Chapter 5

Explaining the Impossible

No one pretends that democracy is perfect or all-wise. Indeed it has been said that democracy is the worst form of government except for all those other forms that have been tried from time to time.

– Winston Churchill

Focus Questions

In this chapter, we'll explore the following questions:

- What is the basic idea behind the proof of Arrow's Theorem?
- Can Pareto's unanimity condition be weakened to solve the problems revealed by Arrow's Theorem?
- What is approval voting? Does it solve any of the problems revealed by Arrow's Theorem?
- What is the intensity of binary independence criterion? How is it related to Arrow's Theorem?

Warmup 5.1. Consider the following mathematical claim:

> *It is impossible for a whole number to be divisible by 2, 11, and 23, and not be greater than 500.*

Is this claim true or false? Give a convincing argument or example to justify your answer.

Suppose we wanted to prove that the claim from Warmup 5.1 is true. How could we do it? One method would be to simply check all of the whole numbers one by one and verify that none of them are divisible by 2, 11, and 23, and not greater than 500. However, this would take quite a long time, wouldn't it? Actually, since there are infinitely many whole numbers, the truth is we'd never be able to check them all.

Of course, we could reduce our work quite a bit if we only considered the whole numbers that are not greater than 500. Then we'd just have to show that none of these numbers are divisible by 2, 11, and 23. But even this seems like an awful lot of work.

Fortunately, there's a much better way to prove that the claim from Warmup 5.1 is true. What if, instead of considering numbers one by one, we constructed some sort of logical argument to establish the truth of the claim? For instance, we might say something like this:

> *The numbers 2, 11, and 23 are all prime numbers. Thus, any whole number divisible by 2, 11, and 23 must be at least as big as $2 \times 11 \times 23 = 506$. So, it is impossible for a whole number to be divisible by 2, 11, and 23, and not be greater than 500.*

At this point, you might be wondering what all of this has to do with voting. Well, as it turns out, the same strategy we just used to prove the claim from Warmup 5.1 can also be used to prove Arrow's Theorem—and we'll do so in this chapter. Our goal will be to see *why* Arrow's Theorem is true and to consider some potential options for resolving the problems that Arrow first brought to light.

Proving Arrow's Theorem

We'll begin by walking step-by-step through a proof of Arrow's Theorem.[1] First, you should be advised that Arrow's Theorem is a significant result, and proving it will require a fair amount of effort and concentration. Nevertheless, we'll be able to tackle and understand the proof if we just take it one step at a time. When we're done, you will have joined a select group of people who know not only the meaning of one of the most important theorems in social choice theory, but also why the theorem is true.

Before we take off into the proof, we first need to mention one more piece of notation that we will be using along the way. Recall that we have used the symbol $\succ$ to represent a preference between two candidates in an election and $\approx$ to represent a tie. Sometimes in this chapter we will have reason to say that a candidate A is either preferred over *or* tied with another candidate B. We will represent this type of relation by writing $A \succsim B$. Note again the analogy to a common symbol ($\geq$) that we use to compare numbers.

Now on to Arrow's Theorem. As strange as it may seem, it is actually easier to prove the strong form of the theorem that we stated later in Chapter 4 (on page 66) than it is to prove the original version that we stated earlier in that chapter (on page 64). So we'll prove the strong form of Arrow's Theorem first, and then look more carefully at why the strong form implies

[1] The strategy we will use to prove Arrow's Theorem was adapted from a paper by Geanakoplos [**22**].

the original form. To make our proof strategy more clear, we'll begin by restating the strong form of the theorem in a slightly different way than how we first stated it on page 66.

Arrow's Theorem (Strong Form). *For an election with more than two candidates, it is impossible for a voting system to satisfy universality, IIA, and unanimity, and not be a dictatorship.*

Question 5.2. Explain why the strong form of Arrow's Theorem as it is stated above is equivalent to how we first stated it on page 66.

If you look carefully, you'll notice that the strong form of Arrow's Theorem, as it is stated above, is similar in style to the mathematical claim we considered in Warmup 5.1. And, just as we thought about doing there, we could try to prove Arrow's Theorem by using a brute-force approach—that is, by checking every possible voting system and verifying that none of them satisfies all four conditions. However, it's hard to imagine how we would go about checking *every* possible voting system. In fact, it's not even clear that we could *identify* all of the possible voting systems, let alone investigate the properties of each one.

A much better approach would be to try to do exactly what we did in our answer to Warmup 5.1. There, we simply assumed that three of the conditions were true (divisible by 2, 11, and 23), and then explained why the fourth condition (not greater than 500) could not also be true. This is where our revised wording comes in particularly handy. It tells us that we can begin our proof by assuming that there is some voting system for an election with more than two candidates that satisfies universality, IIA, and unanimity. To complete the proof, we'll then just need to explain why this voting system must be equivalent to a dictatorship. That is, we'll need to show that there is some voter v in the system such that, for any two candidates A and B, if v prefers A over B, society will also prefer A over B.

