

Chapter 6

Gaming the System

Focus Questions

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

- What is strategic voting?
- What are some common types of strategic voting for various voting systems?
- What does the Gibbard-Satterthwaite Theorem say about manipulation of voting systems?
- How is the Gibbard-Satterthwaite Theorem related to Arrow's Theorem? How can it be proved?

Warmup 6.1. The emergence of a competitive third party has complicated the previously straightforward mayoral elections in Stickeyville. In the upcoming election, three candidates are running: Deion Davis (a Democrat), Rachel Ramirez (a Republican), and Gabriella Gardner (of the Green Party). The preferences of the residents of Stickeyville are as shown in Table 6.1.

Rank	Number of Voters			
	25	15	15	45
1	<i>G</i>	<i>D</i>	<i>D</i>	<i>R</i>
2	<i>D</i>	<i>G</i>	<i>R</i>	<i>D</i>
3	<i>R</i>	<i>R</i>	<i>G</i>	<i>G</i>

TABLE 6.1. Preference schedule for Stickeyville

- (a) If the residents of Stickeyville used instant runoff to decide the outcome of the election, who would win, and what would the resulting societal preference order be?

- (b) Suppose that five of the $D \succ G \succ R$ voters decide that they actually like Gardner best and change their preferences accordingly, to $G \succ D \succ R$. Would this change result in a better or worse outcome for these voters?
- (c) Suppose that six of Ramirez’s 45 voters change their preferences to $G \succ R \succ D$. How would this change impact the outcome of the election?

Throughout most of our investigations so far, we have assumed that voters will vote *sincerely*. In other words, we have assumed that each voter will submit a ballot that represents their true preferences. However, as we saw in Warmup 6.1, there are situations in which it could benefit a voter to misrepresent their preferences. In part (b), the voters who switched to Gardner would have been better off sticking with Davis, even though he was no longer their first choice. In part (c), we observed something even more devious, where just a small number of Ramirez’s voters were able to exploit the sequential nature of the instant runoff system to secure a victory for their candidate. They accomplished this by adopting the unexpected strategy of ranking Gardner—their *last* choice—first. It was a gamble that paid off, as the extra support for Gardner caused Davis to be eliminated in the first round, leading to a solid victory for Ramirez in the second round.

What these examples demonstrate is that instant runoff is *manipulable*. And, as it turns out, most voting systems are. In this chapter, we’ll look at some common strategies for manipulating the voting systems we’ve studied in previous chapters. We’ll also explore an important result that tells us exactly which voting systems are immune to manipulation.

Strategic Voting

In a ranked voting system, it can be hard enough for voters to figure out exactly what their preferences are, particularly if there are a lot of candidates. When a voting system is manipulable, however, the choices become even more difficult. It’s no longer a matter of just deciding which candidates are better or worse; voters must also decide whether to vote in accordance with their actual preferences, or to cast a strategic ballot in the hopes of securing a better election outcome.

There are a number of different tactics that strategic voters might adopt, and we’ve seen examples of several of them already. Some of the most common strategies include:

- **Compromising**, in which a voter casts a ballot for a candidate other than their top choice, usually because they perceive the compromise candidate to be more electable.
- **Pushover** or **mischief voting**, in which a voter ranks a candidate higher than they actually prefer, but not with the intent of electing

that candidate. Instead, the aim is to indirectly benefit the voter's most preferred candidate by harming their opponents.

- **Burying**, in which a voter ranks a strong opponent lower than they actually prefer in order to have a better chance of defeating that opponent.
- **Bullet voting**, in which a voter has the option to vote for multiple candidates (as in approval voting), but votes for only one in order to deny support to opposing candidates.

Question 6.2.* Which of the strategies described above were illustrated in parts (b) and (c) of Warmup 6.1?

Question 6.3. Consider the main voting systems we have studied so far: majority rule, plurality, the Borda count, sequential pairwise voting, instant runoff, and approval voting.

- (a) Which of these systems do you think are most susceptible to compromising, and why?
- (b) Which systems do you think are most susceptible to pushover voting, and why?
- (c) Which systems do you think are most susceptible to burying, and why?

Question 6.4.* Suppose approval voting had been used for the election from Warmup 6.1. Give an example to illustrate how bullet voting could have been a smart strategic choice for some of the voters.

