CHAPTER

1

Social Choice

1.1 INTRODUCTION

In the present chapter we consider the situation wherein a group of voters is collectively trying to choose among several alternatives. When people speak of the area of "social choice," it is typically this context that they have in mind.

In the case where there are only two alternatives, the standard democratic process is to let each person vote for his or her preferred alternative, with the social choice (the "winner") being the alternative receiving the most votes. The situation, however, becomes complicated if there are more than two alternatives. In particular, if we proceed exactly as we did above where we had two alternatives, then we are not taking advantage of some individual comparisons among the several alternatives that could be made.

As a simple example of the kind of complication caused by more than two alternatives, consider the 1980 U.S. Senate race in New York among Alphonse D'Amato (a conservative), Elizabeth Holtzman (a liberal), and Jacob Javits (also a liberal). While we don't have complete information on the "preference orders" of the voters in New York at that time, we can make some reasonable estimates based on exit polls (showing, for example, that Javits's supporters favored Holtzman over D'Amato by a two-to-one margin). At any rate, for the sake of this example, we'll assume that each of the six possible (strict) preference orderings was held by the percentage of voters indicated below.

22%	23%	15%	29%	7%	4%
D	D	H	H	J	J
H	J	D	J	H	D
J	H	J	D	D	H

In addition to reflecting the intuitions confirmed by the exit polls, the above figures yield results coinciding with the known vote tallies of 45% for D'Amato, 44% for Holtzman, and 11% for Javits. The figures in the last two columns reflect the results of the actual exit poll that took place as described above. The middle two columns reflect a similar assumption as to the preference of Javits over D'Amato among Holtzman supporters, although the two-to-one ratio we use was not, to our knowledge, verified by exit polls. The first two columns are based (with no real justification) on the assumption that D'Amato supporters would be roughly evenly split between the two liberal candidates.

With each person voting for his or her top choice, D'Amato emerges as a (close) winner. On the other hand—and this is what is striking—notice that Holtzman would have defeated Javits in a two-person contest 66% to 34%, and she would have defeated D'Amato 51% to 49%. Thus, if we make use of all the information provided by the individual preference rankings, we get conflicting intuitions as to which alternative should reasonably be regarded as the "social choice."

It is precisely this kind of situation that motivates the considerations of the present chapter. The general framework will be as follows. There will be a set A whose elements will be called *alternatives* (or candidates) and typically denoted by a, b, c, etc. There will also be a set P whose elements will be called *people* (or voters) and typically denoted by p_1, p_2, p_3, etc. We shall assume that each person p in P has arranged the alternatives in a list (with no ties) according to preference. As above, we will picture these lists as being vertical with the alternatives displayed from most preferred on top to least preferred on the bottom. Such a list will be called an *individual preference list*, or, for brevity, a *ballot*. A sequence of ballots is called a *profile*. Our concern in this situation will be with so-called social choice procedures, where

a social choice procedure is, intuitively, a fixed "recipe" for choosing an alternative based on the preference orderings of the individuals.

The mathematical notion underlying the concepts to be treated here is a simple one, but one of enormous importance in mathematics. This notion is that of a function. The definition runs as follows:

DEFINITION. Suppose that X and Y are (not necessarily distinct) sets. Then a *function* from X to Y is a procedure that accepts each member of the set X as input and produces, for each such input, a single corresponding output that is a member of the set Y. The set X is called the domain of the function, and we speak of the procedure as being a function on the set X.

A "social choice procedure" is a special kind of function where a typical input is a profile and an output is a single alternative, or a single set of alternatives if we allow ties, or "NW" indicating that there is no winner.

Because of the importance of this notion, we record it here formally as a definition.

DEFINITION. A *social choice procedure* is a function for which a typical input is a sequence of lists (without ties) of some set A (the set of alternatives) and the corresponding output is either an element of A, a subset of A, or "NW."

When discussing social choice procedures, we refer to the output as the "social choice" or "winner" if there is no tie, or the "social choice set" or "those tied for winner" if there is a tie.

