Math 1007: Mathematics and Politics

The George Washington University Fall 2025

Jay Daigle

Contents

Mathematical Reasoning

Defining our terms

A big part of math is coming up with precise definitions so we're consistently talking about the same thing. This is really different from a lot of fields, where definitions can often be a lot fuzzier. Is the United States "a capitalist society"? Is it "a democracy"? Sort of but also it doesn't fit those idealized descriptions but also nothing really does. Categories in general are fuzzy.

In order to "do math" we need to write precise definitions so we don't have fuzzy categories; a thing either does, or does not, meet the definition. And then we can talk about what's true of *every* rule that falls into that category, because there are no edge cases that kinda sorta fall into the category but might be missing some important pieces.

But we also take the definition seriously. If we come up with a definition and it includes things we didn't think about originally, either we include those things in the category even if they "obviously shouldn't" count, or we have to start over and come up with a different definition. We'll do some of both in this course!

This is unlike the way we talk in the sciences, or the social sciences, or humanities, or every-day life. It is similar to the way we do *legal reasoning*, though. A criminal offense will have a specific list of elements, and you commit the crime if and only if you check off all those elements. A regulation will apply to everything meeting certain listed criteria, and sometimes that leads to weird results. Infamously, in California, a bee is a fish (for the purpose of certain environmental laws), because the law code defines "fish" to include "invertebrates" and bees are, in fact, invertebrates.

> Each of these statutes provides that covered species include "native species or subspecies of a bird, mammal, fish, amphibian, reptile, or plant[.]" This portion of the code, however, does not elaborate on what qualifies as a bird, mammal, fish, and so forth. Based only on the qualified species listed above, bees and other land-dwelling invertebrates would not receive protection under the law. The court looked elsewhere in the Fish and Game Code for definitions to help clarify whether bees may qualify for protection under CESA. Importantly, the section 45 of the code defines "fish" as "a wild fish, mollusk, crustacean, invertebrate, amphibian, or part, spawn, or ovum of any of those animals." (Emphasis added). According to the court, the term "invertebrate" under the definition of fish includes both aquatic and terrestrial invertebrates, such as bees. [Price(2022)]

Mathematical reasoning is similar to legal (and philosophical) reasoning, in that we are necessarily concerned with the implications of the precise rules we have set up—and if we don't like those implications, we have to change the rules themselves.

Some of the definitions we're about to lay out will seem obvious. Others will seem quite strange! But it's important to have them all laid out clearly so that we are following the same rules and discussing the same ideas and procedures for the rest of the course. When we want to give a formal definition, we will lay it out in a block like this:

Definition 0.1. A *function* is a rule that assigns exactly one output to every valid input. We call the set of valid inputs the *domain* of the function, and the set of possible outputs the *codomain* or sometimes the *range*. A function must be deterministic, in that given the same input it will always yield the same output.

We will try to use "normal" English in our definitions as much as possible. This will sometimes lead to awkward phrasing or high levels of wordiness as we try to give precise and unambiguous definitions without resorting to too much technical, mathematical jargon. For instance, instead of definition 0.1 we could have written:

Definition 0.2 (Unnecessarily technical definition)**.** Let A,B be sets. We define a function $f : A \to B$ to be a set of ordered pairs $\{(a,b) : a \in A, b \in B\}$ such that for each a in A there is exactly one pair whose first element is a. We call A the domain and B the codomain of f.

That technically conveys the same information, but is much harder to read if you're not used to it; we'll avoid that sort of thing as much as we can.

Proving Theorems

Another big chunk of mathematical reasoning involves making very precise statements, and proving that they are true. We will often set these off in a block of text with a label. We use:

- "Theorem" for a major, important result
- "Proposition" for a less important result that we still care about
- "Lemma" for annoying technical results we mainly want in order to prove something else we actually care about
- "Corollary" for something that follows immediately from something we've already proven.

There are two interesting things going on here, besides the terminology. One is the idea of "proof". When we call something a theorem (or proposition or lemma), we mean that the statement is always true under all possible circumstances. That means we can *never* prove a theorem simply by providing examples of it being true. However, we can very easily *disprove* a theorem by showing a single example where it is false.

For instance, consider the statement "All swans are white". This has the form of a theorem; it claims that something is always true. How could we go about proving it? I could show you a lot of white swans, but no matter how many I show you, that doesn't prove that *all* swans are white; maybe most swans are white, but some are not. On the other hand, if you show me one black swan, you have proven that statement is false.

But it's very important that we really do mean always. Consider the statement "No swans are red". That's very close to being true. If you go out and survey a bunch of swans, you're unlikely to find a red one. But probably some swan, at some point, has accidentally gotten dyed red, and certainly it's possible that a swan *could* be dyed red. So the statement is false.

So how can we possibly have actual theorems? Here's a statement that works: "No black swan is white". That has to be true, always, because *if the swan is white, then it isn't black.* So a theorem almost always has a set of conditions or "hypotheses"; we conclude that any situation satisfying those hypotheses must also satisfy those conclusions.

For instance, the deficit is once again becoming a significant political issue. There are various ways we could try to tackle the deficit. Math can't tell you whether the deficit will increase or decrease. Math can't tell you if the deficit should increase or decrease. But we can say things like "If you increase revenue and decrease spending, then the deficit will go down." The hypotheses necessarily imply the conclusion. And conversely, if the deficit is going up, one of those hypotheses must be failing—either we are decreasing revenue or increasing spending.

Later in the course we will prove proposition 1.49, which says: "Any social choice function that is anonymous and neutral must not be decisive". Don't worry about what those words mean; we'll define them eventually. But the point is that while some methods are "decisive", whatever that means, not all are; and we can show that any time a method is "anonymous" and "neutral", it cannot be "decisive".

So maybe the most useful way to think about theorems is that they show you how something can fail. If we want to have a white swan, we need to look at swans that aren't black. If we want to have a decisive voting method, we need one that either isn't anonymous, or isn't neutral.

Math often gets presented as a set of timeless truths. And that's not *entirely* false. But it works because the truths are all conditional. We can't say that something always has to happen; but we can show that there are two things you can't ever get at the same time. We can show that some tradeoffs will always have to exist.

1 Voting

1.1 Social Choice Functions

1.1.1 Ballots

We want to consider elections with more than 2 candidates. We call the set of candidates the *slate* and the set of voters the *electorate*. We'll generally call the candidates $A, B, C, \ldots$. (Again, instead of "candidates" these could represent bills or policies or pizza toppings, but we'll call them candidates to keep the terminology simple.)

For most of this section we'll assume each voter submits a *preference ballot*, in which they list the candidates in decreasing order of preference. So if there are four candidates $\{A, B, C, D\}$ a voter might submit a ballot that says they prefer $B > D > C > A$, which we might represent

B
D
C
A

We assume that voters are *rational*, in a specific way: we assume that each voter's preferences are *transitive*, meaning that if a voter prefers B to D and prefers D to C they must also prefer B to C. (It's not clear that people work this way in real life! But it's a convenient simplifying assumption for us to make now.)

We can represent the set of all ballots with a *profile*, similar to the profiles of section 1.2.1. But now these profiles need to have more rows, so we might get a profile like

A	B	A	C	A	B	C
B	C	B	A	B	C	A
C	A	C	B	C	A	B

Figure 1.1

In this profile, there are seven voters. We can see the first and third prefer A to B and B to C, while the second and sixth prefer B to C and C to A. In many cases we can summarize this in a *tabulated profile*

3	2	2
A	B	C
B	C	A
C	A	B

Definition 1.1. A *social choice function* for a slate of candidates is a function that takes in a voter profile, and outputs a non-empty subset of the slate.

We think of this subset as the list of "winners" of the election. We must have at least one winner, but we can have multiple winners (which you might interpret as a tie).

There is a tremendously large set of possible social choice functions; even with just three candidates and four voters, the number of possible social choice functions is a thousand-digit number. But most of those functions aren't very interesting.

We're going to start with the most obvious social choice function, which you're probably all familiar with. But then we'll look at a few more sophisticated approaches. How to choose among them?

We'll start by looking at a simple case: the case where there are only two candidates. In section 1.2 we'll think through the ways we could handle an election with only two outcomes, and see that there really is an optimal answer there.

In the process, we'll identify some properties we really want a social choice function to have. In section 1.2.4 we'll look at more of those properties, especially ones that only makes sense with more than two candidates. Then in section 1.3 we'll try to understand how each of the methods we've looked at measure up.

1.1.2 A collection of voting methods

The Plurality Method Let's start with the obvious answer: the method used most places in the United States.

Definition 1.2. A candidate who gets more votes than any other candidate is said to have a *plurality* of the votes.

In the *plurality method* we select as the winner the candidate who is ranked first by the largest number of voters. In the case that there is a tie for the most first-choice votes, we select all the candidates who tie for the most first-choice votes.

Let's apply this to the voting profile in figure 1.1. Candidate A gets 3 first-place votes, while B and C each get 2. So A is the plurality candidate and wins the election.

The plurality method is familiar and simple. It also doesn't use very much information; it only needs to know each voter's top preference (sometimes called a "vote-for-one" ballot.) In some ways, that's an advantage, since it makes voting and ballots much simpler. But it doesn't use much information about voter preferences, so it seems like we could probably produce better social choice methods somehow.

For instance, consider the following tabulated profile 1.2:

5	4	4	4	3
A	B	C	D	E
B	C	B	B	D
C	E	D	E	B
E	D	E	C	C
D	A	A	A	A

Figure 1.2

In this profile, A gets 5 first-choice votes, while B, C, D each get 4 and E gets 3. So A is the plurality winner. However, we see that all the voters who don't rank A first rank them *last*; A seems widely disliked, and probably a bad choice overall. In particular, B seems widely liked, with 4 first-place votes and *thirteen* second-place votes, out of twenty. Maybe we can develop a method that takes that other information into account?

Hare's Method/Instant Runoff Voting Let's start with a method that is currently used in a several jurisdictions. This method is currently used in Australia, and Papua New Guinea. It is also used for general elections in Alaska and Maine, and in for primary elections New York City, like the one that recently nominated Zohran Mamdani for Mayor.

Definition 1.3. *Hare's method* operates as follows. Unless all candidates have the same number of first-place votes, identify the candidate (or candidates) who have the fewest first-place votes, and eliminate them from consideration. Create a new profile where that candidate is removed, and each voter moves up their lower-ranked candidates by one place.

Repeat the process, either until only one candidate remains, or until all remaining candidates have the same number of first-place votes. The remaining candidates are the winners.

Remark 1.4. This method is frequently referred to as *Instant Runoff Voting*, *Single Transferable Vote*, or sometimes just *ranked choice voting*. The last name is common but unfortunate, since there are many voting methods that involve a ranked choice process; we'll see several in this course.

Let's apply this to the following profile:

5	4	4	4	3
B	C	A	D	E
C	A	B	A	A
E	B	E	B	B
D	E	D	E	D
A	D	C	C	C

Figure 1.3

We see that E has the fewest first-place votes, and so can be removed. That generates the profile

5	4	4	4	3
B	C	A	D	A
C	A	B	A	B
D	B	D	B	D
A	D	C	C	C

Now D and C are tied for the fewest first-place votes, with 4 each. We eliminate them both and get the profile

5	4	4	4	3
B	A	A	A	A
A	B	B	B	B

B has 5 first-place votes and A has 15, so B is removed and A is left as the winner.

But now let's go back and apply Hare's method to profile 1.2. Again, E has the fewest first-place votes, and is eliminated. That gives us the following profile:

5	4	4	4	3
A	B	C	D	D
B	C	B	B	B
C	D	D	C	C
D	A	A	A	A

Now A has 5 first-place votes, D has 7, while B and C each have 4. We eliminate B and C simultaneously, giving the profile

5	4	4	4	3
A	D	D	D	D
D	A	A	A	A

So this time D wins. And this happens even though a majority prefer B to D!

Coombs's Method There are other ways to follow up on the same basic idea. Here is another way to follow the instant runoff logic that Hare's method uses:

Definition 1.5. *Coombs's method* operates as follows. Unless all candidates have the same number of last-place votes, identify the candidate (or candidates) with the most last-place votes, and eliminate them. As in Hare's method, when a candidate is eliminated, remove them from the profile and let each voter move the candidates they ranked below the eliminated candidate up a spot.

Repeat the process, eliminating candidates, until either one candidate remains, or all remaining candidates have the same number of last-place votes. The candidate or candidates who remain at the end are the winners.

We can apply this method to profiles 1.2 and 1.3.

5	4	4	4	3
A	B	C	D	E
B	C	B	B	D
C	E	D	E	B
E	D	E	C	C
D	A	A	A	A

$\to^A$

5	4	4	4	3
B	B	C	D	E
C	C	B	B	D
E	E	D	E	B
D	D	E	C	C

$=$

9	4	4	3
B	C	D	E
C	B	B	D
E	D	E	B
D	E	C	C

$\to^D$

9	4	4	3
B	C	B	E
C	B	E	B
E	E	C	C

$\to^E$

9	4	4	3
B	C	B	B
C	B	C	C

$=$

16	4
B	C
C	B

$\to^C$

20
B

and so B wins profile 1.2 by Coombs's method, unlike with Hare's method.

For profile 1.3 we get

5	4	4	4	3
B	C	A	D	E
C	A	B	A	A
E	B	E	B	B
D	E	D	E	D
A	D	C	C	C

$\to^C$

5	4	4	4	3
B	A	A	D	E
E	B	B	A	A
D	E	E	B	B
A	D	D	E	D

$=$

5	8	4	3
B	A	D	E
E	B	A	A
D	E	B	B
A	D	E	D

$\to^D$

5	8	4	3
B	A	A	E
E	B	B	A
A	E	E	B

$=$

5	12	3
B	A	E
E	B	A
A	E	B

$\to^E$

5	12	3
B	A	A
A	B	B

$\to^B$

20
A

and so A wins.

Hare and Coombs have similar structures, but somewhat different incentives. Hare cares about the number of first-place votes, and so rewards candidates with a lot of strong support, even if they have a lot of strong detractors or haters. Coombs cares about the number of *last*-place votes, so it rewards candidates no one hates, even if no one really loves them either; it rewards middle-of-the-road moderation and widespread acceptability.

The Borda Count Method A different approach doesn't work by simulating runoff elections at all. Instead it tries to take into account all the varying levels of support a candidate has at once.

Definition 1.6. The *Borda count method* works as follows. If there are n candidates, give each candidate $n-1$ points for each voter who ranks them first; $n-2$ points for each voter who ranks them second; $n-3$ points for each candidate who ranks them third; and so on, until they get 1 point for each voter who ranks them second-to-last (and 0 points for each voter who ranks them last.)

Add up all the points; the candidate who gets the most points wins. If more than one candidate ties for the most points, all of them win.

Let's try applying this to profile 1.2. There are 5 candidates, so we give 4 points for a first-place vote. We can calculate that:

- A gets $5 \cdot 4 + 4 \cdot 0 + 4 \cdot 0 + 4 \cdot 0 + 3 \cdot 0 = 20$ points;
- B gets $5 \cdot 3 + 4 \cdot 4 + 4 \cdot 3 + 4 \cdot 3 + 3 \cdot 2 = 61$ points;
- C gets $5 \cdot 2 + 4 \cdot 3 + 4 \cdot 4 + 4 \cdot 1 + 3 \cdot 1 = 45$ points;
- D gets $5 \cdot 0 + 4 \cdot 1 + 4 \cdot 2 + 4 \cdot 4 + 3 \cdot 3 = 37$ points;
- E gets $5 \cdot 1 + 4 \cdot 2 + 4 \cdot 1 + 4 \cdot 2 + 3 \cdot 4 = 37$ points.

Thus B wins by the Borda count method.

The Borda method is pretty compelling, and uses all the information from the voter profile. But it does raise some issues. For instance, if we consider figure 1.3, the Borda counts are $A : 49$; $B : 54$; $C : 31$; $D : 28$; $E : 38$. So B is the Borda count winner. B is also the plurality winner, in fact. But we might notice that of the 20 voters, 15 prefer A to B; in a head-to-head matchup between the two, A would win. In fact we've seen this problem a lot; often one candidate wins despite most voters preferring a different one. We could try to build a method that takes those head-to-head matchups into account.

Copeland's Method

Definition 1.7. *Copeland's method* is the social choice function in which each candidate earns one point for every candidate they beat in a head-to-head matchup (using a simple majority method). A candidate earns half a point for every candidate they tie. The candidate with the most points at the end becomes the winner. If there's a tie for the most points, all those candidates are winners.

Again looking at profile 1.2:

5	4	4	4	3
A	B	C	D	E
B	C	B	B	D
C	E	D	E	B
E	D	E	C	C
D	A	A	A	A

We need to look at each pairwise competition. For instance, we see that 5 voters prefer A to B, but 15 prefer B to A, so B wins that matchup 15 to 5, and B gets one point there. We get the following table:

	AB	AC	AD	AE	BC	BD	BE	CD	CE	DE
Votes for first candidate	5	5	5	5	16	13	17	13	13	8
Votes for second candidate	15	15	15	15	4	7	3	7	7	12
Winner	B	C	D	E	B	B	B	C	C	E

We see that B gets 4 points, C gets 3 points, E gets 2 points, and D gets 1 point. So B wins by Copeland's method.

A better way to organize this is to make a *matrix*. We can make a table

	A	B	C	D	E	Total
A	-	0	0	0	0	0
B	1	-	1	1	1	4
C	1	0	-	1	1	3
D	1	0	0	-	0	1
E	1	0	0	1	-	2

In each cell, we put a 1 if the candidate whose row it is wins, a 0 if the candidate whose column it is wins, and a 1/2 if the match is a tie. (We put a - on the diagonal, since it doesn't make sense to ask of A beats A.) Then each candidate gets the sum of their row as a score.

Now let's apply this to profile 1.3:

5	4	4	4	3
B	C	A	D	E
C	A	B	A	A
E	B	E	B	B
D	E	D	E	D
A	D	C	C	C

We get

	A	B	C	D	E	Total
A	-	1	1	1	1	4
B	0	-	1	1	1	3
C	0	0	-	0	0	0
D	0	0	1	-	0	1
E	0	0	1	1	-	2

There are also some methods that are definitely *bad* ideas. These aren't serious suggestions for a democratic election; but they are by-the-definition social choice methods. These are often important as test cases—if you say something is true of "all" social choice methods, you should check in your head whether it still applies to all of these.