The Gibbard-Satterthwaite Theorem

Researchers have spent quite a bit of time studying manipulability of voting systems, in part because avoiding the potential for manipulation seems a perfectly natural and desirable goal. We would like to think that election outcomes can be determined solely by voter preferences and possibly the voting system used to aggregate the ballots. But if voters can sway the results of an election by voting strategically, some will do so. On the other hand, voters who are not as informed about the strategic choices available to them may cast ballots that are not effective—or they might disengage from the political process altogether. What would be ideal is if we could take the effects of strategic voting out of the equation by choosing a system in which all voters would have an incentive to simply be honest and vote according to their actual preferences.

The good news is that some voting systems are in fact immune to manipulation, including some that we studied back in Chapter 1.

Question 6.5. Explain why in each of the following systems, voters have no incentive to vote insincerely.

- (a) Majority rule with exactly two candidates
- (b) Dictatorship with any number of candidates
- (c) Imposed rule with any number of candidates

The voting systems in Question 6.5 are probably not the ones you were hoping for. After all, we already know that elections with only two candidates are pretty straightforward, and there are not many proponents of democracy who would seriously advocate for a dictatorship or imposed rule. So what other options are there?

The short answer to this question is . . . there aren't any. We have philosopher Allan Gibbard and economist Mark Satterthwaite to thank for this not-so-happy result. They proved, independently, in 1973 (Gibbard) and 1975 (Satterthwaite) that a dictatorship is the only non-manipulable voting system for three or more candidates that: (1) works with any set of transitive preferences; (2) always produces a unique winner; and (3) can produce any candidate as the winner.

In order to formally state the Gibbard-Satterthwaite Theorem, we'll need to define a few terms more precisely. First off, the theorem doesn't deal with voting systems exactly as we've defined them in past chapters. Instead, it is concerned only with selecting a winner—not a societal preference order—for each set of voter preferences. To make this distinction clear, we'll adopt the following definition:

Definition 6.6. A **choice function** is a function that receives as input any collection of transitive preference ballots and produces as output a unique winning candidate.

We will be interested in finding choice functions that cannot be manipulated. But what exactly does it mean for a choice function to be manipulable? The next definition provides an answer.

Definition 6.7. A choice function is **manipulable** if there is a situation in which some voter can singlehandedly ensure a more desirable outcome (from their perspective) by submitting a ballot that is different from their true preferences.

When we talk about manipulation, we are always looking at voting situations in which a single voter changes their ballot, while the remaining ballots stay fixed. The idea is that manipulation does not require cooperation—or at least there are some situations where only one voter changing their ballot is enough to change the outcome of the election.

We also need to define what a dictator is with respect to a choice function.

Definition 6.8. For a given choice function, a **dictator** is a voter who can force any candidate to win by ranking that candidate first (without ties) on

their ballot. A choice function is said to be **dictatorial** if some voter is a dictator with respect to it.

Finally, we need to say precisely what it means for a choice function to be able to produce any candidate as the winner. We'll capture this idea with a definition similar to that of citizen sovereignty in Arrow's Theorem.

Definition 6.9. A choice function is **non-imposed** if for every candidate there is a set of ballots for which the choice function will produce that candidate as the winner.

With these definitions in mind, we can now state the Gibbard-Satterthwaite Theorem.

Gibbard-Satterthwaite Theorem. *There does not exist a choice function for an election with more than two candidates that is non-imposed, non-manipulable, and non-dictatorial.*

Proving the Gibbard-Satterthwaite Theorem

If the Gibbard-Satterthwaite Theorem reminds you of Arrow's Theorem, you're wise to observe the similarities between the two. Both are impossibility theorems, both deal with elections having more than two candidates, and both involve non-imposition and non-dictatorship conditions. In light of these similarities, it may not surprise you to learn that the Gibbard-Satterthwaite Theorem is actually a corollary of Arrow's Theorem, meaning that we can use Arrow's Theorem to prove it. For the rest of the chapter, that's what we'll do. The proof is not easy, but—like the proof of Arrow's Theorem—we'll be able to understand it by breaking it down into smaller steps.¹

Our main strategy will be to show that any non-imposed, non-manipulable choice function for an election with three or more candidates must in fact be dictatorial. We'll do so by showing that any such choice function can be extended to a voting system that satisfies universality, unanimity, and IIA. By the strong form of Arrow's Theorem, we will then be able to conclude that the voting system we constructed must be a dictatorship. Finally, we'll use this conclusion to show that the original choice function must also be dictatorial.