Incidentally, it's worth pointing out the similarities between the last step in our proof strategy above and what we did in Chapter 1 when we proved Theorem 1.22. In that theorem, we assumed that we had a voting system that satisfied anonymity, neutrality, and monotonicity, and we needed to show that this voting system was equivalent to a quota system. To do this, we first constructed a process through which we found a potential value for the quota. We then showed that this potential quota actually worked the way that the quota in a quota system is supposed to work.

Our strategy here will be similar. We'll first construct a process through which we will find a potential dictator for our hypothetical voting system. We'll then show that this potential dictator actually *is* a dictator according to our definition.

In order to do all of this, we'll first need to consider a lemma[2] that will help us along the way. Although right now you might not see exactly how we'll use this lemma, be assured that it will play a crucial role in our proof of Arrow's Theorem.

Lemma 5.3. *Assume that, for an election with more than two candidates, a voting system V satisfies universality, IIA, and unanimity. Suppose that B is some candidate in the election, and that every voter ranks B either in first place or last place (without ties) on their individual preference order. Then the societal preference order produced by V must also rank B either first or last—even if, for example, half of the voters rank B first and the other half rank B last.*

Note that, throughout this chapter, when we say that a candidate is ranked first or last in an individual or societal preference order, we rule out the possibility that they are tied with another candidate. You'll want to keep this in mind as you answer the questions that build up our proof of Lemma 5.3 and Arrow's Theorem.

In order to see why Lemma 5.3 is true, let's begin by assuming that every voter in the election does in fact rank B in either first or last place. We'll make no other assumptions about the voters' preferences.

Question 5.4.* Suppose that the societal preference order produced by V does not rank B either first or last. Explain why it must then be the case that, for some other candidates A and C, $A \succsim B$ and $B \succsim C$.

Question 5.5. Given that $A \succsim B$ and $B \succsim C$, what does the transitivity of the societal preference order allow you to conclude about the societal preference between A and C?

Question 5.6.* Now suppose that every voter changes their individual preference order by moving C above A, with no other changes. Will these changes have any effect on the resulting societal preference between A and B or between B and C? Explain. (Hint: Don't forget that every voter ranks B in either first or last place, without ties, on their individual preference order.)

Question 5.7. In Question 5.6, we assumed a change in preferences that resulted in every voter ranking C above A. Given this change, what does unanimity allow you to conclude about the resulting societal preference between A and C?

Question 5.8.* Explain how your answers to Questions 5.5 through 5.7 give rise to a contradiction. What does this contradiction allow you to conclude about the truth of Lemma 5.3? Thoroughly explain your answer.

[2]The word *lemma* means *helping result*. In mathematics, a lemma is typically a result whose main use is to help establish the truth of another more important result.

With the truth of Lemma 5.3 established, let's now proceed to find our potential dictator. From now through Question 5.16, we'll suppose that we have an election with more than two candidates and a voting system V that satisfies universality, IIA, and unanimity (so that Lemma 5.3 applies). For convenience, we'll name the voters in the election $v_1, v_2, v_3, \ldots, v_n$ (where n represents the total number of voters). And, for reasons that will become clear later, we'll begin by considering the special case in which B is ranked last (without ties) by all of the voters in the system, as shown in Table 5.1.

	Voters			
Rank	v_1	v_2	$\cdots$	v_n
First	?	?	$\cdots$	?
$\vdots$	$\vdots$	$\vdots$		$\vdots$
Last	B	B	$\cdots$	B

TABLE 5.1. Candidate B ranked last unanimously

Question 5.9. Based on the fact that B is ranked last by every voter, what can you conclude about the position of B in the societal preference order produced by V? What property allows you to conclude this?

Question 5.10.*

(a) Suppose that *all* of the voters move B from last place to first place in their individual preference orders. How would the resulting societal preference order change, and why would this change occur?

(b) Suppose that only *some* of the voters move B from last place to first place in their individual preference orders. What possible changes could occur in the resulting societal preference order? (Hint: Don't forget about Lemma 5.3!)

(c) Suppose that, one by one, starting with v_1 and proceeding in order, each voter moves B from last place to first place in their individual preference order. Explain why there must be some voter, say v_j, for which this move first causes a corresponding change to occur in the societal preference order.

The voter that we labeled v_j in part (c) of Question 5.10 is a special voter in the following sense: Even if all of the voters before v_j moved B from last place to first place in their individual preference orders (as shown in Table 5.2), there would still be no change to B in the resulting societal preference order. However, as soon as v_j makes the same change (as shown in Table 5.3), suddenly B would move from last to first in the resulting societal preference order. Because of this, we might call v_j a *pivotal* voter.

