PART ONE

1

An Introduction to Social Choice Theory

1.1 Some Intuitions, Terminology, and an Example

In a capitalist democracy there are, according to Nobel Laureate Kenneth J. Arrow (Arrow, 1950), "essentially two ways by which social choices can be made: voting, typically used to make 'political' decisions, and the market mechanism, typically used to make 'economic' decisions." Our concern here is exclusively with the former.

Thus, for us, democratic theory is, in the words of Peter C. Fishburn (Fishburn, 1973, p. 3), "based on the premise that the resolution of a matter of social policy, group choice, or collective action should be based on the preferences of the individuals in the society, group, or collective." And social choice theory is, as William H. Riker put it (Riker, 1986, p. xi), "the description and analysis of the way that the preferences of individual members of a group are amalgamated into a decision of the group as a whole." Arrow, by the way, is an economist, Fishburn a mathematician, and Riker a political scientist.

Let's start with a very simple example. Suppose we have an academic department with ten faculty members, one of whom is serving as chair. They are in the process of filling a position in the department and have interviewed five finalists for the job. Needless to say, the different department members disagree on the ranking of the five, and what is needed is some procedure for passing from the preferences of the individuals in the department to the "preferences," if you will, of the group.

First, let's ask what the ballots could look like. If we were to opt for simplicity, a ballot would have just a single name on it, representing (we presume, perhaps naïvely) the candidate who is that department member's top choice. Or, we could allow a ballot to contain several names, intuitively representing either a group that this department member feels is tied for the top, or those candidates that the department member finds acceptable ("approves of" in the parlance of the voting system called approval voting, which we discuss later).

But none of these ballot types is very expressive, and, in voting situations such as we are describing, one tends to use ballots that allow each department member to rank-order the candidates from best to worst, in his or her opinion, perhaps allowing ties (representing indifference) in the individual ballots and perhaps not.[1] Intuitively, this description of a ballot is fine, but some of what we do in this book requires a bit more precision. So let us momentarily set this example to one side and introduce some notation and terminology.

We make use of the universal quantifier "∀" meaning "for all" and the existential quantifier "∃" meaning "there exists." We do not, however, display or abbreviate the phrase "such that," although it is almost always required in reading an expression such as

$$\forall x \in \mathrm{A}\, \exists y \in \mathrm{A}\; x\mathrm{P}y.$$

We also use the standard abbreviation "iff" for "if and only if."

Set-theoretically, $|\mathrm{A}|$ denotes the number of elements in the finite set A, and $\wp(\mathrm{A})$ is collection of all subsets of A. If n is a positive integer, then $[\mathrm{A}]^n = \{\mathrm{X} \in \wp(\mathrm{A}): |\mathrm{A}| = n\}$. Any subset R of $\mathrm{A} \times \mathrm{A}$ is a *binary relation* on A, and in this case we write "$a\mathrm{R}b$" or we say "$a\mathrm{R}b$ holds" to indicate that $(a, b) \in \mathrm{R}$, and we write "$\neg(a\mathrm{R}b)$" or say "$a\mathrm{R}b$ fails" to indicate that $(a, b) \notin \mathrm{R}$. Finally, if R is a binary relation on A and $v \subseteq \mathrm{A}$, then the *restriction of R to v*, denoted $\mathrm{R}|v$, is the binary relation on v given by $\mathrm{R}|v = \mathrm{R} \cap (v \times v)$.

The binary relations we are most concerned with satisfy one or more of the following properties.

[1] Most readers will assume that the picture we are painting is one in which the notion of indifference is transitive, and we will, in fact, be adopting that convention. Fishburn (1973, pp. 5 and 6), on the other hand, spends considerable time with the case in which indifference is intransitive, and justifiably so. As an example of intransitive indifference in our present context, suppose that Applicant B might receive (and bring along) a large research grant for which he or she has applied, and that we in the department will not know whether or not this grant application is successful before the job offer will be made. We could handle this by pretending to have six applicants instead of five, with Applicant B split into "Applicant B without research support" and "Applicant B with research support." It is now easy to imagine a situation in which a department member might be indifferent between Applicant C and either of these choices. But of course, anyone would (presumably) prefer Applicant B with research support to Applicant B without research support.

