash equilibrium.*

Proof. If (P, Q) is a Nash equilibrium, then we know that $E(P, Q) \geq E(R, Q)$ for every Row strategy R , and $E(P, Q) \leq E(P, S)$ for every Column strategy S . Thus $E(P, Q)$ is the guarantee of both P and Q , which implies that $\bar{r} = \bar{c} = E(P, Q)$. Thus both P and Q are prudent.

Conversely, suppose (P, Q) is doubly prudent, and thus $E(P, Q) = \bar{r} = \bar{c}$. Since P is prudent, every payoff against P must be at least the optimal guarantee \bar{r} . Thus we must have $E(P, S) \geq E(P, Q)$ for every column strategy S . This means that Q is the (or a) best response to P .

Similarly, Q is prudent, so every payoff against Q is at most the optimal guarantee \bar{c} . Then we must have $E(R, Q) \leq E(P, Q)$ for every row strategy R . This means that that P is a best response to Q . Thus (P, Q) is a Nash equilibrium. \square

Example 3.65. Consider the following zero-sum game:

2	-1	-1
-1	2	-1
0	-3	2

It's in general hard to solve these, even for a relatively simple 3×3 game like this. (I had to code a small computer program to do it.) But we can check a solution if one is presented.

Take $P = (1/5, 1/2, 3/10)$ and $Q = (3/10, 3/10, 2/5)$. We claim these give us a Nash equilibrium.

Compute each of Column's expected value payoffs against Row's strategy P :

$$\begin{aligned} E(P, Q_1) &= \frac{1}{5} \cdot (2) + \frac{1}{2} \cdot (-1) + \frac{3}{10} \cdot (0) = -1/10 \\ E(P, Q_2) &= \frac{1}{5} \cdot (-1) + \frac{1}{2} \cdot (2) + \frac{3}{10} \cdot (-3) = -1/10 \\ E(P, Q_3) &= \frac{1}{5} \cdot (-1) + \frac{1}{2} \cdot (-1) + \frac{3}{10} \cdot (2) = -1/10. \end{aligned}$$

Thus P is a neutralizing strategy against Column, with a guarantee of $-1/10$.

Similarly, we can compute each of Row's expected value payoffs against Column's strategy Q :

$$\begin{aligned} E(P_1, Q) &= \frac{3}{10} \cdot (2) + \frac{3}{10} \cdot (-1) + \frac{2}{5} \cdot (-1) = -1/10 \\ E(P_2, Q) &= \frac{3}{10} \cdot (-1) + \frac{3}{10} \cdot (2) + \frac{2}{5} \cdot (-1) = -1/10 \\ E(P_3, Q) &= \frac{3}{10} \cdot (0) + \frac{3}{10} \cdot (-3) + \frac{2}{5} \cdot (2) = -1/10. \end{aligned}$$

Thus Column's strategy Q is a neutralizing strategy against Row, with a guarantee of $-1/10$. Therefore we can conclude this is a Nash equilibrium.

3.3.2 Solving 2 by 2 games

Theorem 3.66. *Any 2-by-2 zero-sum game has a Nash equilibrium.*

Proof. Consider the most abstract and general possible 2-player game:

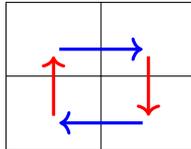
	1	2
1	a	b
2	c	d

In this game, player 1 and player 2 each have two options, labeled 1 and 2. If they both pick 1, the payoff is some number a . If Row picks 1 and Column picks 2, the payoff is some

number b , and so on. The point is that any 2-by-2 game will look like this; the only thing that changes between different games is the specific numbers we put in the grid.

If the game has a saddle point, then by lemma 3.53 it has a Nash equilibrium. So we want to think about the case where there's no saddle point.

Let's start by assuming $a \geq c$. This works "without loss of generality", because if $c > a$ we can flip everything right-to-left and then make the same argument. Then, since there are no saddle points, the flow diagram has to look like this:



In particular, note that we can't have $a = b$ because otherwise there would definitely be a saddle point somewhere. Thus $a > c$. We also can see from the diagram that $d > b$, that $b < a$, and that $c < d$.