As it turns out, v_j is also a dictator. We'll establish this fact in two steps: First, we'll show that for any pair of candidates that does not include

	Voters							
Rank	v_1	v_2	$\cdots$	v_{j-1}	v_j	v_{j+1}	$\cdots$	v_n
First	B	B	$\cdots$	B	?	?	$\cdots$	?
$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
Last	?	?	$\cdots$	?	B	B	$\cdots$	B

TABLE 5.2. Society ranks B last

	Voters							
Rank	v_1	v_2	$\cdots$	v_{j-1}	v_j	v_{j+1}	$\cdots$	v_n
First	B	B	$\cdots$	B	B	?	$\cdots$	?
$\vdots$	$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
Last	?	?	$\cdots$	?	?	B	$\cdots$	B

TABLE 5.3. Society ranks B first

B, say A and C, if v_j prefers A over C, then society will also prefer A over C. We'll then show that the same condition holds for any pair of candidates that does include B.

For the first step, let A and C represent any two candidates other than B. In addition, assume that v_j prefers A over C. We now want to be able to conclude that A is preferred over C in the resulting societal preference order. Because we need to show that A is preferred over C regardless of the preferences of any of the other voters besides v_j, we won't make any assumptions about the preferences of these other voters. The only thing we'll assume is that v_j prefers A over C. For convenience, we'll call the corresponding preference schedule S.

Question 5.11.* Would any of the following changes to S affect the resulting societal preference between A and C? Explain your answer in each case.

(a) Voter v_j moves B between A and C on their individual preference order.

(b) Each of the voters $v_1, v_2, \ldots, v_{j-1}$ (all of the voters listed before v_j) moves B to first place on their individual preference order.

(c) Each of the voters $v_{j+1}, v_{j+2}, \ldots, v_n$ (all of the voters listed after v_j) moves B to last place on their individual preference order.

Question 5.12. Suppose all three changes listed in Question 5.11 are made to S, and call the resulting preference schedule S'.

(a) Explain why each voter's relative ranking of just A and B in S' is the same as their relative ranking of just A and B in the preferences shown in Table 5.2.

(b) Use your answer to part (a) to explain why, given the preference schedule S', A would be preferred over B in the resulting societal preference order.

(c) Use the same reasoning that you used in parts (a) and (b) to explain why, given S', B would be preferred over C in the resulting societal preference order.

(d) Now use your answers to parts (b) and (c) to explain why, given S', A would be preferred over C in the resulting societal preference order.

Question 5.13.* What do your answers to Questions 5.11 and 5.12 allow you to conclude about the societal preference order that would be produced by V given the preference schedule S (instead of S')? Explain.

Question 5.14.* Explain how your answer to Question 5.13 allows you to conclude that whenever v_j prefers A over C, society will also prefer A over C (regardless of the preferences of the other voters in the election).

We're actually very close to being done now. Recall we are trying to show that v_j is a dictator. What we have shown already is that v_j controls the societal preference between any pair of candidates that does not include B. We must now explain why v_j also controls the societal preference between any pair of candidates that does include B.

Here's where we need to be a little sneaky. Remember that we started just before Question 5.9 by assuming B was ranked in last place by every voter. This allowed us to identify a potential dictator, v_j, who controlled the societal preference between any pair of candidates that did not include B. Had we instead started the entire process by assuming that some other candidate, say A, was ranked last by every voter, we would have ended up with another potential dictator, say v_i, who would have controlled the societal preference between any pair of candidates that did not include A.

Question 5.15. Let C be any candidate other than A or B.

(a) Looking back at Question 5.10, clearly explain how v_j could possibly affect the societal preference between B and C. (Hint: Recall that v_j was chosen to be *pivotal* in some sense. You might want to try considering the preferences shown in Tables 5.2 and 5.3 and the resulting societal preference orders.)

(b) What does your answer to part (a) allow you to conclude about the relationship between v_i and v_j? (Hint: Recall that v_i *completely* controls the societal preference between any pair of candidates that does not include A.)

Question 5.16.*

(a) Considering your answer to Question 5.15, is the following statement true or false? Briefly explain how you know.

> There exists a single voter v^* that satisfies all three of the following properties:
>
> - v^* completely controls the societal preference between any pair of candidates that does not include A.
> - v^* completely controls the societal preference between any pair of candidates that does not include B.
> - v^* completely controls the societal preference between any pair of candidates that does not include C.

(b) If there were a voter v^* that satisfied all three of the properties from part (a), what could you conclude about v^*?

Question 5.17. Summarize what you have learned so far in this section by writing a detailed outline of how one could go about proving the strong form of Arrow's Theorem.