Monotonicity and Non-Manipulability

Before we go any further, there is one more useful property that we need to consider. It is a version of the monotonicity property we've studied in previous chapters, but this time with regard to choice functions.

¹¹Our approach is based on that of Vazirani et al. [56] in their book *Algorithmic Game Theory*.

Definition 6.10. A choice function is said to be **monotone** if the only way a single voter can cause the winner to change—say from A to B —is for that voter to change their ranking of A and B from $A \succ B$ to $B \succ A$.

Notice that this definition of monotonicity is a little stronger than the one we've used in the past. In addition to saying that a change favorable to A cannot cause the winner to change from A to B , it says that in order for a voter to singlehandedly force the winner to change from A to B , that voter must make a specific change favorable to B —namely, they must change their ranking of A and B from $A \succ B$ to $B \succ A$. We will use this stronger definition in order to more easily see a relationship between monotonicity and manipulability.

To that end, note that if a choice function is not monotone, this means that there are two sets of ballots for which: (1) the choice function produces different winners—say A for the first set of ballots and B for the second set; (2) the two sets of ballots are identical except for one special voter; and (3) this special voter ranks A and B the same in both sets of ballots (either $A \succ B$ or $B \succ A$).

Question 6.11.* Explain why a choice function that is not monotone must be manipulable.

Worded slightly differently (but equivalently), Question 6.11 establishes the following lemma:

Lemma 6.12. *Every non-manipulable choice function is monotone.*

Choice Functions and Voting Systems

Now let's continue with our proof of the Gibbard-Satterthwaite Theorem. We'll start with any choice function for three or more candidates that is non-imposed and non-manipulable. Let's call that function f (a letter mathematicians often like to use to denote functions). We're now going to try to use f to define a voting system.

Remember that the difference between a choice function and a voting system is that a voting system must output a transitive societal preference order, whereas a choice function just gives us a winner. In other words, we need to *extend* f so that it gives us not just a winner, but a ranking of the candidates. Here's how we'll do it:

- To decide the societal ranking between any pair of candidates, say A and B , we'll first move A and B to the top of every voter's ballot. We won't change any other rankings, and we won't change the voters' relative rankings *between* A and B ; if a voter prefers one to the other, it will stay that way, and if a voter is indifferent between A and B , they will remain tied. What will change is that A and B will both be ranked strictly higher than every other candidate.

- Once we've moved A and B to the top of every voter's ballot, we'll ask f who it will choose as the winner. If f chooses A , we'll say society prefers A to B . If f chooses B , we'll say society prefers B to A .
- After we've considered every possible pair of candidates, we'll combine these pairwise rankings to get a societal preference order.
- We'll let F denote the voting system that sends each set of ballots to the societal preference order constructed as described above.

To illustrate, let's consider an example.

Question 6.13.* Assume that the voters in an election are arranged in some order, say v_1, v_2, \dots, v_n , and let f be the choice function that works as follows:

- If v_1 has a unique top-ranked candidate, then that candidate is declared the winner. Otherwise, the candidates who are not among v_1 's top choices are eliminated.
- If v_2 has a unique top-ranked candidate among those who remain, then that candidate is declared the winner. Otherwise, the remaining candidates who are not among v_2 's top choices are eliminated.
- The process continues in this same manner, with each successive voter eliminating candidates until only one remains. If multiple candidates still remain after the last voter, then the tie is broken by alphabetical order, with the candidate whose name comes first in the alphabet being declared the winner.

Now suppose that there are three voters in the election, with the following preferences:

- $v_1 : A \approx B \succ C \approx D$
- $v_2 : B \succ C \approx D \succ A$
- $v_3 : D \succ C \succ A \succ B$

Given these ballots, what societal preference order would F produce?

The Nitty-Gritty Details

There are a couple of details we need to check in order to make sure that F , as we've defined it above, always produces a transitive societal preference order.