Definition 1.1.1. A binary relation R on a set A is:

reflexive	if	$\forall\, x \in A,\ xRx$
irreflexive	if	$\forall\, x \in A,\ \neg(xRx)$
symmetric	if	$\forall\, x, y \in A$, if xRy, then yRx
asymmetric	if	$\forall\, x, y \in A$, if xRy, then $\neg(yRx)$
antisymmetric	if	$\forall\, x, y \in A$, if xRy and yRx, then $x = y$
transitive	if	$\forall\, x, y, z \in A$, if xRy and yRz, then xRz
complete	if	$\forall\, x, y \in A$, either xRy or yRx

Definition 1.1.2. A binary relation R on a set A is a *weak ordering* (*of A*) if it is transitive and complete and a *linear ordering* (*of A*) if it is also antisymmetric.

If R is a weak ordering of A, then the completeness of R implies (letting $x = y$) that R is also reflexive. Intuitively, a weak ordering corresponds to a list with ties, with xRy being thought of as asserting that x is at least as good as y. A linear ordering corresponds to a list without ties, with xRy now being thought of as asserting that either $x = y$ or x is strictly better than y.

Associated to each weak ordering R of A, there are two so-called derived relations P and I.

Definition 1.1.3. If R is a weak ordering of A, then the derived relations of *strict preference* P and *indifference* I are arrived at by asserting that xPy iff $\neg(yRx)$ and xIy iff xRy and yRx.

If R is a linear ordering of A, then the derived relation I is just equality, and the derived relation P is referred to as a *strict linear ordering of A*. Exercise 1 asks for verification that if R is a weak ordering, then the derived relation P of strict preference is transitive and asymmetric (and thus irreflexive), while the derived relation I of indifference is reflexive, symmetric, and transitive. Thus, I is an equivalence relation, and P is a strict linear ordering of the I-equivalence classes.

The following definition uses the concepts of weak and linear orderings to formalize some election-theoretic terminology.

Definition 1.1.4. If A is a finite non-empty set (which we think of as the set of alternatives from which the voters are choosing), then an *A-ballot* is a weak ordering of A. If, additionally, n is a positive integer (where we think of $N = \{1, \ldots, n\}$ as being the set of voters), then an *(A, n)-profile* is an n-tuple of A-ballots. Similarly, a *linear A-ballot* is a linear ordering of A, and a *linear (A, n)-profile* is an n-tuple of linear A-ballots.

When the set A of alternatives and the integer n are clear from the context, we will use "ballot" in place of "A-ballot" and "profile" in place of "(A, n)-profile." Similarly, we will often use a phrase like "If **P** is a profile" as an abbreviation for the phrase "If **P** is an (A, n)-profile for some set A and some positive integer n." This latter remark is illustrated in the next paragraph.

If **P** is a profile, then R_i denotes its ith component (that is, the ballot of the ith voter), with P_i and I_i denoting the corresponding derived relations of strict preference and indifference for the ith voter. When we have several profiles under consideration at the same time, we use names such as **P**, **P**′, and **P**″, with the understanding that their components and the derived relations also carry the prime, double prime, etc.

If $\mathbf{P} = <R_1, \ldots, R_n>$ is an (A, n)-profile and $X \subseteq A$, then the restriction of **P** to X, denoted $\mathbf{P}|X$, is the profile $<R_1|X, \ldots, R_n|X>$. If $i \in N$, then $\mathbf{P}|N - \{i\}$ is the profile $<R_1, \ldots, R_{i-1}, R_{i+1}, \ldots, R_n>$. This dual use of the vertical bar should cause no confusion.

The following definition collects some additional ballot-theoretic notation we will need.

Definition 1.1.5. Suppose **P** is a linear (A, n)-profile, X is a set of alternatives (that is, $X \subseteq A$), and i is a voter (that is, $i \in N$). Then:

$$
\begin{aligned}
\mathrm{top}_i(\mathbf{P}) = x \quad &\text{iff} \quad \forall y \in A\ xR_iy \\
\max_i(X, \mathbf{P}) = x \quad &\text{iff} \quad x \in X \text{ and } \forall y \in X\ xR_iy \\
\min_i(X, \mathbf{P}) = x \quad &\text{iff} \quad x \in X \text{ and } \forall y \in X\ yR_ix
\end{aligned}
$$

Thus, $\max_i(X, \mathbf{P})$ is the element of X that voter i has most highly ranked on his or her ballot in **P**, and $\min_i(X, \mathbf{P})$ is the one ranked lowest. $\mathrm{Top}_i(\mathbf{P})$ is the alternative that voter i has at the top of his or her ballot in **P**. Hence, $\mathrm{top}_i(\mathbf{P}) = \max_i(A, \mathbf{P})$ and $\max_i(X, \mathbf{P}) = \mathrm{top}_i(\mathbf{P}|X)$.