Before we can officially put Arrow's Theorem to bed, we still have one final detail to think about. Recall that we obtained the strong form of Arrow's Theorem by replacing Arrow's conditions 2 and 4 (monotonicity and citizen sovereignty) with the Pareto condition (unanimity). The following lemma is what makes this replacement possible.

Lemma 5.18. *If a voting system satisfies monotonicity, IIA, and citizen sovereignty, then it also satisfies unanimity.*

Question 5.19. Explain why Lemma 5.18, along with the strong form of Arrow's Theorem, implies the original form of Arrow's Theorem that we stated on page 64.

We've proved a few "if-then" statements now, so you should be getting the hang of it. As you may have noticed, a good first step in any proof is to identify what we can assume and what we need to show. Let's see if you can figure out these two components for the proof of Lemma 5.18.

Question 5.20.* To prove Lemma 5.18, what should we assume? What should we try to show?

So let's assume that we have a voting system V that satisfies exactly the properties you identified in Question 5.20. We must now show that if every voter in the system prefers A over B, then it will also be the case that A

is preferred over B in the societal preference order produced by V. We'll start by assuming that we have some arbitrary preference schedule, say S, in which every voter prefers A over B. We'll then attempt to explain why, given S, A would be preferred over B in the resulting societal preference order. To do this, we'll need to consider two other preference schedules that are related to S. For the first, let S' denote any preference schedule for which A would be preferred over B in the resulting societal preference order.

Question 5.21. Which property that we assumed about V allows us to conclude that such a preference schedule S' actually exists?

For the second, by only moving A on the individual preference orders that make up S', we'll create a new preference schedule in which every voter prefers A over B. Call this new preference schedule S''.

Question 5.22.

(a) With regard only to individual preferences between just A and B, how do the preference schedules S' and S'' differ?

(b) What does your answer to part (a) allow you to conclude about the societal preference between A and B that would be produced by V given S''? Which property that you assumed about V allows you to conclude this?

(c) With regard only to individual preferences between just A and B, how do the preference schedules S and S'' differ?

(d) What does your answer to part (c) allow you to conclude about the societal preference between A and B that would be produced by V given S? Which property that you assumed about V allows you to conclude this?

Question 5.23.* What does your answer to Question 5.22 allow you to conclude about the voting system V and the property of unanimity? Does this finish the proof of Lemma 5.18? Explain.

Potential Solutions

Now that we've seen why both forms of Arrow's Theorem are true, let's look at a few different ways through which we might be able to resolve the problems revealed by Arrow's work.

Weakening the Pareto Condition

As we saw in Chapter 4, unanimity (the Pareto condition) is a fairly strong property. For one thing, it rules out the possibility of two candidates ending up tied in a societal preference order when one of the candidates is unanimously preferred over the other. But, as we've seen, this is exactly the kind of behavior that can occur with plurality or instant runoff when several candidates end up tied with zero first-place votes. Fortunately, Pareto's original unanimity condition can be modified very slightly to allow for such ties.

The Modified Pareto Condition. If there is a pair of candidates in an election, say A and B, such that every voter in the election prefers A over B, then B should not be ranked higher than A in the resulting societal preference order.

Question 5.24.* Does plurality satisfy the modified Pareto condition? What about instant runoff? Sequential pairwise voting? Give a convincing argument to justify each of your answers.

As it turns out, there are voting systems that satisfy Arrow's conditions 1, 3, and 5, and the modified Pareto condition.

Question 5.25. Let V be the voting system in which all candidates are tied in the resulting societal preference order, regardless of the ballots cast. Explain why V satisfies Arrow's conditions 1, 3, and 5, and the modified Pareto condition.

Unfortunately, the voting system from Question 5.25 is not a particularly interesting or useful one. Moreover, thanks to Stanford economist Robert Wilson, we know that there is not much hope of finding a more useful system that satisfies the same conditions. In fact, if we also require neutrality, Wilson's Theorem [**57**] implies that any system that satisfies universality, IIA, and the modified Pareto condition must be a dictatorship, an *inverse* dictatorship (where the societal preference is always the *opposite* of what the dictator wants), or always result in a tie among all of the candidates (as in Question 5.25).

Ballot Restrictions and Approval Voting

In the 1970s, several political analysts independently proposed a new method for deciding the winner of an election with more than two candidates. This method, now commonly referred to as **approval voting**, works as follows:

- Each voter votes to either *approve* or *disapprove* of each candidate in the election. Voters can approve of as many candidates as they wish.
- The societal preference order is determined by the number of approval votes each candidate receives, starting with the candidate who receives the most approval votes and ending with the candidate who receives the fewest (with ties permitted, if candidates receive identical numbers of approval votes).