Definition 1.8. In the *dictatorship method*, one voter is the *dictator*. Their first-choice candidate is the unique winner.

In the *monarchy method*, one candidate is the *monarch*. That candidate is the unique winner regardless of how anyone votes.

In the *all-ties method*, every candidate is selected as a winner.

1.2 Two-Candidate Elections

For right now we want to talk about two-candidate elections. Everything we say will also apply to votes on whether or not to pass a bill, or whether to get pizza or Chinese food; the only thing that matters is that we have exactly two choices, A and B. But for brevity we'll usually refer to them as candidates.

Obviously many elections have more than two options. (We could get Thai!) but it's convenient to start out looking at the two-candidate case. This allows us to deal with some of the issues that come up in general elections without having to think of all of them at once. This is a common technique in mathematical reasoning: whatever we want to study, we often start by working on a much simpler version where some of the difficulty disappears.

Definition 1.9. A *social choice function* for two candidates is a function whose domain is the set of all possible preferences that voters could have, and whose codomain is a set with three options: A wins; B wins; or A and B tie.

It's tempting to require social choice functions to output a definite winner, and remove the possibility of ties. But as we'll see, that can wind up being impractical. We'll return to that idea later.

1.2.1 Some methods of voting

There is a fairly obvious social choice function to start with:

Definition 1.10. The *simple majority method* for a two-candidate election is the social choice function that selects as the winner the candidate who gets more than half of all the votes cast. If each candidate gets exactly half of the votes, then the result is a tie.

So if 6 people vote for A and 4 people vote for B, then A will win. If 5 vote for A and 5 vote for B, then it will be a tie.

Weighted voting methods But note we made an implicit assumption in that last paragraph: we only gave *counts* of votes. What needs to happen for that to work? It needs to not matter who votes, just what the total is. That's true for the simple majority method, but it's not true for every possible social choice function.

Definition 1.11. A *weighted voting method* works as follows. There are n voters, and each voter is assigned a positive number of votes; we will say that voter number i has w_i votes, or a *weight* of w_i.

Set $t = w_1 + w_2 + \cdots + w_n$ be the total number of votes. A candidate who gets more than $t/2$ votes is the winner. If no candidate gets more than $t/2$ votes, the result is a tie.

Example 1.12. Suppose we have ten voters, with the following vote preferences:

A	B	A	B	A	A	A	A	B	B

We call this a *profile* of the electorate. If we don't care about who gives which vote, we can summarize this in a *tabulated profile*:

6	4
A	B

Thus we see that A gets 6 votes and B gets 4 votes. In the simple majority method, A will win.

But now suppose we have a weighted voting method where the first four voters get 5 votes each, the next three get 1 vote each, and the last three get 3 votes each. Then we have $t = 5+5+5+5+1+1+1+3+3+3 = 32$. Adding up the votes we see that A gets $5+5+1+1+1+3 = 16$ votes and B gets $5+5+3+3 = 16$ votes. That election is a tie! But note we could never have figured that out from the tabulated profile.

Poll Question 1.2.1. When does weighted voting make sense?

These weighted voting methods give a preference toward some voters over others. In the limit we can make this preference absolute:

Definition 1.13. In the *dictatorship method*, one of the voters is the *dictator*. Whoever the dictator prefers is the winner.

Example 1.14. Suppose we have a weighted voting method with ten voters, but the first voter gets twelve votes and each other voter gets one vote. Then the preference of the first voter will always win the election.

Obviously the dictatorship method is undemocratic. It is, however, mathematically a social choice function. And it does actually get used in some elections.

Example 1.15. The company Meta has two classes of stock: Class A shares get one vote per share, and Class B shares get 10 votes per share. Founder Mark Zuckerbert holds the overwhelming majority of the Class B stock, which means that while he only owns about 13% of the company, he has more than 50% of the voting power. Facebook shareholder elections are effectively run by the dictatorship method.

Supermajority Methods There's a different way we can tweak voting systems: we can have a bias toward ties.

Definition 1.16. Let p be a number such that $1/2 < p \leq 1$. The *supermajority method* with parameter p selects as the winner the candidate who gets a fraction p or more of all the votes. If there are t votes in total, a candidate must get at least $p \cdot t$ votes to win. If no candidate gets $p \cdot t$ votes, the result is a tie.

Example 1.17. A common value for p is $2/3$. If we look at the same profile above, we have ten voters and thus need $2/3 \cdot 10 = 20/3 \approx 6.66$ votes to win. A gets six votes, and B gets four votes, so the result is a tie.

The number q of votes a candidate needs to win is sometimes called the *quota* and supermajority methods are sometimes called *quota methods*. It's tempting to write that $q = pt$ but that isn't quite true, since no candidate can actually get $20/3$ votes. Instead, the quota is $q = \lceil pt \rceil$, which we read as "the ceiling of the number pt" and means the smallest integer that is greater than or equal to pt.

Poll Question 1.2.2. When would this social choice function make sense?

If $p = 1/2$, this is just the simple majority method. If $p = 1$, this is the *unanimity* method or the *consensus* method, where everyone needs to agree.

Poll Question 1.2.3. Why don't we allow $p < 1/2$?

Status quo methods A weighted voting method favors some voters; a supermajority method favors ties. But sometimes we want to favor some specific result.

Definition 1.18. Start with some social choice function (such as the simple majority method or the supermajority method), which we will call the "base method". A *status quo method* designates one of the two candidates as the status quo, and the other as the challenger. If either candidate wins under the base method, then that candidate also wins under the status quo method. If there is a tie under the base method, then the status quo method names the status quo candidate as a winner.

Example 1.19. Suppose as in example 1.17 we have a supermajority method with $p = 2/3$, and we designate candidate B is the status quo. We have the same voter profile, so A gets 6 votes and B gets 4 votes. The supermajority method would declare this a tie, so candidate B wins.

Poll Question 1.2.4. When does this method get used regularly? When is it a good idea?

Like with the weighted voting method, this can also be taken to an extreme:

Definition 1.20. In the *monarchy method*, one of the candidates is a *monarch*. That candidate wins regardless of how anybody votes.

Mathematically this is a *constant function*: all inputs have the same output.

Note that a monarch is different from a dictator. A dictator doesn't necessarily win the election; they just select who wins the election. A monarch doesn't get any say; they win the election whether they want to or not.

Poll Question 1.2.5. When does a monarchy method make sense?

Bloc Voting

Definition 1.21. A *block voting method* partitions the electorate into n blocs, so that every voter is in exactly one bloc. Then each bloc i is assigned a positive number w_i of votes. Each bloc conducts an election using the simple majority method, possibly with some method for resolving ties. Then the bloc casts all of its w_i votes in the main election for the candidate that won that simple majority election. The winner is the candidate who receives the most votes in the main election, with a tie if both candidates receive the same number of votes.

Example 1.22. Suppose we have the following profile, split into five blocs which each gets one vote:

A	B	A	B	A	B	A	B	B	A	A	A	B	A	B

A wins the first bloc and the fourth bloc, for 2 votes. B wins the second, third, and fifth blocs, for 3 votes. So B wins the election.

If we ignore the blocs, the tabulated voting profile here is:

8	7
A	B

Is that enough information to recover the winner of the bloc election? Who would win this election by the simple majority method?

Poll Question 1.2.6. When would this voting method make sense? Where does it get used in real life?

Testing the limits There are many, many social choice functions, and we can't possibly describe them all. But when we're analyzing social choice functions, it can be really helpful to have a list of options to consider. Especially when we're trying to prove something "always" or "never" happens, it can be useful to consider some very silly social choice functions to check our intuitions.

Two such silly methods are the dictatorship method of definition 1.13 and the monarchy method of definition 1.20. But here are two more to consider:

Definition 1.23. In the *all-ties method*, the election is a tie, no matter how the electorate votes.

Like the monarchy method, this is a constant function.

Definition 1.24. In the *parity method*, if just one candidate gets an even number of votes, then that candidate wins. If both candidates get an odd number of votes, or both candidates get an even number of votes, then the result is a tie.

Example 1.25. Suppose we have the tabulated profile

7	4
A	B

Then B has an even number of votes and A has an odd number of votes, so B wins. If one voter who currently prefers A switches their vote to B, we would get the tabulated profile

6	5
A	B

Now A has an even number of votes so A wins.

1.2.2 Voting method criteria

We've seen a bunch of social choice methods. We haven't said which ones are better than others, and on some level this isn't a mathematical question. We can't "prove" that weighted voting is better or worse than bloc voting, or anything like that. It depends on what you want, and what you mean by "better"!

What math can do is help you figure that out—what are your goals, and how well each method does at achieving your goals. So first we have to come up with some things we care about, and write them down as precisely as we've specified the social choice functions.

Poll Question 1.2.7. What features do you think a good voting system will have?

Followup: what kind of election were you thinking about there?

One way of thinking about our attempt to develop criteria is that there are some social choice methods that we know we don't like. For each of those, we can figure out what we don't like about it, and generate some criterion that expresses that judgment.

Definition 1.26. A method satisfies the *anonymity condition* (or *is anonymous*) if it treats all voters equally.

Another way of expressing this is that an anonymous method will always give the same result if the voters exchange ballots among themselves.

It's clear that the dictatorship method, and weighted voting methods are not anonymous. The simple majority and supermajority methods are anonymous. The monarchy, parity, and all-ties methods are also anonymous.

When we have a result we have established, we will state it as a *proposition*. Major results will be called *theorems*, and minor technical results that are useful mainly to prove other results are *lemmas*.

Proposition 1.27. *A method is anonymous if and only if its outcomes depend only on the tabulated profile.*

Proof. The phrase "if and only if" means that we have to prove two separate things. We need to show that if the method is anonymous, then it depends only on the tabulated profile; and we need to show that if a method depends only on the tabulated profile, then it is anonymous.

First let's suppose we have an anonymous method. We want to show that if two different voter profiles give the same tabulated profile, they will have the same result. Since the two profiles give the same tabulated profile, the same number of voters prefer candidates A and B respectively in each profile. That means we can start with the first profile and exchange ballots around until we have reached the second profile. Since the method is anonymous, this can't change the result; so the outcome depends only on the tabulated profile.

Conversely, suppose we have a social choice method whose outcome depends only on the tabulated profile. If we have two profiles that differ only by exchanging ballots between voters, they will produce the same tabulated profile and thus produce the same result. Therefore, the method must be anonymous.

□

Definition 1.28. A method satisfies the *neutrality criterion* or *is neutral* if it treats both candidates equally.

Status quo methods and monarchy are not neutral. Majority and supermajority methods are neutral, as are weighted voting methods and dictatorships. The parity and all-ties methods are neutral.

So far simple majority and supermajority satisfy both our criteria, as do parity and all-ties. But parity, in particular, is obviously absurd. We should come up with a criterion that explains why parity is a bad voting method. And there is in fact a criterion, which feels so obvious that you might not think to state it if you're not confronted with something as silly as parity. (But we'll see that it can be trickier than it sounds in multi-candidate elections!)

Definition 1.29. A method satisfies the *monotonicity criterion* or *is monotone* if a candidate is never hurt by getting more votes.

That is: suppose the votes are cast and the method selects one candidate as the winner. Suppose the method is applied again after one or more voters change their votes from the losing candidate to the winning candidate. The candidate who was the winner before the change must remain the winner after the change.

Monotonicity is necessary if we want to avoid strategic voting: in a monotone social choice method, it's always reasonable to vote for the candidate you actually prefer. In a two-candidate election any reasonable method will be monotone, but not every method will be.

Proposition 1.30. *The parity method violates monotonicity.*

Proof. This proposition only requires us to present a *counterexample*. To say a method *is* monotone is to say that a certain thing always happens. To show it is not monotone, we don't need to show that the thing never happens; we only need to show that it doesn't always happen. (Note that "it doesn't always rain" is very different from "it never rains"!)

Consider the profile

A	B	A	B	A	A	A	A	B

In this profile, A gets 6 votes and B gets 3 votes, so A wins.

Now suppose the second voter changes their mind and starts to prefer A. We get the following profile:

A	**A**	A	B	A	A	A	A	B

Now A gets 7 votes and B gets 2. By the parity method, B wins this election.

Thus a voter changing their vote from B to A causes B to win, and the parity method violates monotonicity.

□

The all-ties method also seems bad. What's wrong with it? Well, it always produces ties. A voting method should, ideally, give us a result.

Definition 1.31. A method satisfies the *decisiveness criterion* or is *decisive* if it always chooses a winner, that is, never produces a tie.

Poll Question 1.2.8. Which of the methods we've discussed so far are decisive?

This criterion is ideal, but it's actually too strong to be useful. Sometimes a tie is the only really reasonable output, like if both candidates get exactly the same number of votes. (We'll make this claim precise in section 1.2.3). So we want a slightly weaker criterion that isn't too strong but still rules out nonsense like the all-ties method.

Definition 1.32. A method satisfies the *near decisiveness criterion* or is *nearly decisive* if ties can only occur when both candidates receive the same number of votes.

These two criteria aren't exclusive. Any decisive method is also nearly decisive. Note that this isn't normally how we use English; but it is definitely the way that we wrote these definitions. However, many methods are nearly decisive, but not decisive.

1.2.3 May's Theorem

Proposition 1.33. *The simple majority method is nearly decisive.*

Proof. Suppose the number t of voters is odd. No candidate can receive exactly half the votes, since $t/2$ is not an integer. So one must receive more than half; they win a majority and win the election.

Now suppose t is even. If both candidates get $t/2$ votes, they receive the same number of votes, and tie. If not, one gets more than $t/2$ and thus has a majority and wins. So ties only occur when each candidate gets exactly $t/2$ votes. □

So far we haven't found any method that checks all our boxes. The simple majority method checks most of them, but it isn't decisive. However, Kenneth May in 1952 proved that if we want to check off all our criteria, we're asking for too much.

Theorem 1.34 (May's Theorem)**.** *In an election with two candidates, the only voting method that is anonymous, neutral, monotone, and nearly decisive is the simple majority method.*

Proof. Suppose we have an anonymous, neutral, monotone, nearly decisive social choice function for two candidates. Since it is anonymous, we only need to consider tabulated profiles by proposition 3.14. Then let a be the number of voters who support candidate A, and b the number of voters who support candidate B; set $t = a + b$ to be the total number of voters. We want to show that the method we are imagining must be the simple majority method.

First consider the case where t is even. If $a = b = t/2$ then each candidate gets the same number of votes, so by neutrality the result must be a tie, which is the result that the simple majority method would give.

Suppose candidate A has a majority, that is, $a > t/2$. We want to show that for *any* method that is neutral, monotone, and nearly decisive, A must win. By near decisiveness, the result cannot be a tie, so either A wins or B wins.

But B can't win this election. We see that B gets $b = t - a$ votes, and since $a > t/2$ then $b < t/2$. If B wins with $b < t/2$ votes, then by monotonicity, B must win with $t/2$ votes. But we saw that if B gets $t/2$ votes then the election is a tie.

So if A gets a majority, we've seen the election can't be a tie and B can't win; thus A wins. By neutrality, the same is true of B; so in either case, a candidate with a simple majority of votes must win.

Now consider the case where t is odd. In this case, neither candidate can get $t/2$ votes, since it's not an integer; by near decisiveness, some candidate must win.

Suppose A gets $a > t/2$ votes. By near decisiveness, this can't be a tie. We claim that B cannot be the winner. We see that B gets $b = t - a < t/2$ votes, and necessarily $b < a$. If B wins with $b < a$ votes, then by monotonicity B would also have to win with a votes. But by neutrality, that means that A would also win with a votes.

So if A gets $a > t/2$ votes, then the election isn't a tie and B doesn't win; so A must win the election. By neutrality, if B gets more than $t/2$ votes, then B must win as well. So a candidate with a simple majority of votes must always win. □

The important thing about this theorem is it doesn't just prove that we haven't come up with a better method than simple majority. We've shown that any method that satisfies these criteria has to be the simple majority method; there's no possible way to be "more clever" and come up with a better option. In particular, we have the following *impossibility result*:

Corollary 1.35. *It is impossible for a voting system with two candidates to be anonymous, neutral, monotone, and decisive.*

Proof. If a method is decisive, it must be nearly decisive. If a method is anonymous, neutral, monotone, and nearly decisive, it must be the simple majority method by Theorem 1.34. So any method fitting these criteria must be the simple majority method; but the simple majority method is not decisive, so no method fits all these criteria. □

A similar argument to the proof of May's theorem can give us the somewhat more general result below:

Theorem 1.36. *In an election with two candidates, a voting method that is anonymous, neutral, and monotone must be the simple majority method, a supermajority method, or the all-ties method.*

We're not going to prove this right now but I encourage you to think about why that should be true, and how you'd prove it.

1.2.4 Voting System Criteria

Definition 1.37. *Unanimity criterion* or *unanimous* if whenever all voters place the same candidate at the top of their preference orders, that candidate is the unique winner.

Desirable, too strong.

Definition 1.38. A method is *decisive* if it always selects a unique winner.

Also too strong!

Definition 1.39. A method satisfies the *majority criterion* if, whenever a candidate receives a majority of the first-place votes, that candidate must be the unique winner.

Definition 1.40. A method is *anonymous* if the outcome is unchanged whenever two voters exchange their ballots.

Lemma 1.41. *A social choice function is anonymous if and only if it depends only on the tabulated profile.*

Proof. See proposition 3.14. □

Definition 1.42. A method is *neutral* if it treats all candidates the same, in the following sense: suppose we have some profile that names A to be a winner. Now suppose there is another candidate B, and all voters exactly swap their preferences for A and B. In the new profile, B should be a winner.

Also should be the case that voters liking you more should never hurt you.

Definition 1.43. A method is *monotone* if: suppose there is a profile in which candidate A wins, but some voter puts another candidate B immediately ahead of A. If that voter moves A up one place to be ahead of B, then the method must declare A to be a winner in the new profile as well.

It follows that if A moves up any number of places for any number of voters, they still must win. This is extremely desirable because it rules out "tactical" voting: if you like A more than B you should always rank A over B. Ranking A over B can never cause A to lose.

Example 1.44.

A	C	B
B	A	C
C	D	D
D	B	A

Who *shouldn't* win?

Definition 1.45. A method is *Pareto* or satisfies the *Pareto criterion* if whenever every voter prefers a candidate A to another candidate B, then the method does not select B as a winner.