Once we have decided what the ballots will look like, it might well seem natural to ask what we *do* with these ballots to find a winner. That, however, is somewhat getting ahead of ourselves. What we really need to decide first is what kind of outcome our balloting should yield.

For example, should the outcome in the departmental hiring example that we considered earlier be a single winner with the understanding that the chair will call that candidate with the offer, and if he or she refuses, then the chair will reconvene the department and start the balloting process all over again? Or do we allow ties in the outcome of the balloting, with the understanding, perhaps, that either the chair or the dean will be allowed to break the tie?

From the chair's point of view, perhaps the most desirable outcome is neither of these, but instead a list without ties – a linear ordering – that gives the department's "overall ranking" (an intuitive phrase, at best) of the candidates. The chair can then begin at the top, and make the calls one by one until the position is accepted. Similarly, if the outcome is a list with ties – a weak ordering – then we could once again agree to give either the chair or the dean tie-breaking power.

But suppose we are in a situation wherein we proceed to individually rank-order all the candidates, and then, while the chair is handling the accompanying administrative tasks involving the dean and the college affirmative action officer, several candidates notify the department that they have accepted other offers and no longer wish to be considered. To handle such contingencies, we might want the outcome of the election to be a "choice function" that selects one or more "winners" from each non-empty subset of candidates.

This last possibility is an interesting one, and later in this chapter we say something about the historical perspective from the field of economics that apparently played a role in the prominence of so-called choice functions in the formalism. But for now, we conclude the present section with precise definitions of the different kinds of "social choice procedures" alluded to earlier.[2]

Definition 1.1.6. Suppose that A is a non-empty set, n is a positive integer, and V is a function whose domain is the collection of all (A, n)-profiles.[3] Then V is:

(1) a *resolute voting rule for* (*A,* n) if, for every (A, n)-profile **P**, the election outcome V(**P**) is a single element of A,

[2] There is, unfortunately, no common terminology in the literature for the concepts in Definition 1.1.6. Our use of "voting rule" and "social welfare function" is quite common and used, for example, in Moulin (2003). Our use of "social choice function" for a voting rule with a variable agenda is also not without precedent, but one also sees variants of this with, for example, "decision" used in place of "choice" and/or "procedure" used in place of "function." Our use of "resolute" to mean "without ties" would not, however, be considered standard, although it has appeared in the work of Duggan and Schwartz (1993, 2000) and goes back at least to Gärdenfors (1976).

[3] We are building into our formalism a condition known in the literature as "unrestricted scope." Intuitively, this asserts that no voter should be prohibited from submitting any ballot. There is something to be gained (e.g., conceptual simplicity) by burying certain assumptions within the formalism, but there is also a loss. For example, our choice to build in unrestricted scope leads to an omission of a number of important results related to the following question: What conditions can one impose on a profile that will ensure that "bad things" (e.g., opportunities for manipulation or cycles wherein a majority of voters prefer a to b, a majority prefers b to c, and yet a majority also prefers c to a) don't happen? One answer to this, by the way, involves "single-peaked preferences" – see Black (1958), Sen (1966), Fishburn (1973), Taylor (1995), and Shepsle and Bonchek (1997).

(2) a *voting rule for* (*A,* n) if, for every (A, *n*)-profile **P**, the election outcome V(**P**) is a non-empty[4] subset of A,
(3) a *social choice function for* (*A,* n) if, for every (A, *n*)-profile **P**, the election outcome V(**P**) is a "choice function" C that picks out a non-empty subset C(*v*) of *v* for each non-empty subset *v* of A,
(4) a *resolute social choice function for* (*A*, n) if, for every (A, *n*)-profile **P**, the election outcome V(**P**) is a choice function C that picks out a single element C(*v*) from *v* for each non-empty subset *v* of A,
(5) a *social welfare function for* (*A*, n) if, for every (A, *n*)-profile **P**, the election outcome V(**P**) is a weak ordering of A, and
(6) a *resolute social welfare function for* (*A*, n) if, for every (A, *n*)-profile **P**, the election outcome V(**P**) is a linear ordering of A.