Question 5.26.* Three friends, Peter, James, and John, are trying to decide who among them is the greatest. To do so, they ask nine of their friends to cast approval ballots. The results are shown in Table 5.4, with ✓ indicating a vote of approval.

	Number of Voters		
Candidate	4	3	2
Peter	✓		✓
James		✓	✓
John			

TABLE 5.4. Approval voting

(a) Under approval voting, who would be declared the greatest? What societal preference order would be produced?

(b) Do you think that the outcome of the election under approval voting accurately reflects the will of the voters? Why or why not?

Approval voting has been adopted by a number of scientific and technical organizations, such as the American Mathematical Society, the Institute for Operations Research and Management Science, the American Statistical Association, and the Institute of Electrical and Electronics Engineers. Each of these organizations uses approval voting to elect officers and make other important decisions. Approval voting is also used to elect the Secretary-General of the United Nations, as well as new members of the National Academy of Sciences. It is also used for internal elections within political parties in some states.

Many proponents of approval voting have argued that since the method avoids using ranked ballots, Arrow's Theorem does not apply. At first glance, this conclusion seems entirely logical. In reality, however, the situation is a bit more complex. The real question is this: Can approval voting be viewed as a voting system in the sense of Definition 4.13? That is, can approval voting be viewed as a function that receives as input a collection of transitive preference ballots and produces as output a transitive societal preference order? If so, then Arrow's Theorem still applies. If not, then we may have found the perfect voting system after all. Let's examine this issue a little more closely.

Question 5.27.* Consider the two voters represented in Table 5.4 who approved of both Peter and James. Which of the following individual preference orders could be consistent with these two voters' approval ballots?

(a) Peter $\succ$ James $\succ$ John

(b) James $\succ$ Peter $\succ$ John

(c) Peter $\approx$ James $\succ$ John

(d) James $\approx$ Peter $\approx$ John

(e) Peter $\succ$ James $\approx$ John

Question 5.27 suggests that even though approval voting requires a different type of ballot than the voting systems we've considered previously, the underlying preferences of the voters can in fact be viewed in the same way. Admittedly, the correspondence between approval ballots and preference orders is somewhat loose, since each possible approval ballot will likely be consistent with many different preference orders (and, likewise, each possible preference order will likely be consistent with many different approval ballots). However, there are some natural conventions we can adopt to help us make the translation between the two.

Let's look back at Question 5.27. Which of the preference orders in that question do you think *best* represents the approval ballots of the two voters who approved of both Peter and James? It would be reasonable to argue that the one from part (c) (Peter $\approx$ James $\succ$ John) is the best choice because it does not specify any sort of ranking between Peter and James. It accurately reflects the information given in the voters' approval ballots, and avoids making any additional assumptions about the relative rankings of the candidates, other than those that can be directly inferred from the information supplied on the approval ballots. These observations suggest the following way to formally define approval voting.

Definition 5.28. The voting system known as **approval voting** is characterized by the following two conditions:

- The system accepts as input only those preference orders in which the symbol $\succ$ appears exactly once. In other words, the only preference orders allowed are those that have one or more candidates tied for first, followed by all of the remaining candidates tied for last.
- The societal preference order is determined by the number of first-place votes received, starting with the candidate who receives the most and ending with the candidate who receives the fewest (with ties permitted, if candidates receive identical numbers of first-place votes).

Incidentally, we will refer to the first-place votes from Definition 5.28 as *approval votes*, a convention that is completely consistent with our more intuitive understanding of approval voting.

Question 5.29.*

(a) Is approval voting, as described in Definition 5.28, a voting system according to Definition 4.13?

(b) By its very definition, approval voting violates one of the important fairness criteria that we have discussed. Which criterion does it violate, and is this violation acceptable or unacceptable in your opinion? Clearly explain your answers.

We can see from Question 5.29 that approval voting—by definition—violates one of Arrow's conditions. But what about Arrow's other conditions, such as the elusive IIA?

Question 5.30.* Suppose an election is held using approval voting, but that, due to voter irregularities, a revote is necessary. Suppose also that, in this revote, some voters change their ballots, but never in a way that affects their individual preferences between just candidates A and B.

(a) Explain why, in the revote, the difference in the number of approval votes received by A and B will be exactly the same as it was in the original election.

(b) What does your answer to part (a) allow you to conclude about approval voting and IIA?

Question 5.31.*

(a) Is approval voting anonymous? Neutral? Monotone? Clearly explain your answers.

(b) Does approval voting satisfy the Pareto condition? Why or why not?

Question 5.32. Suppose that approval voting was proposed as the method for electing the student body president at your school. Would you support or oppose this proposition? Write a formal letter to the editor of your school's newspaper expressing your views. Incorporate into your argument some of the features of approval voting that we considered in this section, as well as some practical considerations that might be relevant to implementing an approval voting system.