Named after Italian economist Vilfredo Pareto. A choice is "Pareto optimal" if there's no other choice that's better for everyone. This guarantees that a winner is Pareto optimal. (Note there can be many Pareto optimal outcomes.)

Does not mean that if everyone prefers A to B then A must be a winner! Maybe everyone prefers C to both.

7	6
D	C
C	D
A	A
B	B

Definition 1.46. A candidate is a *Condorcet candidate* if they beat every other candidate in a head-to-head. They are an *anti-Condorcet candidate* if they lose to every other candidate in a head-to-head.

Not always a Condorcet candidate. Famously

A	B	C
B	C	A
C	A	B

A beats B, B beats C, but C beats A.

People sometimes call Copeland's method the Condorcet method, but that's not quite right. And people sometimes say that the Condorcet method picks the Condorcet winner, but that's not actually a method—doesn't always give an answer!

Definition 1.47. A method satisfies the *Condorcet criterion* if whenever there's a Condorcet candidate, they're the unique winner.

A method satisfies the *anti-Condorcet criterion* if whenever there's an anti-Condorcet candidate, they don't win.

Definition 1.48. A method is *independent* if:

Suppose there are two profiles where no voter changes their mind about whether candidate A is preferred to candidate B: if a voter ranks A above B in the first profile, they must also rank A above B in the second profile. Suppose that in the first profile, the method makes A a winner but not B. Then the method must not choose B as a winner for the second profile.

The idea here is that to decide if A beats B, it shouldn't matter what voters think about another candidate C. So if we consider the two profiles:

A	C	B	C	C	A	B
B	A	C	B	B	B	A
C	B	A	A	A	C	C

A	C	B	C	C	A	B
B	A	C	B	B	B	C
C	B	A	A	A	C	A

The only difference between the two is that in profile 1, the last voter prefers A to C, while in profile 2 they prefer C to A; preferences between A and B are unchanged. If A wins and B loses in the first profile, that shouldn't change in the second!

Unfortunately, that's much harder than it sounds.

Poll Question 1.2.9. What methods have we looked at violate this?

Proposition 1.49. *Any social choice function that satisfies anonymity and neutrality must violate decisiveness.*

Proof. Since it's anonymous, we only have to look at tabulated profiles. Suppose we have $2n$ voters.

n	n
A	B
B	A

n	n
B	A
A	B

By neutrality, if A wins in the first B must win in the second, and vice versa. So both have to win.

□

Proposition 1.50 (Taylor)**.** *No social choice function involving at least three candidates satisfies both the independence criterion and the Condorcet criterion.*

Proof. Suppose we have an independent Condorcet method. Suppose there are three candidates and three voters, and we have the profile

A	C	B
B	A	C
C	B	A

We see that A can't be a winner. Because if we look at

A	C	C
B	A	B
C	B	A

C is the Condorcet candidate, and so must be the unique winner, and so A is not a winner. But since this doesn't change the relative positions of A and C, by independence A can't win in the first profile either.

We can make the exact same argument to show that B and C also can't be named winners in the original profile. But every profile must have a winner. So no method can be both Condorcet and independent.

□

Proposition 1.51. *If a method is Condorcet then it satisfies the majority criterion.*

Proof. Suppose A has a majority of first-place votes. Then they will win any head-to-head matchup, and thus A is the Condorcet candidate. So any method that satisfies the Condorcet criterion will cause A to win, also satisfying the majority criterion. □

In this sense we can say the Condorcet criterion is stronger than the majority criterion.

1.3 Evaluating Social Choice Functions

To summarize, we have discussed the following social choice functions:

- Plurality
- dictatorship
- monarchy
- all ties
- Borda count
- Hare's method
- Coombs's Method
- Copeland's Method
- Sequential agenda
- Positional methods

We also have several criteria we looked at:

- unanimous
- decisive
- majoritarian
- anonymous
- neutral
- monotone
- Pareto
- condorcet
- anti-condorcet
- independent

We want to think about how each of these criteria apply to each of the social choice functions we've studied.

1.3.1 Plurality method

Proposition 1.52. *The plurality method is majoritarian, monotone, and Pareto, but not Condorcet, anti-Condorcet, or independent.*

Proof. The majority is always a plurality, so the plurality method is majoritarian.

It's monotone because raising a candidate on some preference lists can't reduce their number of first-place votes, or increase any other candidate's. So if a candidate wins before the raising they will still win after.

It's Pareto because if A is ahead of B on every preference list, then B will have no first-place votes at all and thus will not win.

But consider the profile In this profile, B wins the plurality vote. But A beats B in a

2	3	2
A	B	C
C	A	A
B	C	B

Figure 1.4

head-to-head (four to three), A beats C head-to-head (five to four), and C beats B head to head (four to three). So A is the Condorcet candidate, and does not win; B is the anti-Condorcet candidate, and yet wins. So plurality fails both Condorcet and anti-Condorcet.

Now consider the following two profiles, "before" and "after".

A	A	A	A	B	B	B		A	A	C	C	B	B	B
B	B	C	C	A	A	A	→	B	B	A	A	A	A	A
C	C	B	B	C	C	C		C	C	B	B	C	C	C

In the first profile, A wins, with four first-place votes. In the second profile, two voters have changed their minds and now rank C above A; this causes B to win.

(This is exactly the sort of "strategic voting" we see in a lot of primary elections, but also in some general elections. Imagine that A is the Democrat, B is the Republican, and C is a far-left candidate. People will say that voting for C is "throwing away your vote" and just giving the election to B.)

□

1.3.2 Antiplurality method

Proposition 1.53. *The antiplurality method is monotone, but not majoritarian, Condorcet, anti-Condorcet, Pareto, or independent.*

Proof. It's monotone because raising a candidate in a preference list can never increase their number of last-place votes, or decrease the last-place votes of any other candidate. If A wins before the change, they will also win after.

Consider the profile In this profile, B has a majority first-place votes, but also a plurality

C	C	B	B	B
A	A	C	A	A
B	B	A	C	C

Figure 1.5

of last-place votes, and so loses under anti-plurality and so the method is not majoritarian.

Similarly, we see that B is the Condorcet candidate, and A is the anti-Condorcet candidate. But A wins and B loses. That means this method is neither Condorcet nor anti-Condorcet.

We can also construct a profile where A is rated higher than B by all voters, but B still wins. That's a little bit weird, because it means A can't have any last-place votes; but B

can tie with A as a winner by not having any last-place votes either. So this method isn't Pareto.

A	A	A
B	B	B
C	C	C

Finally, we want to think about independence. Consider the following two profiles:

C	A	A	B	B
A	B	B	A	A
B	C	C	C	C

$\rightarrow$

C	A	A	B	B
A	B	B	C	C
B	C	C	A	A

Figure 1.6

In the first profile, A wins the anti-plurality election, and B loses (along with C). In the second profile, we have swapped some voters between C and A, but keep the same relative positions of A and B; now B wins the anti-plurality election, and A And C both lose.

□

1.3.3 Borda count

Proposition 1.54. *The Borda count method is monotone, anti-Condorcet, and Pareto, but not majoritarian, Condorcet, or independent.*

Proof. Monotone because raising A on some preference lists will never decrease A's score, or increase anyone else's. So if A wins before the change, they will also win after.

It's Pareto because if A is ranked above B by every voter, then A will definitely get a higher score than B, and so B will not win.

The anti-Condorcet argument is kind of tricky. (Note this is the first time we're showing something *is* anti-Condorcet, so this will be a new type of argument.) Suppose that A is the anti-Condorcet candidate; we need to show that it's not possible for them to win. That means it's not enough to just give an example; we need to give a totally universal argument.

Instead, we need to make a tricky argument with counting. Let's suppose there are n candidates and m voters. Then every voter gives out $\frac{n(n-1)}{2}$ points, and the total number of points is $\frac{mn(n-1)}{2}$. That means the *average* number of points each candidate can get is $\frac{m(n-1)}{2}$ (the total divided by the number of candidates).

The most total points any candidate can get is $m(n-1)$, since there are m voters and a first-place vote gives $n-1$ points. We want to show that the anti-Condorcet candidate A gets less than half the maximum number of points.

Another way of thinking about the Borda count is that A gets one point for each pair (X,v) of (candidate, voter) such that the voter v ranks candidate X below A. That is, A gets one point for each time a voter ranks a candidate below them. But since A is the anti-Condorcet candidate, they are ranked below each candidate more often than they are ranked above that candidate—that's what it means to lose each round head-to-head. So A gets less than half the maximum number of possible points: less than $\frac{m(n-1)}{2}$. This is precisely the average number of points a candidate gets.

Since A gets less than the average, someone must get more than average. Thus someone gets more points than A, and A cannot win.

Now we want to show that the Borda count isn't majoritarian, Condorcet, or independent. For these we can simply show examples again.

Consider the profile

A	A	A	B	B
B	B	B	C	C
C	C	C	A	A

In this case, A is the Condorcet candidate, with a majority of the first-place votes. But A only gets 6 Borda points, while B gets 7, so B wins the Borda count election.

The profiles in figure 1.6 also show that the Borda count isn't independent. In the first profile A gets 7 points, B gets 6, and C gets 2, so A wins. In the second profile, we've only changed the relative positions of A and C; but now A gets 5 points, B still gets 6, and C gets 4. Now B wins.

□

1.3.4 Hare's method

One of the most popular methods is Hare's method. But it is actually surprisingly weak by our criteria, and in particular it's the first method we're going to talk about that, surprisingly, isn't monotone.

Proposition 1.55. *Hare's method is majoritarian and Pareto, but not monotone, Condorcet, anti-Condorcet, or independent.*

Proof. A candidate at the top of a majority of preference lists will always in in Hare's method, because they will never have less first-place votes than anyone else and so will never be eliminated.

Conversely, if a character is not at the top of any preference lists, they will always be eliminated in the first round. If everyone prefers A to B then B will not appear at the top of any preference lists and so will be immediately resolved.

The following pair of before-and-after preference lists shows that Hare's method violates monotonicity:

6	5	4	2
A	C	B	B
B	A	C	A
C	B	A	C

→

6	5	4	2
A	C	B	A
B	A	C	B
C	B	A	C

Figure 1.7

In the "before" profile, Hare's method will first eliminate C, and then in a head-to-head A will defeat B with eleven votes to six. In the "after" profile, two voters have switched their votes to prefer A to B. this means that in the first round, B is eliminated, and then in the second round C defeats A nine to eight.

To see that Hare violates Condorcet and anti-Condorcet, we can look back to figure 1.4:

2	3	2
A	B	C
C	A	A
B	C	B

In the first round, A and C are eliminated, leaving B as the unique winner. But A is the Condorcet winner and B is the Condorcet loser.

Finally, to consider independence, consider the following tabulated profiles:

In the "before" figure, C is eliminated in the first round, then B is eliminated in the second round, leaving A as the unique winner. In the "after" profile, two voters have switched their preferences on A and C, but this leaves the relative rankings of A and B unchanged. But now, A loses in the first round, and in the second round C is eliminated leaving B as the unique winner. This violates independence.

2	2	1
B	A	A
A	C	B
C	B	C

→

2	2	1
B	C	A
A	A	B
C	B	C

□

1.3.5 Coombs's method

Proposition 1.56. *Coombs's method is Pareto but not majoritarian, monotone, Condorcet, anti-Condorcet, or independence.*

Proof. If candidate A is ahead of candidate B on every preference list, then A will never have any last-place votes until B is eliminated. Therefore A will survive each round, with no last place votes, until B is eliminated; the election cannot come to an end until that happens. So B must lose, and Coombs's method is Pareto.

We saw in figure 1.5 that a candidate can have a majority and still lose an antiplurality election:

C	C	B	B	B
A	A	C	A	A
B	B	A	C	C

Here B is the majority winner, but is eliminated in the first round of Coombs's method and thus does not win. This method violates the majority criterion, and thus also the Condorcet criterion.

It *seems* like Coombs's method should satisfy the anti-Condorcet criterion, since even if they survive to the final round, they will lose the final head-to-head matchup (since they lose any head-to-head matchup). But this is not true! If we look at this same profile 1.5 again, we can see that A is the anti-Condorcet candidate losing to both B and C. But A wins the election, precisely because there is no "final" head-to-head: B and C are both eliminated at the same time, leaving A the unique winner.

The remaining two properties we leave as an exercise for the student. □

Exercise 1.57. *You should prove that Coombs's method isn't monotone or independent. Each of these arguments requires generating an example pair of profiles. You may wish to* *use the proof of proposition 1.55 as a hint or guideline as you construct these.*

1.3.6 Copeland's method

The two elimination methods were very similar; the positional methods are very similar. Copeland's method is structurally quite different and therefore produces a very different set of results. In particular, this system is constructed almost specifically to satisfy the Condorcet property.

Proposition 1.58. *Copeland's method is majoritarian, Condorcet, anti-Condorcet, monotone, and Pareto, but not independent.*

Remark 1.59. We'll see later than while you can different tradeoffs than this, you can't do *better* than this, in a specific way.

Proof. In the Copeland system, a Condorcet candidate gets a perfect, maximum score, since they beat each other candidate head-to-head. Each other candidate loses at least one head-to-head and so does not get a perfect score. So the Condorcet candidate always wins. This also means the system satisfies the majority property, since a majority candidate will be a Condorcet candidate.

An anti-Condorcet candidate loses each head-to-head and thus gets a score of zero. Since some candidate will have a score greater than zero, the anti-Condorcet candidate will always lose.

We now need to prove this method is monotone and Pareto. Again, this will require a bit of a sophisticated argument; we can't just generate an example, but need to provide an argument that will work for any possible profile, no matter how strange.

To prove Pareto: suppose A is above B on every preference list. In this case, B may win some head-to-head matchups, but any matchup B wins, A will also win. This means that A gets a point every time B gets a point, and A gets at least half a point every time B gets half a point. Further, A will defeat B, so A will have at least one more point than B does. So B cannot win.

To prove monotonicity: moving candidate A up on a preference list can never hurt A in any head-to-head matchup, so it cannot reduce A's score. A change that only involves moving A up can't affect any other head-to-head matchup, so it won't ever increase any other candidate's score. Therefore, if A was a winner before the change, they will still win after the change.

But even this argument may suggest why Copeland's method isn't independent. An improvement in A's place can't improve the scores of other candidates; but changing the relative positions of other candidates can change *their* scores. So consider the following two profiles:

B	A	A
C	B	C
A	C	B

→

B	A	C
C	B	A
A	C	B

Figure 1.8

In the first profile, A gets 2 points, and B gets one, so A is the unique winner. In the second profile, we have swapped the positions of A and C for one voter. Now C beats A, getting one point; A beats B, getting one point; and B beats C, getting one point. All three candidates are declared winners.

□

1.3.7 Black's method

Here is a new method that also tries to prioritize Condorcet winners.

Definition 1.60. *Black's method* is the social choice function that chooses the Condorcet candidate as the unique winner if there is a Condorcet candidate, and chooses the Borda count winner if there is not.

We can also think of this as taking the Borda count, and "fixing" the "glitch" where it doesn't satisfy the Condorcet property.

Proposition 1.61. *Black's method is majoritarian, Condorcet, anti-Condorcet, monotone, and Pareto, but not independent.*

Proof. By definition, if there's a Condorcet candidate, Black's method selects them as the unique winner. So Black's method is Condorcet, and therefore majoritarian.

We know that an anti-Condorcet candidate can't be the Condorcet candidate, and we've also see that they can't be a Borda count winner in the proof of proposition 1.54. So an anti-Condorcet candidate can't win in Black's method.

To see Black's method is monotone, we need to split into two cases. Suppose the Black winner was a Condorcet winner. Then moving them up in some rankings will not change that; they will still be the Condorcet winner and will still win. If the Black winner was not a Condorcet candidate, they must have been the Borda count winner. Moving them up in the rankings will still leave them as the Borda count winner. And no *other* candidate can become a Condorcet candidate when the original winner moves up in some rankings. So either the original winner becomes the Condorcet candidate and thus wins, or no one is a Condorcet candidate and the original winner still wins by Borda count. Thus Black's method is monotone.

To see the method is Pareto, imagine that A is favored over B by all voters. Then B cannot be the Condorcet candidate because they lose to A head-to-head, and they can't be the Borda count winner because A will have more points than B does. Thus B cannot win.

Finally, we know that Black's method violates independence, because it is Condorcet. We get a concrete example from figure 1.8: in the before profile, A is the Condorcet candidate and wins, but in the after profile, we get a three-way tie in which each candidate has three Borda points.

□

Proposition 1.62. *The dictatorship method is monotone, Pareto, and independent, but not Condorcet, anti-Condorcet, or majoritarian.*

Proof. If A is the winner, they are at the top of the dictator's preference list. No improvement on other preference lists can change that, so A remains the dictator's top choice, and thus still wins. So the method is monotone.

If A is higher on ever voter's list than B, then in particular A is above B on the dictator's list and so B does not win. So the dictatorship method is Pareto.

If A is at the top of the dictator's preference list in one profile, and then a second profile has A and B in the same relative position, it can't have B at the top of the dictator's preference list. So B can't win in the second profile. Thus the dictatorship method is independent.

But now consider the following profile, where the third voter is dictator.

A	A	B
B	B	A

Figure 1.9

A is the Condorcet and the majority candidate, while B is the anti-Condorcet candidate. But B wins.

□

Proposition 1.63. *The all-ties method and the monarchy method are monotone and independent, but not Condorcet, anti-Condorcet, majority, or Pareto.*

Proof. These methods are both what we call constant functions, meaning that the output to the method ignores the input completely and doesn't care how anyone votes. These methods are all monotone and independent, because no candidate can win on one profile and lose on another. But they also violate Condorcet, anti-Condorcet, majority, and Pareto, because some profile can have a candidate with a Condorcet candidate, an anti-Condorcet candidate, a majority candidate, or a candidate preferred universally to another candidate—and that has no effect on who wins. □

	anon	neu	unan	dec	maj	Con	AC	mono	Par	ind
Plurality	Y	Y	Y	N	Y	N	N	Y	Y	N
Antiplur	Y	Y	Y	N	N	N	N	Y	N	N
Borda	Y	Y	Y	N	N	N	Y	Y	Y	N
Hare	Y	Y	Y	N	Y	N	N	N	Y	N
Coombs	Y	Y	Y	N	N	N	N	N	Y	N
Copeland	Y	Y	Y	N	Y	Y	Y	Y	Y	N
Black	Y	Y	Y	N	Y	Y	Y	Y	Y	N
Dictator	N	Y	Y	Y	N	N	N	Y	Y	Y
All-ties	Y	Y	N	N	N	N	N	Y	N	Y
Monarchy	Y	N	N	Y	N	N	N	Y	N	Y
Bloc	N	Y	Y	N	N	N	N	Y	Y	N
Agenda	Y	N	Y	N	Y	Y	Y	Y	N	N

Figure 1.10: : A summary of the properties of various social choice methods

1.4 Arrow's Theorem

We've seen that it's impossible to produce a "perfect" social choice function. Fundamentally, the problem is cycles like this, first seen in the proof of proposition 1.50:

A	C	B
B	A	C
C	B	A

In this profile, a majority of voters prefer A to B, a majority prefer B to C, and a majority prefer C to A. Even though each individual voter has rational, transitive preferences, the electorate as a whole does not. This is known as the *Condorcet paradox*.