Suppose Column adopts a mixed strategy $Q = (1 - q, q)$. Then the expected payoffs for Row's strategies are

$$E(P_1, Q) = a(1 - q) + bq = (b - a)q + a$$

$$E(P_2, Q) = c(1 - q) + dq = (d - c)q + c.$$

If we look at these as functions of q , each one is a line, with slopes $(b - a)$ and $(d - c)$. Since $c < d$ the second line slopes up, and since $b < a$ the first line slopes down; since $a > c$ and $b < d$, they'll have to cross over somewhere.

Algebraically, we can find this intersection: by solving for q , we get

$$\begin{aligned} (b - a)q + a &= (d - c)q + c \\ a - c &= aq - bq + dq - cq = ((a - c) + (d - b))q \\ q &= \frac{a - c}{(a - c) + (d - b)}. \end{aligned}$$

We can plug this back in to our original equations. Noting that

$$1 - q = 1 - \frac{a - c}{(a - c) + (d - b)} = \frac{(a - c) + (d - b)}{(a - c) + (d - b)} - \frac{a - c}{(a - c) + (d - b)} = \frac{d - b}{(a - c) + (d - b)},$$

we get

$$\begin{aligned}
 E(P_1, Q) &= a(1 - q) + bq \\
 &= a \frac{d - b}{(a - c) + (d - b)} + b \frac{a - c}{(a - c) + (d - b)} \\
 &= \frac{ad - ab + ab - bc}{(a - c) + (d - b)} \\
 &= \frac{ad - bc}{(a - c) + (d - b)} \\
 E(P_2, Q) &= c(1 - q) + dq \\
 &= c \frac{d - b}{(a - c) + (d - b)} + d \frac{a - c}{(a - c) + (d - b)} \\
 &= \frac{cd - cb + ad - cd}{(a - c) + (d - b)} \\
 &= \frac{ad - bc}{(a - c) + (d - b)}.
 \end{aligned}$$

That means that when q takes on this value $q = \frac{a - c}{(a - c) + (d - b)}$, both of Row's pure strategies have the same (expected) payoff, and thus *all* of Row's strategies have the same expected payoff. So $Q = (1 - q, q)$ is a neutralizing strategy for column.

We can run through the exact same argument for a Row strategy $R = (1 - p, p)$. We get the equations

$$\begin{aligned}
 p &= \frac{a - b}{(a - b) + (d - c)} \\
 E(P, Q_1) &= E(P, Q_2) = \frac{ad - bc}{(a - b) + (d - c)}
 \end{aligned}$$

and thus $P = (1 - p, p)$ is a neutralizing strategy for Row.

Thus (P, Q) is a Nash equilibrium by lemma 3.56. □

Remark 3.67. The number $ad - bc$ is sometimes called the *determinant* of the matrix, and winds up being important in many math contexts.

Corollary 3.68. *The solution to a 2-by-2 zero-sum game without saddle points is the pair of mixed strategies*

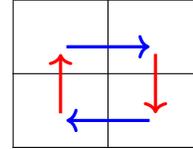
$$\begin{aligned}
 P &= \left(\frac{d - c}{(a - b) + (d - c)}, \frac{a - b}{(a - b) + (d - c)} \right) \\
 Q &= \left(\frac{d - b}{(a - c) + (d - b)}, \frac{a - c}{(a - c) + (d - b)} \right)
 \end{aligned}$$

for Row and Column respectively. The expected value of the game's payout is

$$v = \frac{ad - bc}{(a - b) + (d - c)}.$$

Example 3.69.