Weakening IIA: Intensity of Binary Independence

We'll conclude this section by considering a potential resolution to Arrow's Theorem that was proposed by Donald Saari, a professor of mathematics and economics at the University of California, Irvine. Saari's interpretation of Arrow's Theorem can be summarized as follows:

- We are only considering voting systems that produce transitive societal preference orders. To avoid "garbage in, garbage out" behavior, we must also require individual preferences to be transitive.

- Transitivity forces connections between pairwise comparisons in individual voter preferences. For instance, if a voter prefers A over B and B over C, then transitivity requires the voter to also prefer A over C.
- The IIA criterion requires voting systems to determine the societal preference between any pair of candidates based solely on the individual voters' preferences between those two candidates. This requirement forces voting systems to throw away the connecting information supplied by transitivity, essentially making it impossible for systems that satisfy IIA to distinguish between voters with rational, transitive preferences and voters with irrational, cyclic ones.

According to Saari, the IIA criterion essentially annihilates the assumption that individual preferences are transitive. This then causes perfectly reasonable voting systems to be capable of producing impermissible cyclic societal preferences. Saari's solution to this problem is to weaken IIA by allowing voting systems to take into account not only the pairwise preferences of the voters in the system, but also the *intensity* with which they hold these preferences. We formalize this idea in the next two definitions.

Definition 5.33. Suppose that a voter's preference between two candidates is $A \succ B$. The **intensity** of this preference is the number of candidates listed between A and B on the voter's individual preference order.

Question 5.34.* For each of the following preference orders, what is the intensity of the voter's preference between candidates A and B?

(a) $A \succ B \succ C \succ D$

(b) $A \succ C \succ D \succ B$

(c) $D \succ A \succ C \succ B$

Definition 5.35. A voting system is said to satisfy the **intensity of binary independence** criterion (**IBI** for short) if the societal preference between any two candidates depends only on the individual voters' preferences between those two candidates *and* the intensity with which they hold these preferences.

To state Definition 5.35 in another way, with a voting system that satisfies IBI, if some or all of the voters in an election change their preference ballots, but no voter changes their preference between two specific candidates A and B, or the intensity with which they hold this preference, then the societal preference between A and B will remain unchanged.

Question 5.36.* Suppose Greg, Sharon, Dean, and Carolyn are the last four competitors on the newest reality TV show, *Starvation Island*. According to the rules of the show, the Borda count is used to determine the player eliminated during each round of the contest. Unfortunately, after the ballots are cast by the four competitors, Greg experiences a moment of weakness and begins to eat them. Sharon, Dean, and Carolyn eventually manage to restrain him, but only in time to recover the following information:

- Two ballots contain the partial ordering $G \succ S$.
- One ballot contains the partial ordering $S \succ C \succ G$.
- The remaining ballot contains the full ranking $G \succ D \succ C \succ S$.

Using only this information, what can you conclude about the resulting societal preference between Greg and Sharon?

Question 5.37. Suppose that for a particular pair of candidates in an election, say A and B, you know each of the voters' preferences between these two candidates and the intensity with which they hold these preferences. Suppose also that the voting system used in the election is the Borda count.

(a) Do you have enough information to determine the societal preference between A and B? Give a convincing argument to justify your answer. (Hint: You may find it helpful to consider again the reasoning that you used in Question 5.36.)

(b) Does the Borda count satisfy IBI? Why or why not?

Question 5.38.* Is there a voting system that satisfies universality, IBI, and unanimity, and is not a dictatorship? Would such a system contradict the strong form of Arrow's Theorem as we stated it on page 75?

Concluding Remarks

Throughout the last five chapters, we've learned a lot about voting systems and how they do or do not live up to the standards that we might reasonably expect them to. We've also seen how Arrow's Theorem tells us that certain fairness criteria are incompatible with each other—no matter what voting systems we consider. What Arrow's Theorem doesn't say, however, is that there aren't any good or reasonable voting systems to choose from. Our success in finding a voting system that behaves the way we'd want it to depends on how willing we are to compromise on certain desirable features.

Finally, it's important to note that, while we've been focusing on the mathematical properties of voting systems, practical considerations must also be taken into account. For instance, sequential pairwise voting would likely be expensive and time-consuming in elections with many candidates. And ranked systems, such as the Borda count, could also be difficult to implement in large elections. (See Question 2.32 for an example of an election with 135 candidates; can you imagine trying to rank all of them?) The fact that no voting system is perfect explains why there is so much debate about which systems should be used in various types of elections. Hopefully, our investigations up to this point have prepared you to be able to understand and evaluate the arguments put forth in these kinds of discussions, and to make your own judgments about the best ways to effectively implement democracy.