Your textbook presents the following extension that's even worse: suppose we have the preference order

D	C	B
E	D	C
F	E	D
G	F	E
A	G	F
B	A	G
C	B	A

Imagine these represent a series of different policies we could have on some issue, and that we start with policy D. Someone proposes policy C, and that wins by majority vote over D. Then someone else proposes an amendment to B, and that defeats C by majority vote. Then finally, someone else proposes an amendment to shift to A, and that defeats B. In each case the vote was a clear majority vote. But we have moved from D to A, even though every single voter strongly prefers D to A.

We say in proposition 1.50 that no social choice function can be both independent and Condorcet, due to these loops. But it's worse than that. The insight behind these Condorcet paradox loops lead the economist Kenneth Arrow to formulate the following theorem:

Theorem 1.64 (Arrow)**.** *If a social choice function with at least three candidates satisfies both Pareto and independence, then it must be a dictatorship.*

This is often phrased in the following way:

Theorem 1.65 (Arrow's theorem, alternate version)**.** *It is impossible for a social choice function with at least three candidates to be Pareto, independent, and non-dictatorial.*

Lemma 1.66 (decisiveness lemma)**.** *A social choice function with at least three candidates that satisfies Pareto and independence must be decisive.*

Proof. □

Corollary 1.67. *It is impossible for a method to satisfy Pareto, independence, anonymity, and neutrality.*

Proof of theorem 1.64. □

2 Apportionment

> Article I, Section 2 of the Constitution says that the representatives should be apportioned to the states “according to their respective numbers”. Unfortunately, the Constitution does not offer any further guidance about precisely how this should be done, and every 10 years Congress is left to implement an apportionment that meets the spirit of these words. It is not hard to determine each state’s “fair share” of the House of Representatives. . . . The difficulty is that this fair share need not be a whole number, and yet the number of representatives assigned to each state must be a whole number. So some process is needed to round these fair shares to whole numbers. This is the apportionment problem. What process should we use?
>
> At first, there does not seem to be much substance to this question. Can’t we simply choose the whole number that is nearest to the state’s fair share? It is only upon some reflection that one realizes that the problem is more subtle than this. Rounding the fair shares of all states to their nearest whole number may result in assigning too few or too many seats all told. What should we do? It turns out that Alexander Hamilton and Thomas Jefferson began a debate on this topic in 1792. There are a number of sensible approaches to the problem, but each has its flaws.

2.1 Congressional Apportionment and Hamilton’s Method

Suppose we have three states, and are going to allocate 100 Congressional seats to them. The states have the populations

- A: 4,400,000
- B: 45,300,000
- C: 50,300,000

How many seats should each state get?

The total population is 100,000,000, so it seems like each state should get 1 seat per million people. We would give A 4.4 seats, B 45.3 seats, and C 50.3 seats. But we can’t actually allocate fractions of a representative; every state needs a whole number of seats. It seems like A should get 4 seats, B should get 45 seats, and C should get 50 seats, with one seat left over. The question is: who should get it?

Alexander Hamilton suggested we should allocate "extra" seats in order of their fractional parts. So A gets 4 seats and has 0.4 left over; B gets 45 seats and has 0.3 left over; C has 50 seats and 0.3 left over. We give the extra seat to A, and the final lineup would be:

- A 5 seats
- B 45 seats
- C 50 seats.

And that all seems perfectly reasonable. But there are some bugs!

For instance, what if we instead want to allocate 101 seats? This requires some annoying arithmetic, for which we will need a calculator. When we divide 100,000,000 by 101 we get about 990099. (In fact we get $990099.\overline{0099}$.) Dividing each state's population by this number gives a rough-cut allocation of

- A gets $\frac{4,400,000}{990099} \approx 4.444$ seats;
- B gets $\frac{45,300,000}{990099} \approx 45.753$ seats;
- C gets $\frac{50,300,000}{990099} \approx 50.843$ seats.

Now we give C 50 seats, B 45 seats, and A 4 seats, with two left over. Looking at the fractional parts, the extras should go to B and C, giving a final tally of

- A 4 seats
- B 46 seats
- C 51 seats.

But there's something weird here. When we added an extra seat, B and C both gained seats, while A *lost* a seat. And that seems unfortunately backwards.

Now suppose we return to allocating 100 seats, but we take a new census, which give updated population numbers

- A: 4,500,000
- B: 45,200,000
- C: 49,000,000
- Total population: 98,700,000.

This looks like we should allocate one seat for every 987,000 people. Then our rough cut is

- A gets $\frac{4{,}500{,}000}{987000} \approx 4.559$ seats;
- B gets $\frac{45{,}200{,}000}{987000} \approx 45.795$ seats;
- C gets $\frac{49{,}000{,}000}{987000} \approx 49.645$ seats.

Our rough cut is that C gets 49, B gets 45, and A gets 4, with one left over. Then B has the largest fractional part, so our final allocation is

- A gets 4 seats
- B gets 46 seats
- C gets 50 seats.

Relative to our original situation, A has gained population while B and C have both lost it; but A has lost a seat to B.

Finally, let's think about what happens in the original situation if we add a fourth state D with 1,700,000 people, and correspondingly increase the legislature from 100 to 102. We get

- A: 4,400,000
- B: 45,300,000
- C: 50,300,000
- D: 1,700,000
- Total population: 101,700,000.

We should give one representative for every $\frac{101{,}700{,}000}{102} \approx 997{,}059$ people. So our rough allocations are

- A gets $\frac{4{,}400{,}000}{997{,}059} \approx 4.413$ seats;
- B gets $\frac{45{,}300{,}000}{997{,}059} \approx 45.434$ seats;
- C gets $\frac{50{,}300{,}000}{997{,}059} \approx 50.448$ seats.
- D gets $\frac{1{,}700{,}000}{997{,}059} \approx 1.701$ seats.

We give A 4 seats, B 45 seats, C 50 seats, and D 1 seat; this leaves 2 left over. One goes to D, and the second narrowly goes to C, for a final allocation of

- A gets 4 seats
- B gets 45 seats
- C gets 51 seats
- D gets 2 seats.

We added D to the nation, and added 2 seats for D's representatives, which all makes sense; but this also has the effect of moving one seat from A to C, even though nothing about A or C has changed. So there's a lot of weirdness going on here!

The method we've used here is called *Hamilton's method* and we'll study it more soon.

2.1.1 Defining apportionment

We first need to set up the basic notation we'll be using for the rest of section 2. We assume we have n states, where n is a whole number bigger than 1. (In the USA we've had $n = 50$ since 1959.) We need to allocate h Congressional seats among those states, and this also has to be a positive whole number. In the US House of Representatives, we have $h = 435$.

We list the states in some fixed order, and we use p_k to represent the population of the kth state. Then we use the letter p to represent the total population of the nation, so that

$$p = p_1 + p_2 + \cdots + p_n.$$

We call the collection of numbers $h, n, p_1, \ldots, p_n$ a *census*. (Notice we don't have to tell you the value of p since that's determined by the p_k.)

Definition 2.1. An *apportionment method* is a function whose input is a census $h, n, p_1, \ldots, p_n$, and whose output is a collection of positive integers $a_1, a_2, \ldots, a_n$ that add up to

$$a_1 + a_2 + \cdots + a_n = h.$$

We think of these outputs $a_1, \ldots, a_n$ as the number of congressional seats allocated to each state.

In our opening example, we had $n = 3$, $p_1 = 4{,}500{,}000$, $p_2 = 45{,}200{,}000, p_3 = 49{,}000{,}000$, and $h = 100$. Then in our first allocation we got $a_1 = 4, a_2 = 46, a_3 = 50$.

For the rest of this section, we will regularly find ourselves comparing two numbers and wanting to choose the larger one. This raises the question of what happens if the two numbers are the same size, but we're going to ignore that question completely, and just assume we won't have that sort of tie.

So imagine we put all 50 US states in alphabetical order. Then Maryland is the 20th state. (Your textbook says 19th but I've counted six times and I'm pretty sure it's 20th.) In the 2020 census, Maryland has $p_{20} = 6{,}185{,}278$ people, out of a total population of $p = 331{,}108{,}434$. Since we have $h = 435$, you'd want Maryland to have

$$\frac{p_1 9}{p} = \frac{6{,}185{,}278}{331{,}108{,}434} \approx 0.01868 = 1.868\%$$

of the total seats. In our current allocation, Maryland has $a_{20} = 8$ representatives; we compute that

$$\frac{a_{20}}{h} = \frac{8}{435} \approx 0.01839 \approx 1.839\%$$

of the Congressional representatives. We'd like to make those two percentages as close as possible, but there are two limitations.

The first is that they can never be equal, becasue we need to allocate a whole number of seats. The "ideal" allocation where Maryland has 1.868% of all the seats would give it

$$h \cdot \frac{p_{20}}{p} = 435 \cdot \frac{6{,}185{,}278}{331{,}108{,}434} \approx 435 \cdot 0.01868 \approx 8.126$$

Congressional seats, but we can't give 0.126 of a seat. So there's no perfect answer, and giving Maryland 9 seats would be even further off, giving it 2.069% of the seats for just 1.868% of the population.

The second constraint is that we need to allocate to every state, simultaneously. We could decide to be "generous" and round every state's Congressional delegation up. But that would change h. (Based on 2020 Census figures, h would have to increase to 460.) There's nothing intrinsically wrong with changing h; it's not like we couldn't build 25 more seats in the House of Representatives.

But if we change h to 460, then we'd have *different* ideal allocations. For instance, in the 2020 Census, Kentucky has $p_{17} = 4{,}509{,}342$ people. Its ideal allocation is

$$h \cdot \frac{p_{17}}{p} = 435 \cdot \frac{4{,}509{,}342}{331{,}108{,}434} \approx 435 \cdot 0.0136 \approx 5.924.$$

We could round this up to 6, and in fact the 2020 apportionment did so. But if we rounded *every* state up to get $h = 460$, the new ideal allocation would be

$$h \cdot \frac{p_{17}}{p} = 460 \cdot \frac{4{,}509{,}342}{331{,}108{,}434} \approx 460 \cdot 0.0136 \approx 6.265.$$

Do we round the allocation up *again* to 7? It's not clear this process would ever stop.

2.1.2 Quotas

We've repeatedly calculated the number $h \cdot {p_k}/{p}$, as the "idealized" number of Congressional seats a state should get. We will now call this the stat's *standard quota* $q_k = h \cdot {p_k}/{p}$. Algebraically we could also write this

$$q_k = \frac{p_k}{p/h}$$

and we call that denominator the *standard divisor* $s = \frac{p}{h}$.

If we could assign $a_k = q_k$ for each k then we would automatically have $a_1 + a_2 + \cdots + a_n = h$, which is what we want. Because

$$\begin{aligned} q_1 + q_2 + \cdots + q_n &= h \cdot \frac{p_1}{p} + h \cdot \frac{p_2}{p} + \cdots + h \cdot \frac{p_2}{p} \\ &= \frac{h}{p}(p_1 + p_2 + \cdots + p_n) \\ &= \frac{h}{p} \cdot p = h. \end{aligned}$$

However we (almost certainly) can't assign $a_k = q_k$ because the quota q_k is probably not a whole number, so we need to pick something else. Two obvious choices are the *lower quota*, which is the standard quota rounded down $\lfloor q_k \rfloor$, or the *upper quota*, which is the standard quota rounded up $\lceil q_k \rceil$.

So one way of looking at our next question is: which states do we round up, and which states do we round down? This can actually be a very important question in terms of actual political influence. In the 2010 Cenus, Montana had a standard quota of 1.40. This was in fact rounded down to give Montana one representative, making Montanans the most underrepresented state residents in the country. The average Congressional district had about 710,000 people, but Montana had 994,416.

In the 2020 census, Montana's standard quota rose to about 1.426, and the allocation method used rounded it up to 2 Congressional seats. These seats each have about 543,000 people, and are the smallest and most overrepresented districts in the US.

2.1.3 Hamilton's Method

Alexander Hamilton in 1792 suggested a method guaranteed to assign each state either its lower quota or its upper quota. We call these methods *quota methods*.

In order to talk about these we need to give some notation for arithmetic.

Definition 2.2. The *integer part* of a real number x is the greatest integer less than or equal to x. We will sometimes notate this $\lfloor x \rfloor$, which we read as the "floor" of x.

The *fractional part* of a real number x is the difference between x and its integer part. We can write this as $x - \lfloor x \rfloor$ or sometimes as $\text{frac}(x)$ or $\{x\}$.

Example 2.3. The integer part of 3.14159 is $\lfloor 3.14159 \rfloor = 3$ and the fractional part is $\{3.14159\} = 0.14159$.

The integer part of 8.126 is $\lfloor 8.126 \rfloor = 8$ and the fractional part is $\{8.126\} = 0.126$.

Note that a state's lower quota is the integer part of its standard quota, $\lfloor q_k \rfloor$. Hamilton thought that every state should be guaranteed to get its lower quota, and that excess seats should be allocated based on the size of the fractional part.

So in the 2020 Census, Maryland's standard quota is $q_{20} = 8.126$. Maryland is guaranteed 8 seats under Hamilton's method, but the fractional part $\{q_{20}\} = 0.126$ is small so we're inclined not to give it another representative. But we saw Kentucky has a standard quota of $q_{17} = 5.924$, with lower quota $\lfloor 5.924 \rfloor = 5$. We should guarantee Kentucky 5 seats, but since the fractional part of its quote $\{q_{17}\} = 0.924$ is quite large, we should be strongly inclined to give it a sixth seat, as indeed we do.

Definition 2.4 (Hamilton's method)**.** As a provisional apportionment, assign each state its lower quota $\lfloor q_k \rfloor$. Then assign the seats that remain to the states in decreasing order of the size of the fractional parts of their standard quotas, allocating at most one per state.

Example 2.5. Apportion $h = 10$ seats to $n = 3$ states with populations $p_1 = 264, p_2 = 361, p_3 = 375$.

We get a total population $p = 264 + 361 + 375 = 1000$. The standard divisor is $s = p/h = 1000/10 = 100$. That means we want to allocate roughly one seat per hundred people.

Our standard quotas are

$$q_1 = \frac{p_1}{s} = \frac{264}{100} = 2.64$$
$$q_2 = \frac{p_2}{s} = \frac{361}{100} = 3.61$$
$$q_3 = \frac{p_3}{s} = \frac{375}{100} = 3.75$$

The lower quotas are 2,3,3, so we allocate those 8 seats, and have 2 seats left over. The fractional parts are $0.64, 0.61, 0.75$. We assign the first seat to state 3 with fractional part

0.75, and the second seat to state 1 with fractional part 0.64. Then we are out of seats to assign, so state 2 sticks with its lower quota. The final apportionment is $a_1 = 3, a_2 = 3, a_3 = 4$, and the process is summarized in figure 2.1.

k	p_k	Standard Quota q_k	Lower Quota	Upper Quota	Fractional Part $\{q_k\}$	Hamilton Apportionment
1	264	2.64	2	3	0.64	3
2	361	3.61	3	4	0.61	3
3	375	3.75	3	4	0.75	4

Figure 2.1: Hamilton's Method in example 2.5

2.1.4 "Paradoxes" in Hamilton's method

Hamilton's method is maybe the most obvious way to solve the apportionment problem. It's so obvious that it seems like this should fully resolve the question. But Hamilton's method does have a number of issues it runs into which aren't ideal. In fact, these flaws are so bad that Hamilton's method is almost completely out of the running as an apportionment method. It was used to some degree in the late 19th century and early 20th, and many of these bad behaviors actually occurred in ways that made the method look much too bad.

Example 2.6 (Alabama Paradox)**.** Suppose we have $n = 3$ states and $h = 10$ seats to allocate to them. If they have populations of $p_1 = 1{,}450{,}000, p_2 = 3{,}400{,}000, p_3 = 5{,}150{,}000$, we get a total population of 10,000,000 and a standard divisor of 1,000,000. Then we get the following table:

k	p_k	Standard Quota q_k	Lower Quota	Fractional Part $\{q_k\}$	Hamilton Apportionment
1	1,450,000	1.45	1	0.45	2
2	3,400,000	3.40	3	0.40	3
3	5,150,000	5.15	5	0.15	5

Figure 2.2: The Alabama Paradox Part 1, in example 2.6

So far we have no problems. But now imagine Congress increases the number of seats from 10 to $h = 11$. Which state will gain the additional seat?

Well, now we have a standard divisor of $\frac{10{,}000{,}000}{11} \approx 909091$, and we get the following table:

k	p_k	Standard Quota q_k	Lower Quota	Fractional Part $\{q_k\}$	Hamilton Apportionment
1	1,450,000	1.595	1	0.595	1
2	3,400,000	3.740	3	0.740	4
3	5,150,000	5.665	5	0.665	6

Figure 2.3: The Alabama Paradox Part 2, in example 2.6

The lower quotas haven't changed, but the standard quotas and thus the fractional parts have increased. Now the winners are states 2 and 3: state 1 has actually *lost* a representative because Congress added a representative to the House.

This situation is called the *Alabama paradox* because Alabama would have lost a seat to an increase in h during the 1880 reapportionment. And it seems facially unfair that increasing h should cause any state to lose a seat.

Definition 2.7. An apportionment method is called *house monotone* if an increase in h, while all other parameters remain the same, can never cause any seat allocation a_k to decrease. Thus example 2.6 shows that Hamilton's method is not house monotone.