	A	P
A	60	10
P	0	40



Compute:

$$\begin{aligned}
 p &= \frac{a - b}{(a - b) + (d - c)} \\
 &= \frac{60 - 10}{60 - 10 + 40 - 0} = \frac{50}{90} = \frac{5}{9} \\
 1 - p &= \frac{9}{9} - \frac{5}{9} = \frac{4}{9} \\
 q &= \frac{a - c}{(a - c) + (d - b)} \\
 &= \frac{60 - 0}{60 - 0 + 40 - 10} = \frac{60}{90} = \frac{2}{3} \\
 1 - q &= 1 - \frac{2}{3} = \frac{1}{3} \\
 v &= \frac{ad - bc}{(a - b) + (d - c)} \\
 &= \frac{(60)(40) - (10)(0)}{60 - 10 + 40 - 0} = \frac{2400 - 0}{90} = \frac{240}{9} = \frac{80}{3} = 26\frac{2}{3}.
 \end{aligned}$$

So the optimal strategy is for Player 1 (the US) to play A $4/9$ of the time, and P $5/9$ of the time; Player 2 (Al-Qaeda) to play A $1/3$ of the time and P $2/3$ of the time. The expected payoff is a $26\frac{2}{3}\%$ chance of catching OBL.

3.4 Conflict and Cooperation

In the last section we talked about zero-sum games, where a gain to Row was a loss to Column. These are relatively simpler to analyze, but they leave out any possibility for cooperation—there's no win-win scenario, and also no scenario that represents a loss to both players. In this section we want to talk about *non-zero-sum games*, where each player can win or lose independently.

Definition 3.70. A *bimatrix* is a rectangular array in which each cell has an ordered pair of numbers.

Given a row i and a column j , we can take the cell in that row and column. We write $u_{i,j}$ for the first number in this cell, and $v_{i,j}$ for the second number.

Example 3.71.

0, 3	-1, 2	0, -2
1, 0	0, 1	0, 0
2, -1	-1, -6	0, 3

In this bimatrix, we see that $u_{1,2} = -1$ and $v_{1,2} = 2$.

Now we can define a *bimatrix game* just as we did matrix games in section 3.1. The first player, Row, picks a row; at the same time, the second player, Column, picks a column. We look at the cell of the grid corresponding to that row and column.

But this time, Row gets the first number in the grid in winnings, and Column gets the second number. It's possible for both to win, or both to lose. Or for Row to win a lot and Column to lose only a little. So Row wants to make $u_{i,j}$ as big as possible, but doesn't care about $v_{i,j}$; Column wants to make $v_{i,j}$ as big as possible, but doesn't care about $u_{i,j}$.

Example 3.72. Suppose we take the bimatrix from example 3.71, and Row plays row 3 while Column plays column 3. Then Row will win 0, and Column will win 3.

If Row stays on row 3, but Column instead plays column 2, then Row will lose 1, but Column will lose 6. No one is happy in that situation.

Remark 3.73. Any zero-sum game can be viewed as a bimatrix game where the two numbers in each cell are always opposite of each other. But most non-zero-sum games can't be represented with a single matrix.

3.4.1 Guarantees and Saddle Points

We can carry over a lot of our theory from our discussion of zero-sum games.

Definition 3.74. The *guarantee* of a strategy is the worst payoff a player can get by playing it. In the case of a bimatrix game, Row's guarantee for row k is the lowest value of $u_{k,j}$ in that row; Column's guarantee for column ℓ is the lowest value of $u_{i,\ell}$ in that column.

(Note that both players prefer larger numbers now! They're just looking at different numbers from each other.)

The *prudent strategy* is the strategy with the largest guarantee, and the *prudent strategy method* is the method that recommends playing the prudent strategy.

As before, we can draw a min-max diagram for a game.

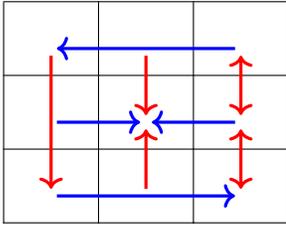
Example 3.75.	0, 3	-1, 2	0, -2	-1
	1, 0	0, 1	0, 0	0 ←
	2, -1	-1, -6	0, 3	-1
				-1 -6 -2
				↑

If both players play their prudent strategy, we wind up in row 2, column 1. Row gets 1 point and Column neither gains nor loses anything. (But both are doing better than their guarantee.)

Definition 3.76. The *best response* to an opponent's strategy is the strategy that gets the best outcome for the player, if the opponent plays that strategy. As before, we can track these with a *flow diagram*.

Example 3.77.