Question 5.39. Considering everything we've learned so far, which voting system do you think is the best? Give a convincing argument to justify your answer, addressing both mathematical and practical considerations. Does your answer depend on the type of election and/or number of candidates involved? If so, explain how.

Questions for Further Study

Question 5.40. Explain where we used each of the following assumptions in our proof of Arrow's Theorem:

- There are more than two candidates in the election.
- The voting system must produce transitive societal preference orders.
- The voting system must satisfy universality.
- The voting system must satisfy IIA.
- The voting system must satisfy unanimity.

Question 5.41. Which of the properties of universality, IBI, and unanimity are satisfied by Black's voting system (see Warmup 4.1), and which are not? Give a convincing argument or example to justify each of your answers.

Question 5.42. Write a short biography of Donald Saari, including his most important contributions both inside and outside of voting theory, and the voting system that he prefers.

Question 5.43. Find a copy of an article written by Donald Saari in which he uses geometric ideas to analyze voting systems, and write a summary of whatever you can understand in the article.

Question 5.44. Find some information about Steven Brams, a professor of politics at New York University, and write a summary of your findings. What voting system does he prefer? What do you think a debate between Brams and Donald Saari would be like? Assuming each made a good case for his preferred system, whose side would you take?

Question 5.45. Find out how voting is conducted for enshrinement into the National Baseball Hall of Fame, and write a detailed summary of your findings. Include in your summary a description of how nominees are selected, who votes, the voting system that is used to decide which nominees will be enshrined, the criteria failed nominees must satisfy to be held over onto the next ballot, and some actual examples for illustration.

Question 5.46.

(a) If approval voting had been used instead of plurality to determine the winner of the the 2016 U.S. presidential election in New Hampshire, who do you think would have won the state? Give a convincing argument to justify your answer.

(b) If approval voting had been used instead of plurality to determine the winner of the the 2000 U.S. presidential election in Florida, who do you think would have won the national election? Give a convincing argument to justify your answer.

(c) If approval voting had been used instead of plurality to determine the winner of the 1998 Minnesota gubernatorial election, who do you think would have won? Give a convincing argument to justify your answer.

Question 5.47. Propose a method of modeling approval voting (that is, a way of associating approval ballots with preference orders) so that Arrow's universality condition *is* satisfied. In your proposed model, which of Arrow's other conditions are satisfied, and which are violated?

Question 5.48. Find an editorial either online or in a popular media source that supports approval voting as a method for deciding elections with more than two candidates. Write a summary of the editorial, and compare the arguments made in it to our investigations in this chapter.

Question 5.49. Find an editorial either online or in a popular media source that supports the Borda count as a method for deciding elections with more than two candidates. Write a summary of the editorial, and compare the arguments made in it to our investigations in this chapter.

Question 5.50. Find an article or book that suggests a potential method for resolving Arrow's Theorem that is different from the three methods we investigated in this chapter. Discuss the pros and cons of the potential resolution you found, and compare it with the three methods we investigated.

Question 5.51. Some have argued that in times of emergency, short periods of dictatorship may be necessary and even desirable. Do you believe that a dictatorship can ever be beneficial to society? If so, under what circumstances? Give a convincing argument to justify your answer.

Question 5.52. Research *single-peaked* preferences, and write a summary of your findings. How can single-peaked preferences be used to resolve Arrow's Theorem, and which economist (already mentioned in a previous chapter) is known for his theorems about single-peaked preferences?

Question 5.53. Consider the following preference schedule for an election with 3 candidates:

	Number of Voters					
Rank	1	1	1	1	1	1
1	A	A	B	B	C	C
2	B	C	A	C	A	B
3	C	B	C	A	B	A

Suppose that approval voting is used to decide the outcome of the election, and suppose also that each voter approves of either one or two candidates. (Approving of no candidates or all three candidates is not a good strategic choice, since doing so has no impact on the outcome of the election.) Show that, depending on where each voter draws the line between approved and disapproved candidates, any societal preference order is possible.

Question 5.54. For an election with four candidates, consider a variation of the Borda count in which the points assigned are 5, 3, 1, and 0 (instead of the usual 3, 2, 1, and 0).

(a) Does this system satisfy IBI? Why or why not?

(b) In general, what conditions must be placed on the points associated with each ranking in order for the resulting system to satisfy IBI?

Question 5.55. Does approval voting satisfy the CWC? Give a convincing argument or example to justify your answer.

Answers to Starred Questions

5.4. Since B is not ranked first in the societal preference order, there must be some other candidate who is ranked above or tied with B. Likewise, since B is not ranked last, there must be some other candidate who is ranked below or tied with B.