Why did this problem pop up? Every state's standard quota increased between figures 2.2 and 2.3, and in fact each standard quota increased by exactly 10%. But since state 1 started with $q_1 = 1.45$ and state 3 started with $q_3 = 5.15$, a 10% increase in the latter is much bigger than a 10% increase in the former.

Example 2.8 (Population paradox)**.** We can run into a similar issue when h stays the same but populations increase. Suppose we start again with the same data and same table as before, with $h = 10$:

But now suppose, keeping $h = 10$, we take a new census, and populations have increased in state 1, while decreasing in states 2 and 3. We get $p_1 = 1{,}470{,}000, p_2 = 3{,}380{,}000, p_3 = 4{,}650{,}000$, for a total population of $p = 9{,}500{,}000$. The standard divisor is now 950,000, and we get the following table of quotas:

State 1 has increased its standard quota q_1, while q_2 and q_3 have decreased. But since q_3 in particular has dropped from 5.15 to 4.89, the fractional part is bigger, and now states 2 and 3 both get extra seats, leaving state 1 with $a_1 = 1$.

k	p_k	Standard Quota q_k	Lower Quota	Fractional Part $\{q_k\}$	Hamilton Apportionment
1	1,450,000	1.45	1	0.45	2
2	3,400,000	3.40	3	0.40	3
3	5,150,000	5.15	5	0.15	5

Figure 2.4: The Population Paradox Part 1 in example 2.8

k	p_k	Standard Quota q_k	Lower Quota	Fractional Part $\{q_k\}$	Hamilton Apportionment
1	1,470,000	1.55	1	0.45	1
2	3,380,000	3.56	3	0.40	4
3	4,650,000	4.89	4	0.15	5

Figure 2.5: The Population Paradox Part 2 in example 2.8

This seems even worse! State 1 *gained* population, while states 2 and 3 lost it. But this leads state 1 to lose a representative.

Definition 2.9. A method is called *population monotone* if a state can never lose a seat when its population increases while no other state's population increases.

In algebraic terms, whenever $a_i' < a_i$ and $a_j' > a_j$, it must be the case either that $p_i' < p_i$ or $p_j' > p_j$.

Example 2.8 shows that Hamilton's method is not population monotone.

Finally we need to consider one more situation.

Example 2.10 (Oklahoma paradox)**.** Start again with the same data, where $h = 10$ and $p = 10{,}000{,}000$.

Imagine the union wants to add one more state, with a population of $p_4 = 2{,}600{,}000$. It seems like we'd expect this state to get 3 seats, so we add a corresponding 3 seats to the house for a new $h = 13$, a new population $p = 12{,}600{,}000$, and a new standard divisor of $s = \frac{12{,}600{,}000}{13} \approx 969{,}231$. We get the following table:

We see that state 4 does indeed get $a_4 = 3$ seats. But in the process, state 1 loses a seat, and state 2 gains one.

k	p_k	Standard Quota q_k	Lower Quota	Fractional Part $\{q_k\}$	Hamilton Apportionment
1	1,450,000	1.45	1	0.45	2
2	3,400,000	3.40	3	0.40	3
3	5,150,000	5.15	5	0.15	5

Figure 2.6: The Oklahoma Paradox Part 1 in example 2.10

k	p_k	Standard Quota q_k	Lower Quota	Fractional Part $\{q_k\}$	Hamilton Apportionment
1	1,450,000	1.50	1	0.50	1
2	3,400,000	3.51	3	0.51	4
3	5,150,000	5.31	5	0.31	5
4	2,600,000	2.68	2	0.68	3

Figure 2.7: The Oklahoma Paradox Part 2 in example 2.10

This is called the *new states paradox* or the *Oklahoma paradox*, due to the impact of adding Oklahoma as a state in 1907. At the time we were using Hamilton's method of apportionment. When Oklahoma was added it was allocated 5 representatives; but if we had redone the allocation New York would also have had to cede a state to Maine, which was seen as highly undesirable.

Hamilton's method is appealingly simple, but the examples 2.6, 2.8, and 2.10 show that it has major flaws. Is there a way to avoid these problems?

The answer is "sort of", and we'll see a new class of approaches in the next section.

2.2 Divisor Methods

2.2.1 Jefferson's Method

In section 2.1 we treated the house size h as a fixed input to our apportionment function. For almost a hundred years the US has worked that way; the House of Representatives has had 435 members since 1930. But prior to that, the size of Congress varied, slowly increasing over time.

This suggests another approach: we could have some rule for how many representatives each state gets, based on its population, and then let h vary.

Jefferson approached things this way. He started with a theory of how large a Congressional district should be. (In his case, he pointed out the Constitution says the "number of Representatives shall not exceed one for every thirty thousand", which suggested that 30,000 people was about the appropriate size for a district.)

In section 2.1.1 we found the standard divisor $s = p/h$, which was the total population divided by the number of seats to be apportioned. We can think of this as the goal size for a Congressional district. But instead we could start by picking a number d, which we think of as an appropriate size for a district. We call this number a *modified divisor*. We can then divide each state's population by d to obtain its *modified quota* p_k/d, just as we found the quota $q_k = p_k/s$ in the last section.

As before, the modified quota is probably not a whole number, so we can't just give every state p_k/d seats. Jefferson suggested we should round each modified quota down to generate the number of seats, and thus set $a_k = \lfloor p_k/d \rfloor$. This has the advantage of being uniform: everyone gets rounded down.

Poll Question 2.2.1. Why did Jefferson suggest we round down, rather than up?

In Jefferson's original method, we start with d, compute each $a_k = p_k/d$, and then find $h = a_1 + a_2 + \cdots + a_n$. But we can also apply his approach to generate an apportionment function in the sense of definition 2.1, which starts with a fixed value of h.

Definition 2.11 (Jefferson's method)**.** Choose a modified divisor d, compute the modified quotas p_k/d, and round these down to obtain $a_k = \lfloor p_k/d \rfloor$. If $a_1 + a_2 + \cdots + a_n = h$, then we have the Jefferson apportionment. Otherwise, choose a new d and try again.

This raises a few of obvious questions.

1. Is there always a d that will work?

2. Is there *more than one* d that will work?

3. If we pick two different ds that both give the same total number of seats, will they give the same apportionment?

The answer to the first question is yes, as long as we ignore the possibility of exact ties in our calculations (which shouldn't happen with large population numbers). We can find the d that works pretty easily. If we get a total number of seats larger than our desired h, that means our d was too small and we need to make it bigger; but if we fewer seats than our desired h, our d is too big and we need to make it smaller. This lets us zero in on the correct d value pretty quickly. (We'll talk about how to do this in practice in section 2.2.2.)

The answer to the second question is also yes: we can pretty much always find multiple d values that will give us the correct number of seats. That makes question 3 especially important; but fortunately it is completely answered by the following proposition.

Proposition 2.12. *Suppose $h, n,$ and $p_1, \dots, p_n$ are given as inputs to our apportionment function. If d and d' are two different divisors, yielding Jefferson apportionments $a_1, \dots, a_n$ and $a'_1, \dots, a'_n$ respectively, then $a_k = a'_k$ for each state k.*

Remark 2.13. We will often be using notation like d' in this section. We read this as "d prime", and use it a lot when we want two different versions of the same number. If we need three, we may write d, d', d''.

There are other solutions to this problem but they're often more annoying. We could write d_1 and d_2, but the as already have subscripts; they could become $a_{1,k}$ and $a_{2,k}$ but that's unpleasant. We could also do something like using Greek letters, making our divisors d and δ and our apportionments a_k and α_k, but then we need to know the Greek alphabet and that still doesn't help when we need three versions.

Proof. Suppose "without loss of generality" that $d \leq d'$. Then, for every state k, the modified quota p_k/d for the first divisor d is at least as big as the modified quota p_k/d' for the second divisor d'. Rounding them down preserves that, so we must have $a_k \geq a'_k$, and that must be true for every state k. That means that

$$a_1 + a_2 + \cdots + a_n \geq a'_1 + a'_2 + \cdots + a'_n.$$

But we know that both apportionments give the same total number of seats, so

$$h = a_1 + \cdots + a_n \geq a_1 + \cdots + a'_n = h$$

and thus the two totals must be the same. □

Example 2.14. Suppose we have $n = 3, h = 10$, and three states with

- $p_1 = 1{,}500{,}000$
- $p_2 = 3{,}200{,}000$
- $p_3 = 5{,}300{,}000$

We need to try some different divisors here.

If we try the standard divisor $s = 1{,}000{,}000$ we get $p_1/s = 1.5$ and thus $a_1 = 1$; $p_2/s = 3.2$ and thus $a_2 = 3$; and $p_3/s = 5.3$ and so $a_3 = 5$. This gives a total number of seats equal to 9. (This shouldn't be a surprise from the work we did last section!)

This makes our house too small, so we need to make d smaller. If we try $d = 900{,}000$ we get modified quotas of $1.\overline{66}$, $3.\overline{55}$, and $5.\overline{88}$, which again would give $a_1 = 1, a_2 = 3, a_3 = 5$ for a total of nine seats.

Making d smaller again we can take $d = 800{,}000$. This gives modified quotas of $1.875, 4, 6.625$ and $a_1 = 1, a_2 = 4, a_3 = 6$ and a total of eleven seats. This is too many, so we need to make d bigger.

Since 800,000 was too small and 900,000 was too big, we might try $d = 850{,}000$. Then we get modified quotas of about $1.76, 3.67, 6.24$ and we get $a_1 = 1, a_2 = 3, a_3 = 6$, and a total of ten seats. This works, so the Jefferson apportionment is $a_1 = 1, a_2 = 3, a_3 = 6$.

		$s = 1{,}000{,}000$		$d = 900{,}000$		$d = 800{,}000$		$d = 850{,}000$	
k	p_k	Quota	Round Down	Quota	Round Down	Quota	Round Down	Quota	Round Down
1	1,500,000	1.50	1	$1.\overline{66}$	1	1.875	1	1.76	1
2	3,200,000	3.20	3	$3.\overline{55}$	3	4	4	3.67	3
3	5,300,000	5.30	5	$5.\overline{88}$	5	6.625	6	6.24	6
Total	10,000,000		9		9		11		10

Figure 2.8: The Jefferson method in example 2.14

Note that this is different from the Hamilton apportionment, which would start with our (unmodified) quotas of $1.5, 3.2$, and 5.3 to generate preliminary allocations of 1,3,5, and then allocate the extra seat to state 1, which has the largest fractional part. So the Hamilton apportionment would be $a_1 = 2, a_2 = 3, a_3 = 5$.

2.2.2 Critical Divisors

We can calculate Jefferson's method with that sort of trial and error approach, but there's a more direct way.

The breakpoints in the number d happen when one of the states would get one additional representative. So for instance, with $d = 850{,}000$, state 3 has a quota of 6.24 which rounds down to 6. Another way of thinking about that sentence is that $\frac{5{,}300{,}000}{6} \geq 850{,}000$, but $\frac{5{,}300{,}000}{7} < 850{,}000$. A little work should convince you that $\frac{5{,}300{,}000}{5} \geq 850{,}000$ and $\frac{5{,}300{,}000}{8} <$ 850,000. So there are exactly six positive integers m so that $\frac{p_3}{m} \geq 850{,}000$, which are 1, 2, 3, 4, 5, 6. And that's the number of seats that state 3 was allocated!

Definition 2.15. We call a number of the form $\frac{p_k}{m}$ for a positive integer k a *(Jefferson) critical divisor* for the state k.

This gives us a better way of thinking about hunting for a good modified divisor. When we choose a modified divisor d, each seat will get a number of seats equal to the number of critical divisors that are greater than or equal to d. So we want to pick a d such that, when we consider all states together, there are exactly h critical divisors greater than or equal to d.

Example 2.16. Let's use this method to look at example 2.14 again. We had populations $p_1 = 1{,}500{,}000$, $p_2 = 3{,}200{,}000$, and $p_3 = 5{,}300{,}000$. We can make a table of the critical divisors:

d	State 1	State 2	State 3
1	1,500,000/1 = 1,500,000	3,200,000/1 = 3,200,000	5,300,000/1 = 5,300,000
2	1,500,000/2 = 750,000	3,200,000/2 = 1,600,000	5,300,000/2 = 2,650,000
3	1,500,000/3 = 500,000	3,200,000/3 = 1,066,667	5,300,000/3 = 1,766,667
4	1,500,000/4 = 375,000	3,200,000/4 = 800,000	5,300,000/4 = 1,325,000
5	1,500,000/5 = 300,000	3,200,000/5 = 640,000	5,300,000/5 = 1,060,000
6	1,500,000/6 = 250,000	3,200,000/6 = 533,333	5,300,000/6 = 883,333
7	1,500,000/7 = 214,857	3,200,000/7 = 457,143	5,300,000/7 = 757,143

Figure 2.9: The Jefferson method with critical divisors from example 2.16

Counting down, we get: 5,300,000, 3,200,000, 2,650,000, 1,766,667, 1,600,000, 1,500,000, 1,325,000, 1,066,667, 1,060,000, and 883,333 makes ten. Then the eleventh number down is 800,000, so we can pick any number between 883,333 and 800,000. 850,000 is of course one of them, and that's why that divisor worked.

Now that we picked $d = 850{,}000$, we can also read that off our table. We see one number bigger than 850,000 in the first column, three in the second, and six in the third. And that gives us $a_1 = 1, a_2 = 3, a_3 = 6$.

This critical divisors approach has a couple advantages. One is that we don't have to do a trial-and-error search for d; we just have a straightforward calculation.

We can actually speed this up a bit more. We know the correct Jefferson divisor will always be smaller than the standard divisor $s = p/h$. So we can start in with our critical divisors there. For instance, if we look at our same example 2.14 one more time, we see that the standard divisor s would allocate nine seats. We can calculate the *next* critical divisor for each state, and we know we need to move past exactly one of them.

We see in figure 2.10 that if the standard divisor allocates a_k seats, we have a critical divisor at p_k/a_k, and the *next* critical divisor for that state will happen at $\frac{p_k}{a_k+1}$. So we can compute $\frac{p_k}{a_k+1}$ for each state, and allocate our next seat to whichever state has the biggest (next) critical divisor.

		$s = 1{,}000{,}000$			$d = 850{,}000$	
k	p_k	Standard Quota	Lower Quota a_k	$\frac{p_k}{a_k+1}$	Modified Quota	Jefferson allocation
1	1,500,000	1.50	1	750,000	1.76	1
2	3,200,000	3.20	3	800,000	3.76	3
3	5,300,000	5.30	5	883,333	6.24	6
Total	10,000,000		9			10

Figure 2.10: The Jefferson method with critical divisors, starting with the standard divisor

But note that if we want to allocate another seat, we may have to do the whole process over. If we allocate a seat to state 3, we will need to compute another critical divisor for state 3. In figure 2.10, if we want to allocate an eleventh seat, the next critical divisor for state 1 is still 750,000, and the next critical divisor for state 2 is still 800,000. But we'll also need to check the next critical divisor for state 3. This is

$$\frac{5{,}300{,}000}{6+1} = \frac{5{,}300{,}000}{7} = 757{,}143.$$

This is less than 800,000, so the eleventh seat goes to state 2. (But note it's bigger than 750,000; if we allocate a twelfth seat that would go to state 3 again, not to state 1.)

2.2.3 Results on Jefferson's Method

Jefferson's method favors large states, relative to Hamilton's method.

Example 2.17. Take $n = 2, h = 10, p_1 = 1{,}800{,}000$, and $p_2 = 8{,}200{,}000$. The standard divisor is $s = p/h = 1{,}000{,}000$. Jefferson's method finds a modified divisor of $d = 910{,}000$. Hamilton's method gives state 1 a second seat; Jefferson's method gives nine seats to state 2.

k	p_k	Standard Quota q_k	Hamilton Apportionment	Modified Quota $d = 910{,}000$	Jefferson Apportionment
1	1,800,000	1.8	2	1.98	1
2	8,200,000	8.2	8	9.02	9

These results can get even more extreme, in a fairly surprising way.

Example 2.18. Suppose we have $n = 4, h = 10$, with populations given in figure 2.11.

k	p_k	Standard Quota q_k	Hamilton Apportionment	Modified Quota $d = 800{,}000$	Jefferson Apportionment
1	1,500,000	1.5	2	1.88	1
2	1,400,000	1.4	1	1.75	9
3	1,300,000	1.3	1	1.62	1
4	5,800,000	5.8	6	7.25	7

Figure 2.11: Jefferson's method violates quota in example 2.18

This time Hamilton's method rounds the fourth state's allocation up from $q_4 = 5.8$ to 6. But Jefferson's method allocates even more representatives, giving the state 7 in total.

Definition 2.19. We say it's a *quota violation* if an apportionment method gives a state more representatives than its upper quota, or less than its lower quota.

An apportionment method satisfies the *quota rule* if it assigns every state either its lower quota or its upper quota.

This seems like a reasonable rule that we would want an apportionment method to satisfy, but Jefferson's method does not. So we're going to have to talk about it.

Sometimes it's useful to split this rule up into two pieces.

Definition 2.20. An apportionment method satisfies the *upper quota rule* if it never assigns a state more than its upper quota.

A violation of this rules is an *upper quota violation*.

Definition 2.21. An apportionment method satisfies the *lower quota rule* if it never assigns a state less than its lower quota.

A violation of this rules is an *lower quota violation*.

Proposition 2.22. *Jefferson's method satisfies the lower quota rule.*

Proof. If we give every state a provisional apportionment equal to its lower quota, that's the Jefferson apportionment using the standard quota as a modified divisor. This will apportion fewer than h seats. Since we want to apportion h seats, we have to lower the modified divisor, which can only give states more seats, not fewer. Thus each state gets at least its lower quota. □

Proposition 2.23. *Jefferson's method satisfies the house monotonicity criterion.*

Modified divisor perspective. An increase in h will lead to a smaller modified divisor d. Decreasing d gives a higher modified quota p_k/d for each state k. Rounding down a *larger* number will never give a *smaller* number, so reducing d can't ever cause a state to lose a seat. □

Critical divisor perspective. Suppose we list all the critical divisors in decreasing order. The critical divisor p_k/m is associated with state k. If we choose the first h divisors in the list, the seats are assigned to the states that are associated to each divisor. If we increase h to $h+1$, then we will choose the first $h+1$ divisors instead, which will include the first h divisors (and then one more). So all the seats we originally allocated will still be allocated to their original states, and then we will allocate one more. So each state will still get at least as many seats as they did when we allocated h seats.

