most of the procedures. Why haven't we presented a number of natural procedures that satisfy *all* of these properties and more? We turn to this question next.

1.7 A GLIMPSE OF IMPOSSIBILITY

In Chapter 7 we shall return to the issue of social choice and present the single most famous theorem in the field: Arrow's impossibility theorem. The natural context for Arrow's theorem, however, is slightly different from the context in which we have explored social choice in the present chapter. Nevertheless, this section previews the kind of difficulty that Arrow's theorem shows is unavoidable. We will do this by stating and proving an impossibility theorem in the context with which we have worked in the present chapter. The proof of this theorem, like that of Arrow's theorem, makes critical use of the voting paradox of Condorcet.

Recall that in **Section 1.4** we introduced five desirable properties of social choice procedures: the always-a-winner condition, the Pareto condition, the Condorcet winner criterion, monotonicity, and independence of irrelevant alternatives. Of the six social choice procedures we looked at, only Condorcet's method and a dictatorship satisfied independence of irrelevant alternatives, and only Condorcet's method and sequential pairwise voting satisfied the Condorcet winner criterion. None of the six procedures satisfied all five of the desirable properties.

Suppose we were to seek a social choice procedure that satisfies all five of our desirable properties. One possibility is to start with one of the six procedures that we looked at and to modify it in such a way that a property that was not satisfied by the original procedure would be satisfied by the new version. For example, there is a very natural way to modify a procedure so that the Condorcet winner criterion becomes satisfied: If there is a Condorcet winner, then it is the social choice; otherwise, apply the procedure at hand.

It is tempting to think that if we modify a dictatorship in the above way, then we will have a social choice procedure that satisfies all five of the desirable properties we discussed in this chapter. This turns out not to be the case (and we will say why in a moment). But maybe there are other ways to alter one or more of the procedures from this chapter so that the result will satisfy all the desirable properties. Or maybe there are procedures that look very different from the ones we presented in this chapter that already satisfy these desirable properties. Or maybe no such procedures have ever been found, but that one will be found a hundred years from now.

No way.

There is no social choice procedure that satisfies all five of the desirable properties that we listed in **Section 1.4**. We are not just saying that none of the six procedures we looked at satisfies all five of the desirable properties—we already know that. We are not just saying that these procedures can't be altered to yield one that satisfies all five of the desirable properties. We are not just saying that no one has yet found a social choice procedure that satisfies all five of the desirable properties. We are saying that no one will *ever* find a social choice procedure that satisfies these five desirable properties. In fact, more is true:

THEOREM. *There is no social choice procedure for three or more alternatives that satisfies the always-a-winner criterion, independence of irrelevant alternatives, and the Condorcet winner criterion.*

We will assume that we have a social choice procedure that satisfies both independence of irrelevant alternatives and the Condorcet winner criterion. We will then show that if this procedure is applied to the profile that constitutes Condorcet's voting paradox (**Section 1.6**), then it produces no winner. Because any procedure satisfying IIA and the CWC fails to satisfy AAW, it follows that no procedure can satisfy all three criteria.

PROOF. Assume that we have a social choice procedure that satisfies both independence of irrelevant alternatives and the Condorcet winner criterion. Consider the following profile from the voting paradox of Condorcet:

$$\begin{array}{ccc} a & c & b \\ b & a & c \\ c & b & a \end{array} \qquad (1)$$

CLAIM 1. The alternative *a* is a nonwinner.

PROOF. Consider the following profile (obtained by moving alternative *b* down in the third preference list from the voting paradox profile):

a	*c*	*c*	
b	*a*	*b*	(2)
c	*b*	*a*	

Notice that *c* is a Condorcet winner for profile (2) (defeating both other alternatives by a margin of 2 to 1). Thus, our social choice procedure (which we are assuming satisfies the Condorcet winner criterion) must produce *c* as the only winner. Thus, *c* is a winner and *a* is a nonwinner for this profile. (We are not done proving the claim because this is not the voting paradox profile.)

Suppose now that the third voter moves *b* up on his or her preference list. The profile then becomes that of the voting paradox (since we just undid what we did earlier). We want to show that *a* is still a nonwinner.

But no one changed his or her mind about whether *c* is preferred to *a* or *a* is preferred to *c*. Thus, because our procedure is assumed to satisfy independence of irrelevant alternatives, and because we had *c* as a winner and *a* as a nonwinner in the profile with which we began the proof of the claim, we can conclude that *a* is still a nonwinner when the procedure is applied to profile (1). This proves the claim.