In **Section 1.2** we begin with the case of two alternatives and a very elegant result (May's theorem) characterizing majority rule. In **Section 1.3** we will present six examples of social choice procedures. The examples are chosen to represent not only viable alternatives for real-world applications (e.g., the Hare system and the Borda count), but at least one extreme position (dictatorship) that will resurface later in two important theoretical contexts (Arrow's impossibility theorem and the Gibbard-Satterthwaite theorem). **Section 1.4**, on the other hand, introduces five apparently desirable properties (including independence of irrelevant alternatives and the Condorcet winner criterion, which are referred to below) that a given social choice procedure may or may not have. **Sections 1.5** and **1.6** then consider the obvious question: Which of the six social choice procedures satisfy which of the five desirable properties? Positive results are presented in **Section 1.5** while negative results are in **Section 1.6**.

In **Section 1.7** we foreshadow one of the cornerstones of social choice theory—Arrow's Theorem—that will be presented in Chapter 7 with the following impossibility theorem: There is no social choice procedure that satisfies independence of irrelevant alternatives, the Condorcet winner criterion, and always produces a winner. In **Section 1.8** we briefly discuss approval voting.

1.2 MAY'S THEOREM FOR TWO ALTERNATIVES

In this section, we consider social choice procedures in which there are only two alternatives. The most common example of this is an election in which there are two candidates. If one alternative is represented by the letter "*a*" and the other by the letter "*b*," then there are only two possible preference lists (or ballots): the one that has *a* over *b* and the one that has *b* over *a*. We can think of the former preference list as a vote for alternative *a* and the latter as a vote for alternative *b*.

Most people would agree that there is really only one social choice procedure that suggests itself in this two-alternative situation: See which of and *a* and *b* has the most votes and declare it to be the winner. Indeed, this social choice procedure—typically called *majority rule* and formalized in the following definition—seems to be the cornerstone of our idea of democracy.

DEFINITION. *Majority rule* is the social choice procedure for two alternatives in which an alternative is a winner if it appears at the top of at least half of the individual preference lists (equivalently, if at least half of the voters vote for that alternative).

In terms of a mathematical analysis of this two-alternative situation, there are two natural questions that suggest themselves: What properties of majority rule make it a compelling choice for democratic decision-making? Are there other social choice procedures in the

two-alternative case that also satisfy these same desirable properties? Both of these questions are answered by the following elegant theorem of Kenneth May.

THEOREM. (May, 1952) *If the number of people is odd and each election produces a unique winner, then majority rule is the only social choice procedure for two alternatives that satisfies the following three conditions:*

1. *It treats all the voters the same: If any two voters exchange ballots, the outcome of the election is unaffected.*
2. *It treats both alternatives the same: If every voter reverses his or her vote (changing a vote for a to a vote for b and vice-versa), then the election outcome is reversed as well.*
3. *It is monotone: If some voter were to change his or her ballot from a vote for the loser to a vote for the winner, then the election outcome would be unchanged.*

We do not give a proof of May's theorem here, but a more general version is proved in Chapter 7. Condition (1) in May's theorem is called *anonymity* and condition (2) is called *neutrality*. In many ways, May's theorem tells us that for two alternatives, the search for a perfect voting system is really quite easy. Alas, things change dramatically when we move to the case of three or more alternatives, as we will now see.

1.3 SIX EXAMPLES OF SOCIAL CHOICE PROCEDURES

We describe in this section six examples of social choice procedures. We have tried to pick a variety that includes some that are well known and often used, some that are inherently interesting, some that illustrate the desirable properties of the next section, and a final one (dictatorship) to illustrate that such procedures need not correspond to democratic choices (although many of the properties introduced in the next section will arise as attempts to isolate exactly such democratic choices). The examples are as follows.

Social Choice Procedure 1: Condorcet's Method

The social choice procedure known as Condorcet's method tries to take advantage of the success enjoyed by majority rule when there are only two alternatives. It does this by seeking an alternative that would, on the basis of the individual preference lists, defeat (or tie) every other alternative if the election had been between these two alternatives. Thus, with Condorcet's method, an alternative x is among the winners if for every other alternative y, at least half the voters rank x over y on their ballots. Although this method is typically attributed to the Marquis de Condorcet (1743–1794), it dates back at least to Ramon Llull in the thirteenth century.