□

2.2.4 Other divisor methods

Jefferson's method is an example of a *divisor method*, in which the populations of the states are divided by modified divisors to obtain modified quotas, which are then rounded to whole numbers.

The way divisor methods differ is in the rounding method. Jefferson's method rounds modified quotas down to the nearest whole number, but there are other options.

Definition 2.24 (Adams's method)**.** Choose a modified divisor d. Compute the modified quotas p_k/d, and round each of these numbers *up* to obtain a_k. If $a_1 + a_2 + \cdots + a_n = h$ then we have the Adams apportionment.

Jefferson's method has a bias toward large states (like Virginia!). Adams's method has a bias toward small states (like much of New England). Consider two situations which each have a total population of $p = 10{,}000{,}000$ and a house size of $h = 10$.

Example 2.25. First, we give state 1 $p_1 = 1{,}800{,}000$ people and give state 2 $p_2 = 8{,}200{,}000$ people. Hamilton's method gives two seats to state 1. In Jefferson's method, we round down the standard quota to get 1 and 8, and then we compute the critical divisors

$$\frac{p_1}{a_1+1} = \frac{1{,}800{,}000}{2} = 900{,}000$$
$$\frac{p_2}{a_2+1} = \frac{8{,}200{,}000}{9} = 911{,}111$$

so state 2 gains another seat at $d = 911{,}111$ and state 1 gains another seat at $d = 900{,}000$. We can take any divisor in between to get the Jefferson apportionment; we use 910,000.

In Adams's method we round the standard quota up to get 2 and 9, which is too many. We compute critical divisors and we get

$$\frac{p_1}{a_1} = \frac{1{,}800{,}000}{1} = 1{,}800{,}000$$
$$\frac{p_2}{a_2} = \frac{8{,}200{,}000}{8} = 1{,}025{,}000$$

State 2 loses a seat when the divisor hits $d = 1{,}025{,}000$ and state 1 loses a seat when the divisor hits $d = 1{,}800{,}000$, so we can pick any divisor between those numbers. (Actually, that's not quite true: state 2 will lose *another* seat at $d = 1{,}171{,}429$ so we need to pick a $1{,}025{,}000 \leq d < 1{,}171{,}429$. We chose 1,100,000.

		$d = 1{,}000{,}000$		$d = 910{,}000$		$d = 1{,}100{,}000$	
k	p_k	Standard Quota	Hamilton a_k	Jefferson Quota	Jefferson a_k	Adams Quota	Adams a_k
1	1,800,000	1.8	2	1.98	1	1.64	2
2	8,200,000	8.2	8	9.02	9	7.45	8

Example 2.26. Now instead we give state 1 $p_1 = 1{,}200{,}000$ people and give state 2 $p_2 = 8{,}800{,}000$ people. Hamilton's method gives one seat to state 1 and two to state 2. In

Jefferson's method, we round down the standard quota to get 1 and 8, and then we compute the critical divisors

$$\frac{p_1}{a_1+1} = \frac{1{,}200{,}000}{2} = 600{,}000$$
$$\frac{p_2}{a_2+1} = \frac{8{,}800{,}000}{9} = 977778$$

so state 2 gains another seat at $d = 977{,}778$ and state 1 gains another seat at $d = 600{,}000$. We need to take a divisor between those to get the Jefferson apportionment; in fact we need one bigger than $8{,}800{,}000/10 = 880{,}000$. In the table below we pick $d = 900{,}000$.

In Adams's method we round the standard quota up to get 2 and 9, which is too many. We compute critical divisors and we get

$$\frac{p_1}{a_1} = \frac{1{,}200{,}000}{1} = 1{,}200{,}000$$
$$\frac{p_2}{a_2} = \frac{8{,}800{,}000}{8} = 1{,}100{,}000$$

State 2 loses a seat when the divisor hits $d = 1{,}100{,}000$ and state 1 loses a seat when the divisor hits $d = 1{,}200{,}000$, so we can pick any divisor between those numbers. We choose 1,150,000.

		$d = 1{,}000{,}000$		$d = 900{,}000$		$d = 1{,}150{,}000$	
k	p_k	Standard Quota	Hamilton a_k	Jefferson Quota	Jefferson a_k	Adams Quota	Adams a_k
1	1,200,000	1.2	1	1.33	1	1.04	2
2	8,800,000	8.8	9	9.78	9	7.65	8

Proposition 2.27. *Adams's method violates the lower quota rule.*

Proof. Exercise □

Adams's method also automatically meets the Constitutional requirement that each state gets at least one Congressional seat, which Hamilton and Jefferson do not.

We saw that Jefferson's method benefits large states by rounding down, while Adams's method benefits small states by rounding up. A reasonable response is to split the difference, and round "normally"—where fractional parts of 0.5 or greater are rounded up, while fractional parts less than 0.5 are rounded down.

Definition 2.28 (Webster's method)**.** Choose a modified divisor d. Compute the modified quotas p_k/d, and round each of these numbers to the nearest whole number, whether up or down, to obtain a_k. If $a_1 + a_2 + \cdots + a_n = h$ then we have the Webster apportionment.

Example 2.29. We can compute apportionments for all of these methods.

		$d = 1{,}000{,}000$		$d = 850{,}000$		
k	p_k	Standard Quota	Hamilton a_k	Jefferson critical divisor	Modified Jefferson Quota	Jefferson a_k
1	3,300,000	3.3	3	825,000	3.88	3
2	5,100,000	5.1	5	850,000	6	6
3	1,600,000	1.6	2	800,000	1.88	1

		$d = 1{,}100{,}000$			$d = 1{,}000{,}000$		
k	p_k	Adams critical divisor	Modified Adams Quota	Adams a_k	Webster critical divisor	Modified Webster Quota	Webster a_k
1	3,300,000	1,100,000	3	3	942,857	3.3	3
2	5,100,000	1,020,000	4.64	5	927,273	5.1	5
3	1,600,000	1,600,000	1.45	2	1,066,667	1.6	2

We can generalize this even further. And we're going to have to if we want to understand the method the US actually uses today.

2.2.5 Other rounding methods

Definition 2.30. A *rounding function* is a function that takes in a real number, outputs an integer, and has the following two properties:

1. If x is an integer, then $f(x) = x$.
2. If $x > y$, then $f(x) \geq f(y)$.

That is, every integer rounds to itself, and a bigger number will never round to a result less than a smaller number rounds to.

We've seen three rounding functions so far.

- The floor $f(x) = \lfloor x \rfloor$ rounds to the greatest integer less than or equal to x. This was used in Jefferson's method.

- The ceiling function $f(x) = \lceil x \rceil$ rounds to the least integer greater than or equal to x. This was used in Adams's method.
- The "regular" rounding where $f(x)$ is the integer nearest to x, used in Webster's method.

We sometimes call this last method *arithmetic rounding*. This is because the cutoff is the "arithmetic mean", where the average of m and $m+1$ is $\frac{m+(m+1)}{2} = m + \frac{1}{2}$.

Instead we can use as the cutoff the so-called *geometric mean*, which is $\sqrt{m(m+1)}$. That is, instead of adding two numbers and then dividing by 2, we multiply the numbers and take the square root. Then our rounding function will give $f(x) = m$, where m satisfies $\sqrt{m(m-1)} \leq x < \sqrt{m(m+1)}$.

This geometric rounding might seem a little weird. It is great for measuring things like average growth rates, which behave multiplicatively. But the main reason it's worth discussing is that we use it in the US today.

Definition 2.31 (Hill's method)**.** Choose a modified divisor d. Compute the modified quotas p_k/d, and round each of these numbers geometrically to obtain a_k. If $a_1 + a_2 + \cdots + a_n = h$ then we have the Hill apportionment.

Joseph Hill was the chief statistician at the Census Bureau from 1909 to 1921. He suggested this method, and in 1941, Congress officially made Hill's method the permanent apportionment method in the US, which it has been ever since.

We can also consider one more rounding function.

Definition 2.32. The *harmonic mean* of two numbers is the reciprocal of the average of their reciprocals

$$\frac{1}{\frac{\frac{1}{x}+\frac{1}{y}}{2}} = \frac{2}{\frac{1}{x}+\frac{1}{y}} = \frac{2}{\frac{y+x}{xy}} = \frac{2xy}{x+y}.$$

If we want to take the harmonic mean of two adjacent numbers, we get $\dfrac{2m(m+1)}{2m+1}$.

The function $f(x) = m$ where m satisfies the condition

$$\frac{2m(m-1)}{2m-1} \leq x < \frac{2m(m+1)}{2m+1}$$

is the *harmonic rounding* function.

Definition 2.33. Choose a modified divisor d. Compute the modified quotas p_k/d, and round each of these numbers harmonically to obtain a_k. If $a_1 + a_2 + \cdots + a_n = h$ then we have the Dean apportionment.

This method was considered in 1830, but was never used.

Because these rounding functions can be hard to work with, it's useful to have a table of cutoffs. This is table 8.7 in your textbook, which is reproduced here:

Figure 2.12: Cutoffs for rounding small numbers according to various methods

	Rounding Function and Method				
	Rounding Up	Harmonic Rounding	Geometric Rounding	Arithmetic Rounding	Rounding Down
	Adams	Dean	Hill	Webster	Jefferson
0–1	0	0	0	0.5	1
1–2	1	1.333	1.414	1.5	2
2–3	2	2.400	2.449	2.5	3
3–4	3	3.429	3.464	3.5	4
4–5	4	4.444	4.472	4.5	5
5–6	5	5.455	5.477	5.5	6
6–7	6	6.462	6.481	6.5	7
7–8	7	7.467	7.484	7.5	8

So for instance, the number 1.42 would be rounded down to 1 by arithmetic rounding (and of course by Jefferson rounding down), but it would be rounded up by geometric rounding and harmonic rounding. The number 5.46 would be rounded up by harmonic rounding, but down by geometric rounding and arithmetic rounding.

Just like with Jefferson's method, we can compute all the divisor apportionments with a critical divisor approach. For the Jefferson method, we computed numbers of the form p_k/m, since we were looking for breakpoints where some state would get a new representative. In effect, we wanted to know when $p_k/d = m$ was an integer, and then we rearranged to get the divisor $d = p_k/m$.

In other methods, these break points don't occur when p_k/d is an integer, but when instead when p_k/d crosses one of the rounding thresholds that we saw in figure 2.12.

For instance, in the Hill method with geometric rounding, our apportionments change when we cross a cutoff of the form $\sqrt{m(m+1)}$. So we want to find out when $p_k/d = \sqrt{m(m+1)}$, and thus rearranging, we compute all the divisors $d = \frac{p_k}{\sqrt{m(m+1)}}$. We call these

Hill critical divisors; they are the values of d at which a state acquires an additional seat.

You might notice a bit of a glitch here: when $m = 0$ we wind up dividing by zero. We're just going to think of this as an infinitely large number, corresponding to the fact that under Hill's method, each state is guaranteed at least one representative. Then we write the rest of our numbers in order from largest to least, and take the h largest (including the infinite ones, one for each state). An appropriate value for the Hill divisor is any number between the hth largest element and the $h+1$st largest element in this list.

It's useful to have summarized what critical divisors look like for each method.

Figure 2.13: Summary of critical divisor formulae

Method	Critical divisor for state k
Adams	$\frac{p_k}{a_k}$
Dean	$\frac{p_k(2(a_k+1))}{2a_k(a_k+1)}$
Hill	$\frac{p_k}{\sqrt{a_k(a_k+1)}}$
Webster	$\frac{p_k}{(a_k+1/2)}$
Jefferson	$\frac{p_k}{a_k+1}$

Example 2.34. Consider a problem with $n = 3, h = 10$, and populations of 1,385,000; 2,390,00; and 6,225,000. The lower quotas are 1,2, and 6, and the Hamilton apportionments will be 1, 3, and 6. If we round arithmetically for Webster's method, all of these will round down and we get 1, 3, and 6 again. We need to lower d to get an apportionment. We can compute the Webster critical divisors:

$$\frac{p_1}{a_1+1/2} = \frac{1{,}385{,}000}{1.5} = 923{,}333$$
$$\frac{p_2}{a_2+1/2} = \frac{2{,}390{,}000}{2.5} = 956{,}000$$
$$\frac{p_3}{a_3+1/2} = \frac{6{,}225{,}000}{6.5} = 957{,}692.$$

We allocate the next seat at 957,692 so we need to pick a d smaller than that, but bigger than 956,000. In the table below we use $d = 957{,}000$. This gives us a final seat allocation $a_1 = 1, a_2 = 2, a_3 = 7$.

If we round geometrically for Hill's method, all three standard quotas also round down. That means we need to lower d again. We compute the Hill critical divisors

$$\frac{p_1}{\sqrt{a_1(a_1+1)}} = \frac{1{,}385{,}000}{\sqrt{1\cdot 2}} = \frac{1{,}385{,}000}{1.414} = 979{,}343$$
$$\frac{p_2}{a_2(a_2+1)} = \frac{2{,}390{,}000}{\sqrt{2\cdot 3}} = \frac{2{,}390{,}000}{2{,}449} = 975{,}713$$
$$\frac{p_3}{a_3(a_3+1)} = \frac{6{,}225{,}000}{\sqrt{6\cdot 7}} = \frac{6{,}225{,}000}{6.481} = 960{,}538.$$

So the next delegate is allocated at $d = 979{,}343$, to state 1. We need a d smaller than that, but bigger than 975,713. In the chart below we use $d = 977{,}000$.

		$d = 1{,}000{,}000$		$d = 957{,}000$		$d = 977{,}000$	
k	p_k	Standard Quota	Hamilton a_k	Webster Quota	Webster a_k	Hill Quota	Hill a_k
1	1,385,000	1.385	1	1.447	1	1.418	2
2	2,390,000	2.390	3	2.497	2	2.446	2
3	6,225,000	6.225	6	6.505	7	6.372	6

We could also just do the calculation by making tables of critical divisors. For Adams we would get:

m	State 1	State 2	State 3
0	$1{,}385{,}000/0 = \infty$	$2{,}390{,}000/0 = \infty$	$6{,}225{,}000/0 = \infty$
1	$1{,}385{,}000/1 = 1{,}385{,}000$	$2{,}390{,}000/1 = 2{,}390{,}000$	$6{,}225{,}000/1 = 6{,}225{,}000$
2	$1{,}385{,}000/2 = 692{,}500$	$2{,}390{,}000/2 = 1{,}195{,}000$	$6{,}225{,}000/2 = 3{,}112{,}500$
3	$1{,}385{,}000/3 = 461{,}667$	$2{,}390{,}000/3 = 796{,}667$	$6{,}225{,}000/3 = 2{,}075{,}000$
4	$1{,}385{,}000/4 = 346{,}250$	$2{,}390{,}000/4 = 597{,}500$	$6{,}225{,}000/4 = 1{,}556{,}250$
5	$1{,}385{,}000/5 = 277{,}000$	$2{,}390{,}000/5 = 478{,}00$	$6{,}225{,}000/5 = 1{,}245{,}000$
6	$1{,}385{,}000/6 = 230{,}833$	$2{,}390{,}000/6 = 398{,}333$	$6{,}225{,}000/6 = 1{,}037{,}500$

The ten largest "numbers" (including infinity three times) are in blue on the table; that gives an allocation of $a_1 = 2, a_2 = 2, a_3 = 6$. This corresponds to any divisor larger than 1,195,000 but less than 1,245,000; indeed we can pick $d = 1{,}200{,}000$ and get lower quotas of 1.154, 1.991, and 5.188, which round up to 2, 2, and 6.

Taking the same approach for Hill's method gives us the table:

m	State 1	State 2	State 3
0	$1{,}385{,}000/0 = \infty$	$2{,}390{,}000/0 = \infty$	$6{,}225{,}000/0 = \infty$
1	$1{,}385{,}000/\sqrt{2} = 979{,}343$	$2{,}390{,}000/\sqrt{2} = 1{,}689{,}985$	$6{,}225{,}000/\sqrt{2} = 4{,}401{,}740$
2	$1{,}385{,}000/\sqrt{6} = 565{,}424$	$2{,}390{,}000/\sqrt{6} = 975{,}713$	$6{,}225{,}000/\sqrt{6} = 2{,}541{,}346$
3	$1{,}385{,}000/\sqrt{12} = 399{,}815$	$2{,}390{,}000/\sqrt{12} = 689{,}934$	$6{,}225{,}000/\sqrt{12} = 1{,}797{,}003$
4	$1{,}385{,}000/\sqrt{20} = 309{,}695$	$2{,}390{,}000/\sqrt{20} = 689{,}934$	$6{,}225{,}000/\sqrt{20} = 1{,}391{,}952$
5	$1{,}385{,}000/\sqrt{30} = 252{,}865$	$2{,}390{,}000/\sqrt{30} = 436{,}352$	$6{,}225{,}000/\sqrt{30} = 1{,}136{,}524$
6	$1{,}385{,}000/\sqrt{42} = 213{,}710$	$2{,}390{,}000/\sqrt{42} = 368{,}785$	$6{,}225{,}000/\sqrt{42} = 960{,}538$

The tenth number on this list is 979,343, and the eleventh is 975,713, so we can pick any number between them. In the table above we used 977,000.

2.3 Evaluating Apportionment Methods

We want a method to treat all states the same, only taking their populations into account.

Definition 2.35. An apportionment method is *neutral* if permuting the populations of states permutes the resulting numbers of seats in the same way.

This doesn't prevent a bias against large states, or small states; Adams's and Jefferson's methods are both neutral in this sense. But it prohibits a bias against Western states, or states that start with a vowel. (Arguably our current method is not neutral, if you include Washington D.C. on your list of states to be apportioned.)

Definition 2.36. An apportionment method is *proportional* if it produces the same result for two censuses with the same house size, and the same relative populations p_k/p.

The idea here is that we shouldn't care about the absolute numbers of people, just their relative share of the population. We say the *population distribution* is the list $p_1/p, p_2/p, \ldots, p_n/p$, which tells you what fraction of the total population each state has. A proportional apportionment method can be computed from just the population distribution, without knowing the absolute numbers of people.

The core intuition here is that if *every* state doubles its population, this shouldn't affect the apportionment at all.

Proposition 2.37. *Hamilton's method, and every divisor method, is proportional.*

Proof. Hamilton's method depends only on the standard quotas $q_k = p_k/s$. But the standard divisor $s = p/h$, so we can rewrite this as $p_k = h\frac{p_k}{p}$, so the standard quota depends only on the house size h and the population distribution p_k/p.