First, in constructing F , we implicitly assumed that whenever two candidates are moved to the top of every voter's ballot, f will choose one of those two candidates as the winner. But what if this didn't happen? What if, for example, we moved A and B to the top of every ballot, but f chose some other candidate C as the winner? In this case, we wouldn't be able to write down a ranking of A and B . Fortunately, however, the next lemma rules out this possibility.

Lemma 6.14. *If f is a non-imposed, non-manipulable choice function, and if there is a set S of candidates that are at the top of every voter's preference order (meaning that every voter prefers every candidate in S to every candidate not in S), then f must choose a candidate in S as the winner.*

In the next question, we'll develop an argument to show why Lemma 6.14 is true. It may take a little while to fully understand this argument, but it's important to put in the effort, since we'll use similar arguments throughout the rest of the proof of the Gibbard-Satterthwaite Theorem. To make matters somewhat simpler, we'll look at the special case of the lemma in which the set S has exactly two candidates—say A and B —in it.

Question 6.15.* Let $b_1, b_2, b_3, \dots, b_n$ be a set of ballots where candidates A and B are at the top of every ballot.

- (a) Explain why there must be some other set of ballots, say $b_1^*, b_2^*, b_3^*, \dots, b_n^*$, such that given those ballots, f would choose A as the winner.
- (b) What would happen if we gave f the ballots $b_1, b_2^*, b_3^*, \dots, b_n^*$? Could f choose a candidate other than A or B as the winner? Why or why not? (Hint: Only one voter's ballot is different from those in part (a). What would monotonicity require in this situation for the winner to change from A to some other candidate C ?)
- (c) What would happen if we gave f the ballots $b_1, b_2, b_3^*, \dots, b_n^*$? What possible winners could f choose?
- (d) If we continue this process, changing one ballot at a time, we will eventually get back to our original set of ballots: $b_1, b_2, b_3, \dots, b_n$. Explain why this observation establishes that, given these ballots, f must choose A or B as the winner.

Lemma 6.14 tells us that our method for extending a choice function f to a voting system F will, at the very least, produce a ranking of every pair of candidates. Now all we need to do is combine those pairwise rankings to get a societal preference order. But what if the pairwise rankings don't fit together nicely? What if we end up with $A \succ B$, $B \succ C$, and $C \succ A$ —a (gasp!) cycle? Apart from having some serious déjà vu from our investigations of Condorcet's paradox (see Question 3.4), this would really derail our efforts. After all, voting systems are supposed to produce *transitive* societal preference orders, and cycles are exactly the opposite of that. The good news is that we're lucky again: We'll never actually end up with a cycle when we follow the method we've been using. The next question shows why.

Question 6.16. Suppose that for some set of ballots, say b_1, b_2, \dots, b_n , the rankings produced by F combine to give the cycle $A \succ B \succ C \succ A$. If we move A , B , and C to the top of each of these ballots, we'll get a new set of

ballots, which we'll denote by

$$b_1^{ABC}, b_2^{ABC}, \dots, b_n^{ABC}.$$

Lemma 6.14 tells us that, given these new ballots, f must choose A , B , or C as the winner. For the sake of our argument, we'll suppose f chooses A . (We could make a similar argument if f chose B or C instead.)

- (a) Suppose that instead of moving A , B , and C to the top of the first ballot, we had just moved A and C to the top. We'll call that ballot b_1^{AC} . What winner would f choose if we gave it the ballots $b_1^{AC}, b_2^{ABC}, \dots, b_n^{ABC}$? (Hint: How would b_1^{AC} and b_1^{ABC} have to be different in order for f to no longer choose A ? Is this possible?)
- (b) Suppose we continue changing ballots, one by one (as in Question 6.15), until we obtain $b_1^{AC}, b_2^{AC}, \dots, b_n^{AC}$. What winner will f choose for this set of ballots?
- (c) What does your answer to part (b) tell you about how F (the voting system) will rank A and C ? Why is this a problem given our assumption that $A \succ B \succ C \succ A$?

The type of argument we used in Question 6.16 is called a *proof by contradiction*. If we assume that F produces a cycle, $A \succ B \succ C \succ A$, then we can show that it must also be the case that $A \succ C$. But we can't have both $C \succ A$ and $A \succ C$ at the same time—that would be a contradiction! And since we can never derive a contradiction from a true statement (as long as our logic is correct), this means that our original assumption that F can produce a cycle must be false.