To illustrate this idea of one-on-one competitions, suppose the preference lists are:

c	b	b	a	c
b	a	c	b	a
a	c	a	c	b

Then b defeats a in a one-on-one contest by a score of 3 to 2, since the first three voters rank b over a, while the last two voters rank a over b. The reader can also check that b defeats c by a score of 3 to 2, and that c defeats a by a score of 3 to 2.

Because b defeats each of the other alternatives in a one-on-one contest, it is the (unique) winner for this profile when Condorcet's method is used.

Social Choice Procedure 2: Plurality Voting

Plurality voting is the social choice procedure that most directly generalizes the idea of simple majority vote from the easy case of two alternatives to the complicated case of three or more alternatives. The idea is simply to declare as the social choice(s) the alternative(s) with the largest number of first-place rankings in the individual preference lists.

Social Choice Procedure 3: The Borda Count

First popularized by Jean-Charles de Borda in 1781, the social choice procedure known as the Borda count takes advantage of the information regarding individual intensity of preference provided by

looking at how high up in the preference list of an individual a given alternative occurs. More precisely, one uses each preference list to award "points" to each of n alternatives as follows: the alternative at the bottom of the list gets zero points, the alternative at the next to the bottom spot gets one point, the next one up gets two points and so on up to the top alternative on the list which gets $n - 1$ points. For each alternative, we simply add the points awarded it from each of the individual preference lists. The alternative(s) with the highest "Borda score" is declared to be the social choice.

Social Choice Procedure 4: The Hare System

The social choice procedure known as the Hare procedure was introduced by Thomas Hare in 1861, and is also known by names such as the "single transferable vote system" or "instant runoff voting." In 1862, John Stuart Mill spoke of it as being "among the greatest improvements yet made in the theory and practice of government." Today, it is used to elect public officials in Australia, Malta, the Republic of Ireland, and Northern Ireland.

The Hare system is based on the idea of arriving at a social choice by successive deletions of less desirable alternatives. More precisely, the procedure is as follows. We begin by deleting the alternative or alternatives occurring on top of the fewest lists. At this stage we have lists that are at least one alternative shorter than that with which we started. Now, we simply repeat this process of deleting the least desirable alternative or alternatives (as measured by the number of lists on top of which it, or they, appear). The alternative(s) deleted last is declared the winner.

Notice that if, at any stage, some alternative occurs at the top of more than half the lists, then that alternative will turn out to be the unique winner. However, an alternative occurring at the top of exactly half the lists—even if it is the only one doing so—is not necessarily the unique winner (although it must be among the winners).

Social Choice Procedure 5: Sequential Pairwise Voting with a Fixed Agenda

One typically thinks of an agenda as the collection of things to be discussed or decided upon. In the context of social choice theory, however, the term *agenda* refers to the *order* in which a fixed set of things will be discussed or decided upon. Thus, when we speak of a "fixed agenda," we are assuming we have a specified ordering $a, b, c, \ldots$ of the alternatives. (This ordering should not be confused with any of the individual preference orderings.) Sequential pairwise voting can be thought of as a series of one-on-one competitions among the alternatives as in Condorcet's method.

The procedure known as sequential pairwise voting with a fixed agenda runs as follows. We have a fixed ordering of the alternatives $a, b, c, \ldots$ called the agenda. The first alternative in the ordering is pitted against the second in a one-on-one contest. The winning alternative (or both, if there is a tie) is then pitted against the third alternative in the list in a one-on-one contest. An alternative is deleted at the end of any round in which it loses a one-on-one contest. The process is continued along the agenda until the "survivors" have finally met the last alternative in the agenda. Those remaining at the end are declared to be the social choices.

Social Choice Procedure 6: A Dictatorship

Of the six examples of social choice procedures we'll have at hand, this is the easiest to describe. Choose one of the "people" p and call this person the dictator. The procedure now runs as follows. Given the sequence of individual preference lists, we simply ignore all the lists except that of the dictator p. The alternative on top of p's list is now declared to be the social choice.

We shall illustrate the six social choice procedures with a single example that is somewhat enlightening in its own right.

Example:

Suppose we have five alternatives a, b, c, d, and e, and seven people who have individual preference lists as follows:

a	a	a	c	c	b	e
b	d	d	b	d	c	c
c	b	b	d	b	d	d
d	e	e	e	a	a	b
e	c	c	a	e	e	a

For each of our six procedures, we shall calculate what the resulting social choice is.