5.6. Because every voter ranks B either first or last on their individual preference order, moving C above A with no other changes will not affect any voter's individual preference between A and B or between B and C. Thus, by IIA, the resulting societal preference between A and B and between B and C must remain unchanged.

5.8. The contradiction is that $A \succsim C$ (by Question 5.5) and $C \succ A$ (by Question 5.7). These two relations cannot both be true at the same time. Remember that we assumed that B was not ranked first or last in the societal preference order produced by V (even though every individual voter ranked B first or last). Since this assumption led to a contradiction, it must be the case that B is in fact ranked either first or last in the societal preference order produced by V. This, however, is exactly the conclusion of Lemma 5.3.

5.10. (a) If all of the voters moved B from last place to first place, then by unanimity, B would have to be ranked first in the resulting societal preference order.

(b) If only some of the voters moved B from last place to first place, then by Lemma 5.3, B would have to remain last or move to first in the resulting societal preference order.

(c) Since the societal preference order would have to change—from ranking B last if every voter ranked B last to ranking B first if every voter ranked B first—then as the voters moved B from last place to first place on their individual preference orders, there would have to be some particular voter for which this change in the resulting societal preference order would first occur.

5.11. By IIA, none of the changes listed would affect the resulting societal preference between A and C.

5.13. Since the societal preference between A and C is the same given either S or S' (by Question 5.11), and A is preferred over C in the societal preference order produced from S' (by Question 5.12), we can conclude that A would be preferred over C in the societal preference order produced from S.

5.14. We assumed before Question 5.11 only that v_j prefers A over C. This assumption by itself allowed us to conclude (with some work) that A must also be preferred over C in the resulting societal preference order. Since we assumed nothing about the preferences of any of the other voters besides v_j, we could make this same conclusion even if some or all of the other voters opposed v_j's preference of A over C.

5.16. The statement from part (a) is true, since v_j satisfies all three conditions. (We established this for the first two conditions, and a similar argument would work for the third condition.) We can conclude that v^* is v_j—a dictator.

5.20. To prove Lemma 5.18, we should assume that we have a voting system that satisfies monotonicity, IIA, and citizen sovereignty. We should then try to prove that this voting system also satisfies unanimity.

5.23. We can conclude that V satisfies unanimity, just as we wanted. This indeed finishes the proof of Lemma 5.18.

5.24. Plurality satisfies the modified Pareto condition. If every voter in an election preferred A over B, then B would definitely not receive any first-place votes. Thus, A could not receive fewer first-place votes than B, and so under plurality B could not be ranked higher than A in the resulting societal preference order (though A and B could be tied). Instant runoff also satisfies the modified Pareto condition, although sequential pairwise voting does not (as demonstrated by Question 4.21).

5.26. (a) Since Peter received 6 votes of approval, James 5, and John 0, the societal preference order produced by approval voting would be Peter $\succ$ James $\succ$ John.

5.27. Each of the individual preference orders in parts (a), (b), and (c) could be consistent with the two voters' approval ballots. However, the preference orders in parts (d) and (e) could not, since each involves at least one tie between a pair of candidates, one of whom was approved and the other not.

5.29. Approval voting is a voting system according to Definition 4.13, although it violates universality by placing restrictions on the kinds of preference orders that can be provided as input.

5.30. (a) Suppose that, in the original election, a voter had approved of both A and B (i.e., placed both A and B in first place on his or her individual preference order). In the revote, the voter would have to either again approve of both A and B, or else disapprove of both A and B. In the first case, both A and B would *keep* the approval vote they had earned from the voter in the original election. In the second case, both A and B would *lose* the approval vote they had earned from the voter in the original election. In either case, both candidates would be affected in exactly the same way, and so the overall difference in their approval votes would not change. The same would be true if the voter had originally disapproved of both A and B, or if the voter had originally approved of one and disapproved of the other.

(b) Part (a) allows you to conclude that approval voting satisfies IIA.

5.31. Approval voting is anonymous, neutral, and monotone, *and* it satisfies the Pareto condition.

5.34. (a) Since no candidates are listed between A and B, the intensity of the $A \succ B$ preference is 0.

(b) Since two candidates (C and D) are listed between A and B, the intensity of the $A \succ B$ preference is 2.

5.36. On each of the two ballots containing the partial ordering $G \succ S$, Greg would receive 1 more point than Sharon. On the ballot containing the partial ordering $S \succ C \succ G$, Sharon would receive 2 more points than Greg. On the remaining ballot, Greg would receive 3 more points than Sharon. Thus, the total point differential between Greg and Sharon would be 3 points in favor of Greg, yielding the societal preference $G \succ S$.

5.38. The Borda count satisfies universality, IBI, and unanimity, and is not a dictatorship. This does not contradict the strong form of Arrow's Theorem though, since IBI is a weaker criterion than IIA.