Let's pause for a moment to catch our breath and take stock of what we've accomplished and where we're headed next. We've shown that we can extend any non-imposed, non-manipulable choice function f to a voting system F , and we've ensured that F always produces a transitive societal preference order. Combined with the fact that f , as a choice function, accepts *any* set of transitive preference ballots, we have actually shown that our voting system F satisfies Arrow's universality condition. Our goal is to show that F also satisfies unanimity and IIA, and is therefore—by Arrow's Theorem—a dictatorship.

Question 6.17. To prove that F satisfies unanimity, we'll assume that every voter ranks A above B on their ballot.

- (a) Suppose we move A and B to the top of every voter's ballot. Explain why A must then be the unique top-ranked candidate on every ballot.
- (b) In light of the observation from part (a), what does Lemma 6.14 allow us to conclude about the candidate chosen by f ?
- (c) Use your answer to part (b) to explain why F produces the societal ranking $A \succ B$ and therefore satisfies unanimity.

Question 6.18. Now let's show that F satisfies IIA. Consider two sets of ballots: b_1, b_2, \dots, b_n and $b_1^*, b_2^*, \dots, b_n^*$. Suppose also that both sets of ballots agree on their rankings of two candidates, say A and B . That is, suppose that the ranking of A and B on any ballot b_i is the same as on b_i^* .

- (a) Now move A and B to the top of each ballot to obtain

$$b_1^{AB}, b_2^{AB}, \dots, b_n^{AB}$$

and

$$(b_1^*)^{AB}, (b_2^*)^{AB}, \dots, (b_n^*)^{AB}.$$

Explain why these new sets of ballots must still agree on their rankings of A and B . In other words, explain why the ranking of A and B on b_i^{AB} is always the same as on $(b_i^*)^{AB}$.

- (b) Explain why, given either set of ballots— $b_1^{AB}, b_2^{AB}, \dots, b_n^{AB}$ or $(b_1^*)^{AB}, (b_2^*)^{AB}, \dots, (b_n^*)^{AB}$ — f must choose either A or B as the winner.
- (c) We want to argue that f chooses the same candidate when given $b_1^{AB}, b_2^{AB}, \dots, b_n^{AB}$ as when given $(b_1^*)^{AB}, (b_2^*)^{AB}, \dots, (b_n^*)^{AB}$. To do so, we will sequentially (one voter at a time) change the first set of ballots into the second, as we have done in previous questions. To start, explain why f will choose the same winner when given $(b_1^*)^{AB}, b_2^{AB}, \dots, b_n^{AB}$ as when given $b_1^{AB}, b_2^{AB}, \dots, b_n^{AB}$. (Hint: Only one ballot is potentially different between the two sets, and the difference is limited by your answer to part (a).)
- (d) Now continue changing ballots, one at a time, from b_i^{AB} to $(b_i^*)^{AB}$. Explain why each such change has no impact on the winner chosen by f .
- (e) Use part (d) to argue that f will choose the same winner given either $b_1^{AB}, b_2^{AB}, \dots, b_n^{AB}$ or $(b_1^*)^{AB}, (b_2^*)^{AB}, \dots, (b_n^*)^{AB}$. Explain why this implies that F ranks A and B the same way given either set of ballots, and thus F satisfies IIA.

We have now shown that F is a voting system that satisfies universality, unanimity, and IIA. But we know from Arrow's Theorem that any such voting system must be a dictatorship. This means that there is some voter, say v , whose ballot always agrees with the societal preference order given by F . Does this make v a dictator for f as well? Let's find out.

Question 6.19.* Let v be the dictator for the voting system F , and suppose that A is v 's unique top-ranked candidate. We want to see if f will always choose A as the winner. To do so, suppose that there is some set of ballots in which v ranks A first but f chooses another candidate, say B , instead.

- (a) What kind of proof are we setting up by making this assumption?
- (b) Suppose we move A and B to the top of each ballot. Explain why f will choose A in this situation. (Hint: Remember that v is a

dictator for F and ranks A above B . Also remember how F is defined.)