For divisor methods, we are dividing p_k by a divisor d. If all populations increase by a factor of c, we can increase the divisor d by that same factor and get the same modified quotas, with the same rounded result.

Alternatively, if we think in terms of critical divisors, we are writing the critical divisors in order from greatest to least. These critical divisors look like $p_k/f(m)$ for some function of the whole numbers m. If we instead compute $\frac{p_k/p}{f(m)}$ those numbers will be in exactly the same order and give the same result; so the output depends only on the population distribution. □

There's another obvious sense of fairness: smaller states should definitely not get *more* seats.

Definition 2.38. An apportionment method is *order-preserving* if, whenever $a_i > a_j$, then $p_i > p_j$.

Note this just says that a state with more seats must have a larger population. The opposite is not true; a state with a larger population may not have more seats. It can't have *fewer* seats, by this rule, but it's possible for two states with different populations to have the same number of seats. (Indeed, that's very hard to avoid!)

We *could* call this property "population monotone", but we're going to save that for the stronger idea first stated in definition 2.9.

We can also recall the ideas from definitions 2.19 and 2.20:

Definition. We say it's a *quota violation* if an apportionment method gives a state more representatives than its upper quota, or less than its lower quota.

An apportionment method satisfies the *quota rule* if it assigns every state either its lower quota or its upper quota.

An apportionment method satisfies the *upper quota rule* if it never assigns a state more than its upper quota.

A violation of this rules is an *upper quota violation*.

Hamilton's method obviously satisfies the quota rule. We saw that Jefferson's method violates the upper quota rule, and similarly Adams's method violates the lower quota rule. We'll see in section 2.3.5 that *any* divisor method has to have quota violations.

2.3.1 House Monotonicity

Recall definition 2.7:

Definition. An apportionment method is called *house monotone* if an increase in h, while all other parameters remain the same, can never cause any seat allocation a_k to decrease. Thus example 2.6 shows that Hamilton's method is not house monotone.

The Alabama paradox in 1880 is a violation of house monotonicity, and shows that Hamilton's method is not house monotone. Had h increased from 299 to 300, Alabama's apportionment would have dropped from 8 seats to 7. See also example 2.6 where we saw an example of Hamilton's method failing house monotonicity.

Divisor methods do better. We saw in proposition 2.23 that Jefferson's method is house monotone. That's just a special case of the following result:

Proposition 2.39. *All divisor methods are house monotone.*

Proof. Consider any divisor method applied to states with populations $p_1, p_2, \ldots, p_n$, and suppose a divisor d gives an apportionment of h seats.

If we want to increase h, we will need to decrease d, which will increase the modified quota p_k/d for each k. Rounding a larger number can never give a smaller result (by definition 2.30). So no state can ever get a smaller apportionment from a larger h. □

2.3.2 Population Monotonicity

Recall definition 2.9:

Definition. A method is called *population monotone* if a state can never lose a seat when its population increases while no other state's population increases.

In algebraic terms, whenever $a_i' < a_i$ and $a_j' > a_j$, it must be the case either that $p_i' < p_i$ or $p_j' > p_j$.

We know that Hamilton's method doesn't satisfy population monotonicity, from example 2.8. However, divisor methods do better here:

Proposition 2.40. *All divisor methods satisfy population monotonicity.*

Proof. Suppose $a_i' < a_i$ and $a_j' > a_j$. By definition 2.30 of rounding functions, this must mean that the modified divisor p_i/d of state i decreased, and the modified divisor of p_j/d increased. That gives us the two inequalities

$$p_i'/d' < p_i/d \qquad\qquad p_j'/d' > p_j/d.$$

We can rearrange these inequalities to get

$$p_i' < p_i \frac{d'}{d} \qquad\qquad p_j' > p_j \frac{d'}{d}.$$

Now we think about two possibilities. First, let's suppose that $d' \leq d$ and so $d'/d \leq 1$. That implies that

$$p_i' < p_i \frac{d'}{d} \leq p_i$$

and thus $p_i' < p_i$, which satisfies the definition of population monotonicity.

If that's not true, then we must have $d' > d$ and so $d'/d > 1$. That would imply that

$$p_j' > p_j \frac{d'}{d} > p_j$$

and thus $p_j' > p_j$, also satisfying the definition of population monotonicity.

□

Corollary 2.41. *Hamilton's method is not a divisor method.*

This corollary is more interesting that it seems. Obviously we didn't *define* Hamilton's method as a divisor method. But conceivably there could be some weird rounding function that would always give the same result as the Hamilton algorithm. But in fact that's not the case; any divisor method will be population monotone, and Hamilton's method is not, so there's no possible rounding function that will give those results.

We can leverage our ideas here to make proposition 2.39 kind of superfluous..

Proposition 2.42. *Any method that is population monotone is also house monotone.*

Proof. Suppose a method is population monotone, and consider a situation in which our house size changes from h to $h' = h + 1$. And suppose that no populations change, which we can write as $p_i' = p_i$ and $p_j' = p_j$.

Increasing the house size means that at least one state will gain a seat. So assume (without loss of generality) that state j gains a seat, and thus $a_j' > a_j$.

Now imagine that some state loses a seat, meaning that $a_i' < a_i$. By population monotonicity, we must have either $p_i' < p_i$, or $p_j' > p_j$. But we know that $p_i' = p_i$ and $p_j' = p_j$, so that's not possible.

Therefore no state can lose a seat, and so the method is house monotone. □

An interesting note is that this proof works because the definition of population monotonicity doesn't have an all-else-equal clause; the property holds even if h changes. All of that work is loaded into the premise that one states gains seats and another loses seats.

Proposition 2.43. *Any method that is population monotone and neutral must be order-preserving.*

Proof. Suppose a method is neutral and population monotone, and suppose we get a census with $p_j > p_i$. We can construct another census by swapping the populations of states i and j—that is, we write $p'_i = p_j$ and $p'_j = p_i$, and then $p'_k = p_k$ for all other states k.

This results in a (hypothetical) increase in population for state i, and a decrease in population for state j. Since $p'_i \geq p_i$ and $p'_j \leq p_j$, it follows that either $a'_i \geq a_i$, or $a'_j \leq a_j$. (Or both!)

But by neutrality, swapping the populations should swap the apportionments, so $a'_i = a_j$ and $a'_j = a_i$. Thus we can make at least one of the following two arguments:

$$a_j = a'_i \geq a_i$$
$$a_j \geq a'_j = a_i.$$

In either case, we see that $a_j \geq a_i$.

Thus if $p_j > p_i$, we must have $a_j \geq a_i$, which is the definition of order preserving. □

2.3.3 Relative population monotonicity

There are actually two ways of thinking about population change. If a state goes from an original population p_k to a later population p'_k, the most obvious way to measure the change is with the *absolute change* $\Delta p_k = p'_k - p_k$. We can also think of this as the arithmetic population change. It's the amount by which the population has grown.

Remark 2.44. Δ is the Greek letter capital Delta, which is a capital D. It's often to refer to a change in a quantity over time; you can think of it as standing for "distance". You may recognize it from algebra, where we sometimes write the slope of a line as $\frac{\Delta y}{\Delta x}$, the change in y divided by the change in x.

But this isn't actually how we talk about population growth most of the time. Generally we talk about growth in percentage terms. This makes sense, because gaining 10,000 people makes a much bigger difference for a state with 500,000 people than to a state with 20,000,000 people.

Instead, we often want to talk about the *relative change*

$$\frac{\Delta p_k}{p_k} = \frac{p'_k - p_k}{p_k}.$$

This measures how big the growth is relative to the starting population.

k	1	2	3
p_k	10,000	10,000	100,000
p'_k	11,000	20,000	110,000
Δp_k	1,000	10,000	10,000
$\Delta p_k/p_k$	0.1	1.0	0.1
%	10%	100%	10%

Example 2.45. Suppose we have the following table of populations:

State 2 and State 3 have the same absolute population growth; each gains 10,000 people between the two censuses. But that change is much more dramatic for state 2 than state 3, because state 2 started with many fewer people. State 1 and state 3 have the same *relative* population growth, of 10%; state 2, rather, grows 100% between the two censuses.

Definition 2.46. An apportionment method is *relative population monotone* if, whenever we consider states with positive population and $a'_i < a_i$ and $a'_j > a_j$, then $\frac{\Delta p_j}{p_j} > \frac{\Delta p_i}{p_i}$.

This is essentially the same as definition 2.9 that we studied in section 2.3.2, but instead it focuses on relative changes. This makes it better at comparing states of very different populations. It also makes the criterion simpler, because we don't need to worry about whether the states are growing or shrinking; we can just make one statement.

We need to consider the case where a state has zero population because we want to talk about introducing new states from one census to another; we can think of this as the state having zero population in the first census. But we can't talk about relative population growth from a baseline of zero population, since that would involve dividing by zero.

This property is stronger than regular population monotonicity. In particular, if a method is relative population monotone, then it's population monotone.

Proposition 2.47. *If an apportionment method is relative population monotone, then it is population monotone.*

Proof. Suppose $a'_i < a_i$ and $a'_j > a_j$, and that all populations are positive. By relative population monotonicity, we conclude that $\frac{\Delta p_j}{p_j} > \frac{\Delta p_i}{p_i}$. We can conclude, in particular, that either $\Delta p_j/p_j$ is positive, or $\Delta p_i/p_i$ is negative. (Or both!)

If $\Delta p_j/p_j$ is positive, that means that Δp_j is positive, and thus $p'_j < p_j$. If $\Delta p_i/p_i$ is negative, that means that Δp_i is negative, and thus $p'_i < p_i$. And that's the definition of population monotonicity. □

But the property isn't *too* strong. We can achieve it with reasonable methods.

Proposition 2.48. *All divisor methods are relative population monotone.*

Proof. Suppose we care computing apportionments of two censuses using a divisor method, and suppose that $a'_i < a_i$ and $a'_j > a_j$. We know that the modified quota p_i/d must have decreased and the modified quota of p_j/d must have increased. Another way of writing that is:

$$\frac{p'_i}{d'} < \frac{p_i}{d} \qquad \frac{p'_j}{d'} > \frac{p_j}{d}$$

$$\text{thus} \qquad \frac{p'_i}{p_i} < \frac{d'}{d} \qquad \frac{p'_j}{p_j} > \frac{d'}{d}$$

We can link these two inequalities together, since they both have the d'/d term.

$$\begin{aligned} \frac{p'_i}{p_i} &< \frac{p'_j}{p_j} \\ \frac{p'_i}{p_i} - 1 &< \frac{p'_j}{p_j} - 1 \\ \frac{p'_i}{p_i} - \frac{p_i}{p_i} &< \frac{p'_j}{p_j} - \frac{p_j}{p_j} \\ \frac{p'_i - p_i}{p_i} &< \frac{p'_j - p_j}{p_j}. \end{aligned}$$

And that's just the statement that $\frac{\Delta p_i}{p_i} < \frac{\Delta p_j}{p_j}$, which is what we needed to prove.

□

And in fact relative population monotone isn't all that much stronger than population monotone at all. In particular, if we assume proportionality, they're the same.

Proposition 2.49. *If an apportionment method is proportional and population monotone, then it's relative population monotone.*

Proof. Suppose an apportionment method is proportional, but not relative population monotone. That means we have some pair of states i and j where $a'_i < a_i$ and $a'_j > a_j$, but $\frac{\Delta p_j}{p_j} \leq \frac{\Delta p_i}{p_i}$.

For notational convenience we set

$$r = 1 + \frac{\Delta p_i}{p_i} \qquad s = 1 + \frac{\Delta p_j}{p_j}$$

and observe that $0 < s$ since the relative growth rate can't be less than -100%, and $s \leq r$ by our hypothesis.

Now we compute

$$rp_i = \left(1 + \frac{\Delta p_i}{p_i}\right) = p_i + \Delta p_i = p_i + \left(p_i' - p_i\right) = p_i'$$
$$sp_j = \left(1 + \frac{\Delta p_j}{p_j}\right) = p_j + \Delta p_j = p_j + \left(p_j' - p_j\right) = p_j'$$

Now suppose we have a third census (which we'll denote with double-prime letters, like p''), where for every state k we set $p_k'' = p_k'/r$. This is a pure rescaling of the second census, where every state's population is scaled by a factor of r. By proportionality, this can't change the apportionment, so we have $a_k'' = a_k'$ for every state k. But we can compute

$$p_i'' = p_i'/r = p_i$$
$$p_j'' = p_j'/r = \frac{s}{r}p_j' \leq p_j.$$

If we compare our original census to this third census, we must have $a_i'' = a_i' < a_i$ and $a_j'' = a_j' > a_j$; but we have $p_i'' = p_i$ and $p_j'' \leq p_j$, meaning neither $p_i'' < p_i$ nor $p_j'' > p_j$. So this apportionment system is also not population monotone.

□

Since we're generally only going to consider proportional methods, we can use population monotonicity and relative population monotonicity interchangeably; in the future we won't really specify.

2.3.4 The New States Paradox

An interesting specific situation is when a new state is introduced entirely to the union; this has happened many times in the history of the US (most recently with Hawai'i in 1959). We can interpret this as a situation where a state's population increases from 0 to a positive number. (Even the Adams method doesn't assign a seat to a state with a population of zero!)

Definition 2.50. Suppose state k is joining the union as a new state, and thus $p_k = 0$ and $p_k' > 0$. Suppose there are other states i and j whose populations are unchanged. We say a *new states paradox* or *Oklahoma paradox* occurs if $a_i' < a_i$ and $a_j' > a_j$.

Remark 2.51. There's nothing weird about $a_i' < a_i$, or $a_j' > a_j$, in isolation. Maybe adding the new state means some old states have to lose representation; that will certainly happen if we don't increase h. And if we do increase h to accommodate the new state, then maybe other states will get increased representation too.

But if both happen at once, that means that adding the new state has caused other states to transfer seats among themselves; that's the odd bit.

We don't have a ton of results to prove here, though. If an apportionment method is population monotone, then in particular it can't suffer from the new states paradox. Thus by proposition 2.40, no divisor method can suffer from the new states paradox. On the other hand, example 2.10 shows that Hamilton's method is vulnerable to the new states paradox. (This also furnishes another proof that Hamilton's method is not a divisor method.)

2.3.5 The Impossibility Theorem

Let's return to one of our earliest criteria, the quota criterion. Hamilton's method satisfies the quota rule; we've seen that Jefferson's method does not. (We will see that no divisor method can satisfy the quota rule, in fact.) There are other quota methods, such as Lowndes's method:

Definition 2.52 (Lowndes's method)**.** As a provisional apportionment, assign every state its lower quota. Then assign the remaining seats to the states, at most one per state, in decreasing order of $\frac{\{q_k\}}{\lfloor q_k \rfloor}$, the ratio of the fractional part of the standard quota to the lower quota.

This has the same basic approach as Hamilton's method: we give every state its lower quota, then use the fractional part of the standard quota to decide which states to round up to the upper quota. But it gives an extra bonus to smaller states, since $\frac{\{q_k\}}{\lfloor q_k \rfloor}$ will be bigger when $\lfloor q_k \rfloor$ is smaller.

The underlying logic is that the extra representative gives more "extra representation" in the smaller state; cutting a seat from a smaller state will create bigger districts than cutting a seat from a larger state will.

For example, if we have one state with quota $q_i = 2.3$ and another state with quota $q_j = 7.9$, then Hamilton's method would give an extra representative to state j rather than state i. But if we deny state i its extra district, that makes each district $\frac{.3}{2} = 15\%$ bigger than the desired standard divisor; if we instead deny state j, that makes each district $\frac{.9}{7} \approx 12.86\%$ bigger than the desired size. So Lowndes argues that it's better to give the extra district to state i.

Exercise 2.53. *Lowndes's method is not house monotone or population monotone, and is vulnerable to the new states paradox.*

So far we've seen quota methods, which fail various monotonicity criteria; and we've seen divisor methods, which we haven't shown can avoid quota violations. (Adams's method avoids upper quota violations, and Jefferson's method avoids lower quota violations, but that's more because they have strong biases in the other direction than because they particularly respect quotas.)

In fact, there's no way to completely solve both these problems at once.

Theorem 2.54 (Balinski and Young)**.** *No apportionment rule that is neutral and population monotone can satisfy the quota rule.*

Proof. We can prove this by constructing a specific example where there's no neutral way to apportion seats that satisfies both population monotonicity and the quota rule. So consider a pair of censuses where we want to allocate $h = 10$ seats:

$$\begin{aligned} p_1 &= 69{,}000 & p_1' &= 68{,}000 \\ p_2 &= 5{,}200 & p_2' &= 5{,}500 \\ p_3 &= 5{,}000 & p_3' &= 5{,}600 \\ p_4 &= 19{,}900 & p_4' &= 5{,}700. \end{aligned}$$

In the "before" situation, we have a total population of 100,000, so the standard divisor is $s = 10{,}000$ and the standard quotas are

$$q_1 = 6.99 \qquad q_2 = 0.52 \qquad q_3 = 0.50 \qquad q_4 = 1.99.$$

Since the method satisfies the quota rule, we know that state 1 has 7 seats or less, and state 4 has 2 seats or less, so states 2 and 3 have to get at least one seat between them. Because this method is population monotone and neutral, proposition 2.43 shows it must be order-preserving; so state 2 must receive at least one seat.

Now let's consider the after census. The total population is 84,800 so the standard divisor is 8,480. We can compute the standard quotas for this after census, and we get

$$q_1' = 8.02 \qquad q_2' = 0.65 \qquad q_3' = 0.66 \qquad q_4' = 0.67.$$

Again by the quota rule, state 1 has to get at least 8 seats, leaving at most two seats for the other three states. By the order-preserving property, we can't give a seat to state 2 if states 3 and 4 don't have one each, and since there are only two seats to go around, so state 2 cannot get any seats at all.

So if we have a neutral, population monotone method that satisfies the quota rule, it must give state 1 at most 7 seats in the before situation, and at least 8 in the after situation; it must give state 2 at least one seat in the before situation, and cannot give state 2 any in the after situation. Thus state 1 will gain seats, and state 2 will lose seats.

But the population of state 1 had declined while the population of state 2 has increased. That's a violation of population monotonicity, so we cannot have such a method.