0, 3	-1, 2	0, -2
1, 0	0, 1	0, 0
2, -1	-1, -6	0, 3



In this case the horizontal arrows are looking for the largest *second* number in each row, while the vertical arrows are looking for the largest *first* number. Remember the Column only cares about second numbers, while Row only cares about first numbers.

Definition 3.78. A *saddle point* is an outcome in which each player's strategy is a best response to the opponent's strategy. Thus saddle points happen in cells where every arrow of the flow diagram points inward.

Example 3.79. Let's think about the logic of our bimatrix game. Suppose Column thinks that Row will select row 1. Then Column would want to select column 1, for the payoff of 3. (This isn't great for Row, though not awful.)

But if Column will select the first column, then Row would want to select row 3, to get a payoff of 2. (This gives Column a -1 .)

But then if Row is playing row 3, Column would want to play column 3, for the payoff of 3. This leaves Row getting nothing. In this case Row *could* switch to either other row; if Column picks column 3 then Row gets zero regardless. But they have no active incentive to switch. So $(3, 3)$ is a saddle point, with the payoffs $(0, 3)$.

But now imagine that Column is selecting column 2. Then Row would want to select row 2, to get the payoff of 0 (rather than -1). Now neither player has an incentive to switch; if Row moves they'll get -1 instead of 0, and if Column moves they'll get 0 instead of 1. Thus (2,2) is also a saddle point, with payoffs (0,1).

But something interesting happened here! In theorem 3.19 we saw that, in a zero-sum game, all saddle points had to have the same payoff. But here they don't. We have two saddle points; Row gets 0 in either one; but in one Column gets 1, while in the other Column gets 3.

3.4.2 Cooperation and Competition

Can list a game in order of preference. Look at the last game:

0, 3	-1, 2	0, -2
1, 0	0, 1	0, 0
2, -1	-1, -6	0, 3

Row's payoffs are 2, 1, 0, -1. Column's are 3, 2, 1, 0, -1, -2, -6. Can order these.

3rd, 1st	4th, 2nd	3rd, 6th
2nd, 4th	3rd, 3rd	3rd, 4th
1st, 5th	4th, 7th	3rd, 1st

Definition 3.80. A *coordination game* is a game where the two players' preference orders are exactly the same. A *strictly competitive game* is one where the preference orders are exactly opposite. A *mixed motive game* is any other game, which combines some coordination and some competition.

Proposition 3.81. Any zero-sum or constant-sum game is strictly competitive.

3.4.3 Some important games

Meet in New York

	Times Square	Grand Central	Empire State
Times Square	1, 1	0, 0	0, 0
Grand Central	0, 0	1, 1	0, 0
Empire State	0, 0	0, 0	1, 1

Coordination game, not even mixed. But hard!

Schelling improvement:

	Times Square	Grand Central	Empire State
Times Square	1, 1	0, 0	0, 0
Grand Central	0, 0	2, 2	0, 0
Empire State	0, 0	0, 0	1, 1

Battle of the Sexes

	Hockey	Ballet
Ballet	0, 0	10, 5
Hockey	5, 10	-5, -5

Mixed motives game.

Consider naive strategy? Look at prudent—same as naive.

2 saddle points, both better than either non-saddle. But which? Threaten bad move to get your way.

Payoff polygon

Pareto optimal, pareto principle

Chicken

	Don't Swerve	Swerve
Don't Swerve	-10, -10	10, -5
Swerve	-5, 10	0, 0

Mixed motives. Both prefer lower-right to upper-left, but other preferences on corners.

Examples: war, cuban missile crisis

	Keep bases	Withdraw
Invade	Nuclear War	US Victory
Blockade	Soviet Victory	Cold War

Prisoner's Dilemma

	Confess	Don't Confess
Confess	-10, -10	10, -5
Don't Confess	-5, 10	0,0

Mixed again

Stag Hunt

	Stag	Rabbit
Stag	10, 10	1, 5
Rabbit	5, 1	3,3

Almost a coordination game! But the corners are the issue again.

What is prudent strategy? It's a problem.

3.5 Nash Equilibria