- (c) Now, starting with the ballots from part (b), change one ballot at a time to put A and B back where they were before we moved them to the top. Explain why none of these changes can cause f to choose B rather than A .
- (d) Explain why part (c) leads to a contradiction.

Question 6.19 finishes the proof of the Gibbard-Satterthwaite Theorem. Under the assumption that our choice function f is non-imposed and non-manipulable, we were able to find a dictator for f . This means that f cannot possibly be non-imposed, non-manipulable, *and* non-dictatorial.

Concluding Remarks

Like Arrow's Theorem, the Gibbard-Satterthwaite Theorem can be viewed as a disappointing result for the practice of democracy. In some sense, it confirms what we've seen in a variety of other settings—namely, that voting systems can sometimes produce unexpected outcomes. Naturally, there are times when the quirks of a particular system can be exploited by voters to give an advantage to their preferred candidates.

Also like Arrow's Theorem, there are resolutions to the Gibbard-Satterthwaite Theorem that make its conclusions seem somewhat less dire. Consider, for example, approval voting. If voters are really only interested in getting one of their approved candidates selected—and they don't like certain approved candidates better than others—then the only reasonable strategy is to vote for all of their approved candidates. The Gibbard-Satterthwaite Theorem doesn't apply in this setting, since we are restricting voters to having *dichotomous* preferences, as we did in Chapter 5. In other words, we are viewing approval voting as a sort of choice function that only accepts certain types of ballots—those that rank one group of approved candidates above one group of disapproved candidates, with no preferences expressed among the candidates within each group. Since our definition of a choice function doesn't allow this kind of restriction (analogous to the fact that our model of approval voting from Chapter 5 violates universality), we avoid Gibbard-Satterthwaite's unpleasant conclusions.

If we assume, however, that voters have additional preferences underlying their approval ballots, then a version of the Gibbard-Satterthwaite Theorem applies,² and voters must make the strategic decision of where to draw the line between approved and disapproved candidates. There might be multiple approval ballots that could all be viewed as sincere representations of the voter's preferences. For example, a voter with preferences

²²Gibbard's version of the theorem allows for a more complex model of ballots and strategies that accommodates mixing preference rankings with approval voting strategies.

$A \succ B \succ C$ could reasonably approve only of A , or of A and B , or of A , B , and C (although the latter would be a wasted ballot, since it would have no impact on the outcome of the election). Voting for only A (bullet voting) increases the chances of electing the voter's most preferred candidate but may also lead to a less desirable outcome if voters who prefer $C \succ B \succ A$ adopt a similar strategy. In this case, a compromise candidate such as B may receive very few approval votes but actually be viewed as acceptable by the entire electorate, had they voted sincerely.

Questions for Further Study

Question 6.20. Write a short biography of Allan Gibbard, including some of his contributions apart from the Gibbard-Satterthwaite Theorem.

Question 6.21. Write a short biography of Mark Satterthwaite, including some of his contributions apart from the Gibbard-Satterthwaite Theorem.

Question 6.22. Give an argument to explain why approval voting with dichotomous preferences is immune to manipulation.

Question 6.23. Give an example to show how sequential pairwise voting is subject to pushover voting. In your example, does it make a difference if you allow voters to change their ballots between rounds of voting?

Question 6.24. Explain why burying is not an effective strategy in elections decided by instant runoff. Give specific examples to illustrate your argument.

Question 6.25. Explain why pushover voting is not an effective strategy in elections decided by the Borda count. Give specific examples to illustrate your argument.

Question 6.26. Suppose that in an election decided by approval voting, you prefer candidate A over candidate B . Is there ever a situation in which it makes sense to approve of B and disapprove of A ? Give a convincing argument or example to justify your answer.

Question 6.27. Research Duverger's Law, and write a brief summary of your findings. Be sure to explain how Duverger's Law is related to manipulation and strategic voting.

Question 6.28. Research the participation criterion, and write a brief summary of your findings. Be sure to explain how the participation criterion is related to manipulation and strategic voting.

Question 6.29. Where in the proof of the Gibbard-Satterthwaite Theorem did we use the assumption that there are three or more candidates in the election? Why is this assumption essential to the theorem?

Question 6.30. Find an article or online source arguing that instant runoff is less manipulable than approval voting. Then find a source that argues the opposite. Summarize the key points made by each, and explain which argument you find more persuasive.