□

Corollary 2.55. *No divisor method satisfies the quota rule.*

This shows that even though both the quota rule and population monotonicity are intuitively appealing, but we cannot have them both; we have to choose. (In the US today we use Hill's method, which is a divisor method and thus violates the quota rule.)

But theorem 2.54 doesn't mention *house* monotonicity. And it turns out it is in fact possible to get a method that satisfies the quota rule and is still house monotone.

2.4 Balinski and Young Apportionment

Jefferson's method is very prone to quota violations, but it's straightforward to calculate and seems like it has a fair amount of constitutional support. Can we tweak it to also satisfy the quota rule?

Theorem 2.54 says we can't get everything we want. If it satisfies the quota rule, it can't possibly be population monotone. But it turns out we can maintain the house monotonicity and still keep the quota rule. And we can do this fairly straightforwardly, by just instituting a rule that we never allow an upper quota violation.

We will define this apportionment method iteratively, or inductively. We're going to assign seats one at a time, in order; so that the way we assign 10 seats is to assign the first 9 seats, and then assign one more.

Recall we can look at Jefferson's method itself that way. We start for $h = 0$, in which case we obviously apportion 0 seats to each state, $a_k = 0$. (In formal "mathematical induction" we call this the base case.)

If we've already apportioned h seats, and have $a_1, a_2, \ldots, a_n$, we can always apportion the next seat. We compute the critical divisors $\frac{p_k}{a_k+1}$, which represent how large a Congressional district in state k would be if we assign it one more seat; the larger this number is, the less overrepresented state k would be if we give them another seat. Thus the $h + 1$st seat goes to the state with the largest critical divisor. We saw this worked out in example 2.16.

This method is straightforward, and obviously house monotone. But it is unfortunately prone to upper quota violations; in most reasonable situations large states will get more than their upper quotas. For instance, with 2020 census data, California has a standard quota of 51.99, giving it an upper quota of 52 and a lower quota of 51. Jefferson's method would allocate it 54 seats. Jefferson's method would also give New York and Texas one seat more than their upper quotas; it would give Vermont and Wyoming no seats at all.

Definition 2.56 (Balinksi and Young method)**.** We define the method of Balinski and Young inductively.

If $h = 0$, then set $a_k = 0$ for every k.

Suppose we have an apportionment for some fixed h, given by $a_1, a_2, \dots, a_n$ such that $a_1 + \cdots + a_n = h$.

For each state k, compute the quotient $\frac{p_k}{a_k+1}$ and call this the strength of the kth state's claim for the next seat. We "want" to give the next seat to the state with the strongest claim, but we don't want to have any upper quota violations.

So we say a state is *eligible* if $a_k + 1 \leq \left\lceil (h+1)\frac{p_k}{p} \right\rceil$, so that giving the state another seat would not give an upper quota violation.

Then we assign the $h+1$st seat to the eligible state with the strongest claim.

Poll Question 2.4.1*.* Why do we use Jefferson rather than Adams or Hill or Webster as the base for this method?

This method will sort of "obviously" avoid upper quota violations, since we simply refuse to allocate a seat when it would cause an upper quota violation. But that maybe leads to a question of what happens if no state is eligible. Fortunately that can't happen.

Proposition 2.57. *At each inductive stage of the method of Balinski and Young, at least one state is eligible to receive the next seat.*

Proof. The basic idea of this proof is that for a state to be ineligible, we have to have given it a lot of seats, relative to the total number of seats h. But we can't give every state a lot of seats relative to h, because h is the total number of seats we can allocate. So at least one state must have space.

Consider a census with populations $p_1, p_2, \dots, p_n$, and suppose that Balinski and Young have apportioned $a_1, a_2, \dots, a_n$ seats for a total of h seats to the various states.

After we apportion the $h+1$st state, the standard divisor will be $s = \frac{p}{h+1}$, so the standard quota for each state will be $\frac{p_k}{s} = (h+1)\frac{p_k}{p}$. A state will be ineligible if giving it one more seat will put it over this upper quota; thus it's ineligible if $a_k + 1 > \left\lceil (h+1)\frac{p_k}{p} \right\rceil$.

But since those are both whole numbers, the only way we can have $a_k + 1$ *bigger* than the upper quota is if a_k is *at least as big* as the upper quota. (If we have fractions that wouldn't be true; we can have $3 < 3.5$ but $3 + 1 > 3.5$. But for whole numbers, if $a_k + 1$ is bigger than some whole number than a_k must be at least the same size.) So if a state is ineligible, we must have

$$a_k \geq \left\lceil (h+1)\frac{p_k}{p} \right\rceil \geq (h+1)\frac{p_k}{p}.$$

Now if *every* state is ineligible, we have the following series of inequalities:

$$\begin{aligned} a_1 &\geq (h+1)\frac{p_1}{p} \\ a_2 &\geq (h+1)\frac{p_2}{p} \\ &\vdots \\ a_n &\geq (h+1)\frac{p_n}{p}. \end{aligned}$$

If we add all these inequalities, we get

$$\begin{aligned} a_1 + \cdots + a_n &\geq (h+1)\frac{p_1}{p} + (h+1)\frac{p_2}{p} + \cdots + (h+1)\frac{p_n}{p} \\ &= (h+1)\left(p_1 + p_2 + \cdots + p_n\right)\frac{1}{p} \\ &= (h+1)(p)\frac{1}{p} = h+1. \end{aligned}$$

That is, once we add up all these allocations, we see that $a_1 + \cdots + a_n \geq h + 1$, we have to have already allocated at least $h + 1$ seats. But we're at the step where we have allocated exactly h seats, so that can't be true. So that means at least one state must have

$$a_n < (h+1)\frac{p_n}{p}$$

and be eligible to receive the next seat.

□

Let's work through an example comparing Jefferson's method to Balinski and Young's method. Suppose we have states A, B, C with populations $p_1 = 7, p_2 = 22, p_3 = 71$, for a total $p = 100$. Suppose we want to attain a house size of $h = 15$. We can think of Jefferson's method as working iteratively like this, as seen in figure 2.14

Now let's see what this would look like in the method of Balinski and Young. The basic approach will be the same, but at each step we remove ineligibles, and avoid awarding seats to any state that isn't eligible for another seat. We see this process in figure 2.15

h	Jefferson Critical Divisor			Jefferson Apportionment		
	$\frac{p_1}{a_1+1}$	$\frac{p_2}{a_2+1}$	$\frac{p_3}{a_3+1}$	a_1	a_2	a_3
0				0	0	0
1	7	22	**71**	0	0	1
2	7	22	**35.5**	0	0	2
3	7	22	**23.67**	0	0	3
4	7	**22**	17.75	0	1	3
5	7	11	**17.75**	0	1	4
6	7	11	**14.2**	0	1	5
7	7	11	**11.83**	0	1	6
8	7	**11**	10.14	0	2	6
9	7	7.33	**10.14**	0	2	7
10	7	7.33	**8.875**	0	2	8
11	7	7.33	**7.89**	0	2	9
12	7	**7.33**	7.1	0	3	9
13	7	5.5	**7.1**	0	3	10
14	**7**	5.5	6.45	1	3	10
15	3.5	5.5	**6.45**	1	3	11

Figure 2.14: Jefferson's method worked iteratively

2.4.1 The Quota Rule for Balinski and Young

It's clear that the Balinski-Young method is house monotone, because of the way we produce it iteratively: the seats are given away one at a time, and we never back up and change an earlier apportionment. It's also clear that it satisfies the upper quota rule, because we simply never assign a seat that would violate the upper quota rule.

The trick is to show that it also satisfies the lower quota rule. The Jefferson method, of course, satisfies the lower quota rule, but that's not realy for any deep reason. We need to check that the Balinksi and Young method doesn't ever accidentally ruin that.

Proposition 2.58. *The Balinski and Young method satisfies the lower quota rule.*

h	Jefferson Critical Divisor			Jefferson Apportionment			Standard Quotas			Balinski and Young Apportionment		
	$\frac{p_1}{a_1+1}$	$\frac{p_2}{a_2+1}$	$\frac{p_3}{a_3+1}$	a_1	a_2	a_3	q_1	q_2	q_3	a_1	a_2	a_3
0				0	0	0						
1	7	22	**71**	0	0	1	0.07	0.22	0.71	0	0	1
2	7	22	**35.5**	0	0	2	0.14	0.44	1.42	0	0	2
3	7	22	**23.67**	0	0	3	0.21	0.66	2.13	0	0	3
4	7	**22**	7.75	0	1	3	0.28	0.88	.84	0	1	3
5	7	11	**17.75**	0	1	4	0.35	1.1	3.55	0	1	4
6	7	11	**14.2**	0	1	5	0.42	1.32	4.26	0	1	5
7	7	**11**	~~11.83~~	0	1	6	0.49	1.54	~~4.97~~	0	2	5
8	7	7.33	**11.83**	0	2	6	0.56	1.76	5.68	0	2	6
9	7	7.33	**10.14**	0	2	7	0.63	1.98	6.39	0	2	7
10	7	7.33	**8.88**	0	2	8	0.7	2.2	7.1	0	2	8
11	7	**7.33**	~~7.89~~	0	2	9	0.77	2.42	~~7.81~~	0	3	8
12	7	5.5	**7.89**	0	3	9	0.84	2.64	8.52	0	3	9
13	7	5.5	**7.1**	0	3	10	0.91	2.86	9.23	0	3	10
14	**7**	5.5	~~6.45~~	1	3	10	0.98	3.08	~~9.94~~	1	3	10
15	3.5	5.5	**6.45**	1	3	11	1.05	3.3	10.65	1	3	11

Figure 2.15: Apportionment by the Balinski and Young method

Proof. □

Corollary 2.59. *The Balinski and Young method is not population monotone.*

Proof. Since the Balinski and Young method satisfies the quota rule, Theorem 2.54 (proven by Balinski and Young!) shows it can't be population monotone. □

2.5 Why choose different rounding functions?

In sections 2.2.4 and 2.2.5 we covered a collection of rounding methods, which all give almost, but not quite, the same results. What's the point of having all that variety?

We talked about this briefly at the time, but each rounding method is "the best" at a slightly different sort of thing. We can view each rounding function as finding the optimal solution to a specific question; they differ in what question they answer.

2.5.1 Degree of Representation and Webster's Method

Definition 2.60. The *degree of representation* of a state k is the number $\frac{a_k}{p_k}$, which measures the fraction of a congressional seat each individual citizen is allocated.

In the 2020 census, Maryland had 6,185,278 people and got allocated 8 congressional seats. This means that each citizen of Maryland is represented by $\frac{8}{6{,}185{,}278} \approx .000{,}001{,}293$ of a Congressperson.

Because every state gets a whole number of representatives, we can't give every state exactly the same degree of representation. But Webster's method, the divisor method with arithmetic rounding does the best it can.

Proposition 2.61. *Webster's method is the unique apportionment method with the following property: the difference between the degrees of representation of any two states cannot be decreased by transferring a seat from the better represented state to the worse represented state.*

2.5.2 District Size and Dean's Method

We also can't make it so that districts have the same number of people in each state. But Dean's method, the divisor method with harmonic rounding, does the best it can.

Proposition 2.62. *Dean's method is the unique apportionment method with the following property: the difference between the sizes of districts in any two states cannot be made smaller by transferring a seat from the smaller-district state to the bigger-district state.*

2.5.3 Average district size and Hill's Method

There is more that one way to compare how close two numbers are. One is to subtract them from each other and see how big the result is—or in other words, how far it is from 0. But this doesn't work well when our numbers have radically different scales. It might not really capture what we care about to say that 1 is closer to 1,000,000 than 1,000,000 is to 2,000,000.

Another way of measuring the distance between two numbers is to divide one by the other, and see how close we are to 1. By this method, 100 and 200 are closer than 10 and 30 are, and 1,000,000 is much closer to 2,000,000 than it is to 1.

It makes sense to say that we want the districts in different states to be as close in size as possible. If we measure closeness by subtraction—arithmetically—then Dean's method does this best. But if we measure closeness geometrically, by multiplication, then Hill's method, the divisor method with geometric rounding, does the best.

It also makes sense to say we want the states to have degrees of representation as close as possible. Once again, we saw that Webster's method does that, if we measure closeness arithmetically. But if we measure closeness geometrically, once again Hill's method does the best we can.

Proposition 2.63. *Hill's method is the unique apportionment method with the following property: the ratio between the average sizes of districts in any two states (expressed as number greater than 1) cannot be made smaller by transferring a seat from the smaller-district state to the bigger-district state.*

Further, it is also the unique apportionment method with the property: the ratio between the degrees of representation in any two states cannot be made smaller by transferring a seat from the better-represented state to the worse-represented state.

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 ith row and the jth column will be written $u_{i,j}$ or sometimes u_{ij}. We call this matrix the *payoff matrix*.

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

Example 3.4. In World War II, he 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	Japan: South		
United States	North	2	2	2	←
United States	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.)

	North	South
North	2	2
South	1	3

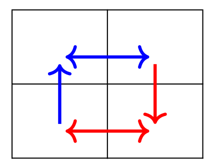

Figure 3.4: The flow diagram for the Battle of the Bismarck Sea

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

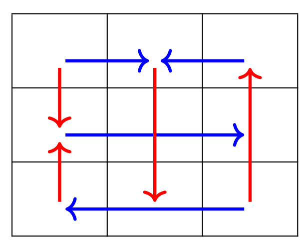

Figure 3.5: The flow diagram for Rock Paper Scissors

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

$$u_{k,\ell} \leq u_{k,j} \qquad \text{for any column } j$$
$$u_{k,\ell} \geq u_{i,\ell} \qquad \text{for any row } i.$$

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 and 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. □

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. □

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. □

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:

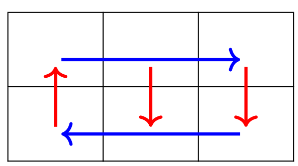

Figure 3.7: The flow diagram for example 3.29

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 $p_1, p_2, \ldots, p_n$. 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 $p_1 + p_2 + \cdots + p_n = 1$; 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 $P = (p_1, p_2, \ldots, p_n)$, 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 $P = (1/n, 1/n, \ldots, 1/n)$.

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 $P = (p_1, p_2, \ldots, p_n)$. 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 $x_1, x_2, \ldots, x_n$.

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 $P = (p_1, p_2, \ldots, p_n)$. 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 $P = (p_1, \ldots, p_m)$ 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 $Q = (q_1, \ldots, q_n)$ on their set of n pure strategies.

Example 3.39. Suppose Column plays Rock Paper Scissorswith 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 $Q = (q_1, \ldots, q_n)$, then the expected value of the payoff is*

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

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

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

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 $P = (p_1, \ldots, p_m)$ and the other by a probability distribution $Q = (q_1, \ldots, q_n)$. 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 ScissorsRow 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 $P = (p_1, \ldots, p_m)$ and Column plays $Q = (q_1, \ldots, q_n)$. 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

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

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 $P = (p_1, \ldots, p_m)$ and Column plays $Q = (q_1, \ldots, q_n)$. Then the expected value of the payoff can be computed either by*

$$E(P,Q) = p_1E(P_1, Q) + \cdots + p_mE(P_m, Q)$$
$$E(P,Q) = q_1E(P,Q_1) + \cdots + q_nE(P, Q_n).$$

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 $E(P,Q)$. □

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:

$$\begin{aligned} E(P_1, Q) &= \frac{2}{7} \cdot (0) + \frac{4}{7} \cdot (-1) + \frac{1}{7} \cdot (1) = -3/7 \\ E(P_2, Q) &= \frac{2}{7} \cdot (1) + \frac{4}{7} \cdot (0) + \frac{1}{7} \cdot (-1) = 1/7 \\ E(P_3, Q) &= \frac{2}{7} \cdot (-1) + \frac{4}{7} \cdot (1) + \frac{1}{7} \cdot (0) = 2/7 \end{aligned}$$

(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 $P = (p_1, \ldots, p_m)$, then

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

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 Scissorsone 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. □

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 $\overline{r} = 0$ and $\overline{c} = 0$.

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* $\overline{r} = \overline{c}$*. 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 $\overline{r} = \overline{c} = E(P,Q)$. Thus both P and Q are prudent.

Conversely, suppose (P,Q) is doubly prudent, and thus $E(P,Q) = \overline{r} = \overline{c}$. Since P is prudent, every payoff against P must be at least the optimal guarantee $\overline{r}$. 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 $\overline{c}$. 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. □

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:

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

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:

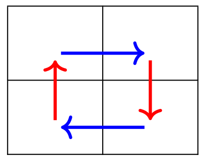

In particular, note that we can't have $a = b$ because otherwise there would definitely be a saddle point somewhere. Thus $a > c$. We also can see from the diagram that $d > b$, that $b < a$, and that $c < d$.

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

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

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

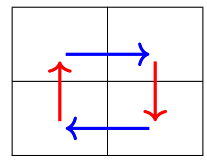

Compute:

$$\begin{aligned}
p &= \frac{a-b}{(a-b)+(d-c)} \\
&= \frac{60-10}{60-10+40-0} = \frac{50}{90} = \frac{5}{9} \\
1-p &= \frac{9}{9} - \frac{5}{9} = \frac{4}{9} \\
q &= \frac{a-c}{(a-c)+(d-b)} \\
&= \frac{60-0}{60-0+40-10} = \frac{60}{90} = \frac{2}{3} \\
1-q &= 1 - \frac{2}{3} = \frac{1}{3} \\
v &= \frac{ad-bc}{(a-b)+(d-c)} \\
&= \frac{(60)(40)-(10)(0)}{60-10+40-0} = \frac{2400-0}{90} = \frac{240}{9} = \frac{80}{3} = 26^2/_3.
\end{aligned}$$

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 $26^2/_3$% 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

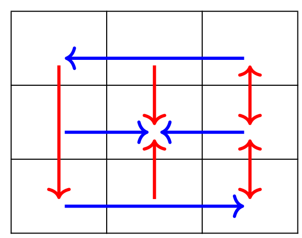

In this case the horizontal arrows are looking for the largest *second* number in each row, while the vertical arrows are looking for the largest *first* number. Remember the Column only cares about second numbers, while Row only cares about first numbers.

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

References

[Price(2022)] Anna Price. When a bee is a fish in the eyes of the law. *Library of Congress Blogs*, 2022. URL https://blogs.loc.gov/law/2022/06/when-a-bee-is-a-fish-in-the-eyes-of-the-law/.