The set *v* of alternatives occurring in (3) and (4) is called an *agenda*.[5] In general, V is called an *aggregation procedure* if it is any one of (1)–(6). As before, we suppress the reference to the pair (A, *n*) whenever possible.

In point of fact, our primary concern is with the first three aggregation procedures given in Definition 1.1.6:

(1) resolute voting rules (Chapter 3 and 7)
(2) (non-resolute) voting rules (Chapters 4 and 8)
(3) social choice functions (Chapters 5 and 8).

With each kind of aggregation procedure, there are two contexts in which we work: the one in which only linear ballots are considered and the other in which we allow ties in the ballots.

But before we press on with any additional notation and terminology, let us pause to give a quick historical overview of the field of social choice theory. This will, at the same time, provide an informal introduction to a number of aggregation procedures, most of which are rigorously defined in Section 1.4.

[4] By saying "non-empty" we are disallowing the possibility of an election resulting in no alternative being chosen. Fishburn (1973, p. 3) justifies this by the observation that one can always include alternatives such as "delay the decision to a later time" or "maintain the status quo."

[5] There are at least three different ways the term "agenda" is used in voting-theoretic contexts: (1) as the set of alternatives from which a choice is to be made (our use here), (2) as an ordering in which alternatives will be pitted against each other in one-one-one contests based on the ballots cast, and (3) as an ordering in which alternatives will be pitted against the status quo until one defeats the status quo.

1.2 A Little History

Jean Charles Chevalier de Borda (1733–99) was, according to Duncan Black (1958, p. 156), "the first thinker to develop a mathematical theory of elections." In Borda's 1781 paper (apparently the only one of his that we now possess) he introduced the aggregation procedure that is known today as the *Borda count*. It selects a winner (or winners) from among k alternatives by assigning each alternative k-1 points for each ballot on which it appears first, k-2 points for each ballot on which it occurs second, and so on. The points are then summed, with the winner (or winners) being the alternative with the most points. If ties in the ballots are allowed, the procedure can be suitably modified. These so-called Borda scores can also be used to produce a list, perhaps with ties, as the outcome of an election. Interestingly, recent historical work by McLean and Urken (1993) and Pukelsheim (unpublished) reveals that Borda's system had been explicitly described in 1433 by Nicholas of Cusa (1401–64), a Renaissance scholar interested in the question of how German kings should be elected.

In that same 1781 paper, Borda pointed out a very nice equivalent version of the Borda count, not often referred to today, that goes as follows: Each alternative is pitted one-on-one against each of the other alternatives, based on the ballots cast. Having done this, one doesn't just look for the alternative that defeats the most other alternatives – this would be quite a different social choice procedure, one known today as *Copeland's function* and introduced in an unpublished 1951 note by A. H. Copeland (Fishburn, 1973, p. 170). Instead, one looks for the alternative with the greatest total score from these one-on-one contests. For example, if one of four alternatives defeats two others by scores of 4–3 and 5–2, but loses to the third by a score of 6–1, then that alternative's total score is $4 + 5 + 1 = 10$.

In fact, this latter characterization of the Borda count gives rise to an easy way to hand-calculate Borda scores given a sequence of ballots: Given an alternative a, one simply counts the total number of occurrences of other alternatives below a, proceeding ballot-by-ballot (Taylor, 1995). It is easy to see that this is the same as Borda's equivalent, the difference being that what we are suggesting here is a ballot-by-ballot enumeration instead of an alternative-by-alternative enumeration.

But Borda was not alone in his election-theoretic ponderings, as a systematic theory of elections was, as Black (1958, p. 156) again informs us, "part of the general uprush of thought in France in the second half of the eighteenth century." For example, Borda's "method of marks" arose again in 1795 in the writings of Pierre-Simon, Marquis de Laplace (1749–1827), who derived the method via some probabilistic considerations.

However, no one's contributions at that time were more important than the observations of Marie Jean Antoine Nicolas Caritat, Marquis de Condorcet (1743–94). In a 1785 publication (Condorcet, 1785) he explicitly discussed what we now call the "Condorcet voting paradox" wherein we find that if a group of voters is broken into three equal-size groups with preferences for three alternatives as shown below, then a majority prefers a to b, a majority prefers b to c, but (somewhat paradoxically) a majority also prefers c to a.