Answers to Starred Questions

6.2. Part (b) illustrated how *compromising*, by voting $D \succ G \succ R$, would have resulted in a better outcome for the Gardner converts than voting according to their true preferences, $G \succ D \succ R$. Part (c) gave an example of *pushover voting*, since the Ramirez voters who voted for Gardner did so not to help her win, but rather to harm Ramirez's main opponent, Davis. The strategy was effective, since Davis ended up being eliminated in the first round, ensuring an easy victory for Ramirez in the second round.

6.4. Suppose every voter approves of their top two choices. Then Davis will win the election with 100 approval votes, followed by Ramirez with 60 and Gardner with 40. However, if the Ramirez voters employ a bullet voting strategy and only approve of Ramirez, then Ramirez will win with 60 approval votes, followed by Davis with 55 and Gardner with 40.

6.11. Continuing the argument from the preceding paragraph, suppose that the special voter prefers B over A in both ballots. Since the first set of ballots produces A as the winner, whereas the second set produces B as the winner, this voter could singlehandedly ensure a better election outcome (from their perspective) by voting the second ballot instead of the first. Likewise, if the voter prefers A over B in both ballots, they can ensure a better election outcome by voting the first ballot instead of the second.

6.13. We'll need to consider each pair of candidates. For A and B , moving these candidates to the top of each voter's ballot yields the following preferences:

- $v_1 : A \approx B \succ C \approx D$
- $v_2 : B \succ A \succ C \approx D$
- $v_3 : A \succ B \succ D \succ C$

Now we need to see which candidate f will choose with these ballots. Since A and B are the top-ranked candidates on v_1 's ballot, f eliminates C and D before moving on to v_2 . Since v_2 prefers B to A , f chooses B as the winner. Therefore, F produces a societal ranking of $B \succ A$.

Looking at another pair, say C and D , we get the following ballots:

- $v_1 : C \approx D \succ A \approx B$
- $v_2 : C \approx D \succ B \succ A$
- $v_3 : D \succ C \succ A \succ B$

For these ballots, f eliminates A and B first, since they are not among v_1 's top-ranked candidates. Moving on to v_2 , f is unable to choose a winner since v_2 is indifferent between C and D . Finally, v_3 breaks the tie, resulting in a societal ranking of $D \succ C$.

Repeating this process for every possible pair of candidates yields the rankings $B \succ A$, $A \succ C$, $A \succ D$, $B \succ C$, $B \succ D$, and $D \succ C$. Combining these rankings yields the societal preference order $B \succ A \succ D \succ C$.

- 6.15.** (a) Since f is non-imposed, there must be some collection of ballots for which f chooses A as the winner.
- (b) The only difference between the ballots in part (a) and those in part (b) is that the first ballot is b_1 instead of b_1^* . Because f is monotone, for the winner to change from A to C , the first voter would have had to change their ranking of A and C from $A \succ C$ to $C \succ A$. But it can't be the case that $C \succ A$ in b_1 , since A and B are top-ranked in b_1 . Therefore f can't choose C as the winner.
- (c) By the same argument as in part (b), since we only changed one ballot—from b_2^* to b_2 —and A and B are both top-ranked in b_2 , f couldn't choose any winner other than A or B .
- (d) Each time we change a ballot from b_i^* to b_i , f still must choose either A or B as the winner. By the time we have completed the sequence of changes, we'll have the ballots b_1, b_2, \dots, b_n , and f will still have to choose either A or B as the winner.
- 6.19.** (a) We are setting up a proof by contradiction.
- (b) Moving A and B to the top of each ballot will not change the fact that v ranks A above B . Since the societal preference order produced by F must agree with v , it must yield a ranking of $A \succ B$. But because of the way F is defined, this means that f must choose A .
- (c) Moving A and B back to where they began on a ballot does not change the relative ranking of A and B on that ballot—only how A and B compare to the remaining candidates. By monotonicity, it is therefore impossible for such a change to cause f to choose B , rather than A , as the winner.
- (d) If we move A and B back, one ballot at a time, to their original positions, we will eventually end up with our original set of ballots. But none of the changes we make along the way will cause f to choose B as the winner. This is a contradiction, since we assumed that f would choose B as the winner given the original set of ballots.