Condorcet's method: If we look at a one-on-one contest between alternatives a and b, we see that a occurs over b on the first three ballots and b occurs over a on the last four ballots. Thus, alternative b would defeat alternative a by a vote of 4 to 3 if they were pitted against each other. Similarly, alternative b would defeat alternative c (4 to 3, again) and alternative e (6 to 1). Thus, we have so far determined that neither a nor c nor e is the winner with Condorcet's method. But alternative d would defeat alternative b by a score of 4 to 3, and so b is not a winner either. This leaves only alternative d as a possibility for a winner. But alternative c handily defeats alternative d (5 to 2) and so d is also a non-winner. Hence, there is no winner with Condorcet's method.

Plurality: Since a occurs at the top of the most lists (three), it is the social choice when the plurality method is used.

Borda count: One way to find the Borda winner is to actually make a vertical column of values 4, 3, 2, 1, 0 to the left of the preference rankings. (Another way is to count the number of symbols occurring below the alternative whose Borda score is being calculated.) For example, alternative a receives a total of 14 points in the Borda system: four each for being in first place on the first three lists, none for being in last place on the fourth list and the seventh list, and one each for being in next to last place on the fifth and sixth lists. (Or, scanning the columns from left to right, we see that the number of symbols below a is $4 + 4 + 4 + 0 + 1 + 1 + 0$.) Similar calculations, again left for the reader, show that b gets 17 points, c and d each gets 16 points, and e gets only 7 points. Thus, the social choice is b when the Borda count is used.

Hare system: We decide which alternative occurs on the top of the fewest lists and delete it from all the lists. Since d is the only alternative not occurring at the top of any list, it is deleted from each list leaving the following:

a	a	a	c	c	b	e
b	b	b	b	b	c	c
c	e	e	e	a	a	b
e	c	c	a	e	e	a

Here, b and e are tied, each appearing on top of a single list, and so we now delete both of these from each list leaving the following:

a	a	a	c	c	c	c
c	c	c	a	a	a	a

Now, a occurs on top of only three of the seven lists, and thus is eliminated. Hence, c is the social choice when the Hare system is used.

Sequential pairwise voting with a fixed agenda $a\ b\ c\ d\ e$: We begin by pairing a against b in a one-on-one contest. Since b occurs higher up than a on a total of four of the seven lists (the last four), a is eliminated (having lost this one-on-one contest to b by a score of 4 to 3). Now, b goes against c and again emerges victorious by a score of 4 to 3, and so c is eliminated. Alternative b now takes on d, but winds up losing this one-on-one by a score of 4 to 3. Thus, b is eliminated and the final round pits d against e, which the reader can check is an easy win for d. Thus, d emerges as the social choice under sequential pairwise voting with this particular fixed agenda.

A dictatorship: We could pick any one of the seven people to be the dictator, but let's make it person number seven. Then the social choice is simply the alternative on top of the last list, which is e in this case.

Thus, our six examples of social choice procedures yield six different results when confronted by these particular preference lists. This raises the question of whether some procedures might be strictly better than others. But better in what ways? This we investigate in the next section.

1.4 FIVE DESIRABLE PROPERTIES OF SOCIAL CHOICE PROCEDURES

The phrase *social choice* suggests that we are primarily interested in procedures that will select alternatives in a way that reflects, in some sense, the will of the people. A meaningful comparison of different procedures will require our having at hand some properties that are, at least intuitively, desirable. The social choice theory literature is not at all lacking in this regard. We shall, however, limit ourselves to the introduction of five such properties; more are introduced in the exercises. It should be noted that our choice of which properties to consider

has been influenced, at least in part, by a desire to provide familiarity with some of the important ideas underlying major theorems in Chapter 7. In particular, a version of the property called "independence of irrelevant alternatives" will play a key role in Arrow's impossibility theorem.

The five properties are the following.

The Always-A-Winner Condition (AAW)

A social choice procedure is said to satisfy the *always-a-winner condition (AAW)* if every sequence of individual preference lists produces at least one winner.