Group #1	Group #2	Group #3
a	b	c
b	c	a
c	a	b

The Condorcet Voting Paradox

For a given sequence of ballots, a candidate that would defeat each of the other candidates in a one-on-one contest – based on the ballots – is called a *Condorcet winner* for that election. For example, if we regard the U.S. presidential race of 2000 in the state of Florida as one with four candidates (Bush, Gore, Nader, and Buchanan), then it is almost certainly true that Gore was the Condorcet winner. Of course, Bush won using what is known as plurality voting, wherein one simply looks for the alternative with the most first-place votes.

The name "Condorcet's method" is often applied to the voting procedure in which the Condorcet winner, if there is one, is the unique winner of the election and there is no winner otherwise. Like Borda, Condorcet was anticipated by several centuries. Indeed, very recent research by Pukelsheim et al. (unpublished) shows that Condorcet's method can be traced back at least to Ramon Llull (1232–1316), a Catalan philosopher and missionary who was involved in devising election schemes for selecting the abbess of a convent. (See Pukelsheim's amusing "Spotlight" on page 418 of COMAP, 2003.)

Condorcet knew of Borda's work and Condorcet pointed out in his writings that Borda's method of marks, like plurality voting, can result in a Condorcet winner not being elected. Nevertheless, even Condorcet would probably have been surprised that this "defect" would one day determine a presidential election, as it did in the United States in the year 2000.

The nineteenth century saw a few small election-theoretic results from people like Issac Todhunter (1820–84), M. W. Crofton (1826–1915), E. J. Nanson (1850–1936), and Francis Galton (1822–1911). But it was the Reverend Charles Lutwidge Dodgson (1832–98) – better known by the pseudonym Lewis Carroll – who made the most significant contributions at the time, beginning with his rediscovery in 1874 of the Condorcet voting paradox. Dodgson was the

Mathematical Lecturer at Christ Church, and he even published a monograph entitled *Elementary Treatise on Determinants* between the appearance of *Alice's Adventures in Wonderland* (1865) and *Through the Looking Glass, and What Alice Found There* (1872). The mathematical biographer E. T. Bell spoke of Dodgson as having in him "the stuff of a great mathematical logician" (Black, 1958, p. 195), and Duncan Black characterized Dodgson's understanding of the theory of elections and committees as "second only to that of Condorcet" (Black, 1958, p. 212).

In an 1873 pamphlet entitled "A Discussion of the Various Methods of Procedure in Conducting Elections" (Black, 1958, p. 214), Dodgson proposes – without claiming to have discovered them himself – several "Methods of Procedure" for the case where an election is necessary. The description of each that follows is taken verbatim from that pamphlet, although we do not reproduce his examples showing why he finds fault with each. Our own comments are added in brackets.

(1) The Method of a Simple Majority: In this Method, each elector names the one candidate he prefers, and he who gets the greatest number of votes is taken as the winner. [This is known today as *plurality voting*.]
(2) The Method of an Absolute Majority: In this Method, each elector names the one candidate he prefers; and if there be an absolute majority for any one candidate, he is the winner. [Dodgson offers no provision for the case where no one has more than half the votes.]
(3) The Method of Elimination, where the names are voted on by two at a time: In this Method, two names are chosen at random and proposed for voting, the loser is struck out from further competition, and the winner taken along with some other candidate, and so on, til there is only one candidate left. [This procedure is essentially what Straffin (1980) calls "sequential pairwise voting with a fixed agenda" (see also Taylor, 1995). Here, "agenda" refers to an ordering of the alternatives.]
(4) The Method of Elimination, where the names are voted on all at once: In this Method, each elector names the one candidate he prefers: the one who gets the fewest votes is excluded from further competition, and the process is repeated. [This is the procedure introduced in 1861 by Thomas Hare, and known today by various names including the "Hare system" and the "single transferable vote system." In 1862, John Stuart Mill (Mill, 1862) spoke of it as being "among the greatest improvements yet made in the theory and practice of government." It is currently used to elect public officials in Australia, Malta, the Republic of Ireland, and Northern Ireland. The Hare system was essentially the method used to choose Sydney, Australia, as the site of the 2000 Summer Olympics. In this election, Beijing would have been the plurality winner, but after Istanbul, Berlin, and Manchester were eliminated (in that order), Sydney defeated Beijing by a vote of 45 to 43.]
(5) The Method of Marks: In this Method, a certain number of marks is fixed, which each elector shall have at his disposal; he may assign them all to one candidate,