CLAIM 2. The alternative *b* is a nonwinner.

PROOF. Consider the following profile (obtained by moving alternative *c* down in the second preference list from the voting paradox profile):

a	*a*	*b*	
b	*c*	*c*	(3)
c	*b*	*a*	

Notice that *a* is a Condorcet winner for profile (3) (defeating both other alternatives by a margin of 2 to 1). Thus, our social choice procedure (which we are assuming satisfies the Condorcet winner criterion) must produce *a* as the only winner. Thus, *a* is a winner and *b* is a nonwinner for this profile. (We are again not done proving the claim because this is not the voting paradox profile.)

Suppose now that the second voter moves *c* up on his or her preference list. The profile then becomes that of the voting paradox (since we just undid what we did earlier). We want to show that *b* is still a nonwinner.

But no one changed his or her mind about whether *a* is preferred to *b* or *b* is preferred to *a*. Thus, because our procedure is assumed to satisfy independence of irrelevant alternatives, and because we had *a* as a winner and *b* as a nonwinner in the profile with which we began the proof of the claim, we can conclude that *b* is still a nonwinner when the procedure is applied to profile (1). This proves the claim.

CLAIM 3. The alternative *c* is a nonwinner.

PROOF. We leave this for the reader (see Exercise 40).

The above three claims show that when our procedure is confronted with the voting paradox profile, it produces *no* winner. Thus, any social choice procedure satisfying IIA and the CWC fails to satisfy AAW. This completes the proof.

This is only part of the remarkable story of the difficulty with "reflecting the will of the people." More of the story will be told in Chapter 7.

1.8 APPROVAL VOTING

The voting systems we have considered so far are social choice procedures: a collection of individual preference lists (without ties) is the input, and the output is a single or possibly a collection of alternatives. There are, however, other types of voting systems. Here we consider one of the most popular alternative methods—approval voting. Approval voting was explicitly proposed in the 1971 Ph.D. thesis of Robert Weber at Yale University. Since then, Steven Brams, a political scientist at NYU, and Peter Fishburn, a former researcher at Bell Laboratories, have done much more research on and promotion of approval voting. Under approval voting, given a set *A* of alternatives, each voter votes for (or "approves of") as many alternatives as he or she chooses. The voters do not rank the alternatives. The social choice is the alternative (or set of alternatives) with the largest number of votes.

For example, suppose that there are three alternatives and five voters. The ballots might look as follows, where each column consists of the set of all alternatives approved of by the corresponding voter. Remember that the ordering within the column is arbitrary (alphabetical in this case); no ranking of alternatives is indicated.

a	a	a	b	c
c	b		c	
	c			

In this example, three voters approve of alternative a, two voters approve of alternative b, and four voters approve of alternative c; alternative c is therefore the social choice.

Many professional societies—including the American Mathematical Society, the Mathematical Association of America, and the National Academy of Sciences—use approval voting for some elections. Since 1996, approval voting has been used by the United Nations to elect the Secretary-General; it has also been used in government elections in Pennsylvania, Oregon, Eastern Europe, and the Soviet Union.

Supporters of approval voting argue that it is much easier to understand than some other procedures. It allows individual voters to equally value two or more alternatives unlike the social choice procedures we have looked at previously which do not allow ties. Because voters essentially need only say yes or no for each alternative, approval voting may be easier for the voters than other procedures which require the voters to rank each alternative. Opponents, however, argue that since approval voting does not use as much information about the voters' preferences, the resulting social choice does not as accurately reflect the will of the people.

Another major argument in support of approval voting is that it will reduce negative campaigning. Negative campaigning is more effective in a plurality system, since only first-place votes matter. If a candidate is not a voter's first choice, then it makes no difference whether that alternative is second or last in the voter's opinion. It doesn't matter, therefore, if negative campaigning further lowers a candidate's status in a voter's eyes. When voters can vote for more than one alternative

though, candidates have a major incentive to remain well respected by as many voters as possible since a voter may decide to vote for his top two, three, four, or more candidates. It is quite possible that negative campaigning would therefore decrease with approval voting because negative campaigning is often looked down upon by voters. Note that this argument applies not only to approval voting, but to any system in which candidates can benefit by being high on (even if not on top of) a voter's preference list.