The Condorcet Winner Criterion

An alternative x is said to be a *Condorcet winner* if it is the unique winner when Condorcet's method is used. Thus, x is a Condorcet winner provided that for every other alternative y, one finds x occurring above y on strictly more than half the lists. This defines what we mean by a Condorcet winner. For the definition of the "Condorcet winner criterion," we have the following:

A social choice procedure is said to satisfy the *Condorcet winner criterion* (CWC) provided that—if there is a Condorcet winner—then it alone is the social choice.

A sequence of preference lists often will not have a Condorcet winner, as we saw in the example in the last section. For those sequences of preference lists that do have a Condorcet winner, it always turns out to be unique; the Condorcet winner criterion is saying that, in this case, the unique Condorcet winner should be the unique winner produced by the social choice procedure. We should also point out that there are weaker versions of the Condorcet winner criterion that have been considered in the literature; see Fishburn (1973) or Nurmi (1987).

The Pareto Condition

A social choice procedure is said to satisfy the *Pareto condition* (or sometimes, for brevity, just *Pareto*) if the following holds for every pair x and y of alternatives:

If everyone prefers x to y, then y is not a social choice.

If we were not allowing ties, we could have said "the" social choice instead of "a" social choice in the statement of the Pareto condition (named after economist Vilfredo Pareto, who lived during the early part of the twentieth century). With ties, however, what we are saying is that if everyone finds x strictly preferable to y (recall that we are not allowing ties in the individual preference lists), then alternative y should not be the social choice and should not even be among the social choices if there is a tie.

Monotonicity

A social choice procedure is said to be *monotone* (or *monotonic*) provided that the following holds for every alternative x:

If x is the social choice (or tied for such) and someone changes his or her preference list by moving x up one spot (that is, exchanging x's position with that of the alternative immediately above x on his or her list), then x should still be the social choice (or tied for such).

The intuition behind the monotonicity condition is that if x is the social choice and someone changes his or her list in a way that is favorable to x (but not favorable to any other alternative) then x should remain the social choice. Monotonicity has also been called "nonperversity" in the literature. Indeed, a social choice procedure that is not monotone might well be regarded as perverse.

Independence of Irrelevant Alternatives

A social choice procedure is said to satisfy the condition of *independence of irrelevant alternatives* (IIA) provided that the following holds for every pair of alternatives x and y:

If the social choice set includes x but not y, and one or more voters change their preferences, but no one changes his or her mind about whether x is preferred to y or y to x, then the social choice set should not change so as to include y.

The point here is that if a preference list is changed but the relative positions of x and y to each other are not changed, then the new list can be described as arising from upward and downward shifts of alternatives other than x and y. Changing preferences toward these other alternatives should, intuitively, be irrelevant to the question of social preference of x to y or y to x.

Of course, if we start with x a winner and y a nonwinner, and people move some other alternative z around, then we cannot hope to conclude that x is still a winner. After all, everyone may have moved z to the top of their list. Independence of irrelevant alternatives is simply saying that y should remain a nonwinner.

To feel comfortable with these properties, one needs to see some specific social choice procedures that provide illustrations of the properties themselves and—perhaps more importantly—examples of their failure. This occurs in the next two sections.

1.5 POSITIVE RESULTS—PROOFS

From the previous two sections we have at hand six social choice procedures (Condorcet's method, purality, Borda, Hare, sequential pairwise, and dictatorship) and five properties (always a winner, the Condorcet winner criterion, Pareto, monotonicity, and independence of irrelevant alternatives) pertaining to such procedures. Which procedures satisfy which properties? The answer is given in the following table (where a "yes" indicates the property holds for the given procedure).

	AAW	CWC	Pareto	Mono	IIA
Condorcet		Yes	Yes	Yes	Yes
Plurality	Yes		Yes	Yes	
Borda	Yes		Yes	Yes	
Hare	Yes		Yes		
Seq Pairs	Yes	Yes		Yes	
Dictator	Yes		Yes	Yes	Yes

Our goal in this section is to prove the nineteen positive results in the chart. The first five positive results—that all but Condorcet's method always produce at least one winner—are collected together in Proposition 1 and treated in a somewhat dismissive manner. Each of the other results will be stated as a proposition and provided with a complete proof that emphasizes the structural aspects of the definitions of the properties. That is, we clearly indicate that we are dealing with an arbitrary sequence of preference lists, that we are making