Another argument in support of approval voting is that it eliminates the effect of *spoiler candidates*, candidates who cannot feasibly win an election but sometimes alter the outcome of an election. For example, many Gore supporters in 2000 blamed Ralph Nader voters when George W. Bush was elected. Since there is reason to believe that the majority of Nader voters preferred Gore to Bush, those voters would likely have voted for Gore had Nader not been an alternative. It is possible then that the presence of Nader as a candidate caused Bush to win over Gore. Supporters of third-party candidates often face the difficult dilemma of voting for their true first-choice candidate, or strategically voting for their second-choice candidate since their first choice is unlikely to win. With approval voting, voters have the option of voting for both; they are able to express their support for their desired candidate while preventing that support from throwing the election to their least favorite candidate. Again, it is worth noting that the social choice procedures which use the voter's full ranking of the candidates also reduce the effect of spoiler candidates.

Approval voting allows more flexibility than plurality. Under approval voting, a voter still has the option of voting solely for their first-place alternative, but has the flexibility to vote for more. Opponents of approval voting argue though that this flexibility is a drawback; one can show that depending on where the voters draw the line between approval and disapproval, almost any candidate can win. For example, before an election using approval voting, the president of the Mathematical Association of America issued the following statement to the voters: "Suppose there are three candidates of whom two are outstanding. Suppose the third is a person you believe is not yet ready for office but whom you decide to vote for as a means of encouragement (in addition to voting for your favorite). If enough voters reason that way, you will elect that person now." (L. Gilman, FOCUS). While this may be alarming, it may be reasonable to assume that if the voters understand the system, this situation would not occur. One might argue that if a voter truly believes someone not ready for office, then he or she does not "approve" of and therefore should not vote for that candidate.

1.9 CONCLUSIONS

We began the chapter by looking at the 1980 U.S. Senate race in New York where Alphonse D'Amato defeated Elizabeth Holtzman and Jacob Javits, even though reasonable assumptions suggest that Holtzman could have beaten either D'Amato or Javits in a one-on-one contest. (In terminology from later in the chapter, Holtzman was a Condorcet winner.) This introduction was meant to suggest some potential difficulties in producing a "reasonable" social choice when there are three or more alternatives. In terms of mathematical preliminaries, we introduced the notion of a function and defined a social choice procedure to be a special kind of function where a typical input is a sequence of preference lists and the corresponding output is either a single alternative (the social choice), a collection of alternatives, or the symbol NW indicating no winner.

The chapter introduced six social choice procedures—Condorcet's method, plurality voting, the Borda count, the Hare system, sequential pairwise voting with a fixed agenda, and a dictatorship—and five apparently desirable properties that pertain to such procedures—the always-a-winner condition, the Pareto condition, the Condorcet winner criterion, monotonicity, and independence of irrelevant alternatives. Asking the thirty obvious questions about which procedures satisfy which properties produced both affirmative answers and negative answers. Among the negative answers were some striking results: the Hare procedure fails to satisfy monotonicity, sequential pairwise voting with a fixed agenda fails to satisfy the Pareto condition, and only a dictatorship (among those considered here) satisfied both independence of irrelevant alternatives and the always-a-winner condition.

In **Section 1.7** we gave a concrete preview of some inherent difficulties when dealing with three or more alternatives by proving that it is impossible to find a social choice procedure that satisfies the always-a-winner condition, independence of irrelevant alternatives, and the Condorcet winner criterion. Finally, we concluded in **Section 1.8** with a brief look at approval voting.

EXERCISES

The purpose of the first two exercises is to help the reader gain some familiarity with the idea of a "function from a set X to a set Y" (as defined in **Section 1.1**). For a procedure to be a function from X to Y, it must assign to each object in X and unique object in Y. In each of the following, sets X and Y are specified as is a procedure. Determine if the given procedure is or is not a function from X to Y.

1. Let X and Y both be the set of non-negative integers: $0, 1, 2, 3, \ldots$
   (a) The procedure corresponding to taking the square root of the input.
   (b) The procedure corresponding to doubling the input.
   (c) The procedure that, given input x, outputs y if and only if y is two units away from x on the number line.
   (d) The procedure that, given input x, outputs the number 17.
2. Let X be the set of (finite) non-empty sequences of nonnegative integers and let y be the set of nonnegative integers.
   (a) The procedure that, given a finite sequence, outputs y if and only if y is the seventh term of the sequence.
   (b) The procedure that outputs n if and only if n is twice the length of the sequence.
   (c) The procedure that outputs y if and only if y is greater than the last term of the sequence.
   (d) The procedure that outputs the number 17 regardless of the input.
3. For each of the six social choice procedures described in this chapter, calculate the social choice or social choices resulting from the following sequence of individual preference lists. (For sequential pairwise voting, take the agenda to be *abcde*. For the last procedure,