Philosophy, Politics, and Economics

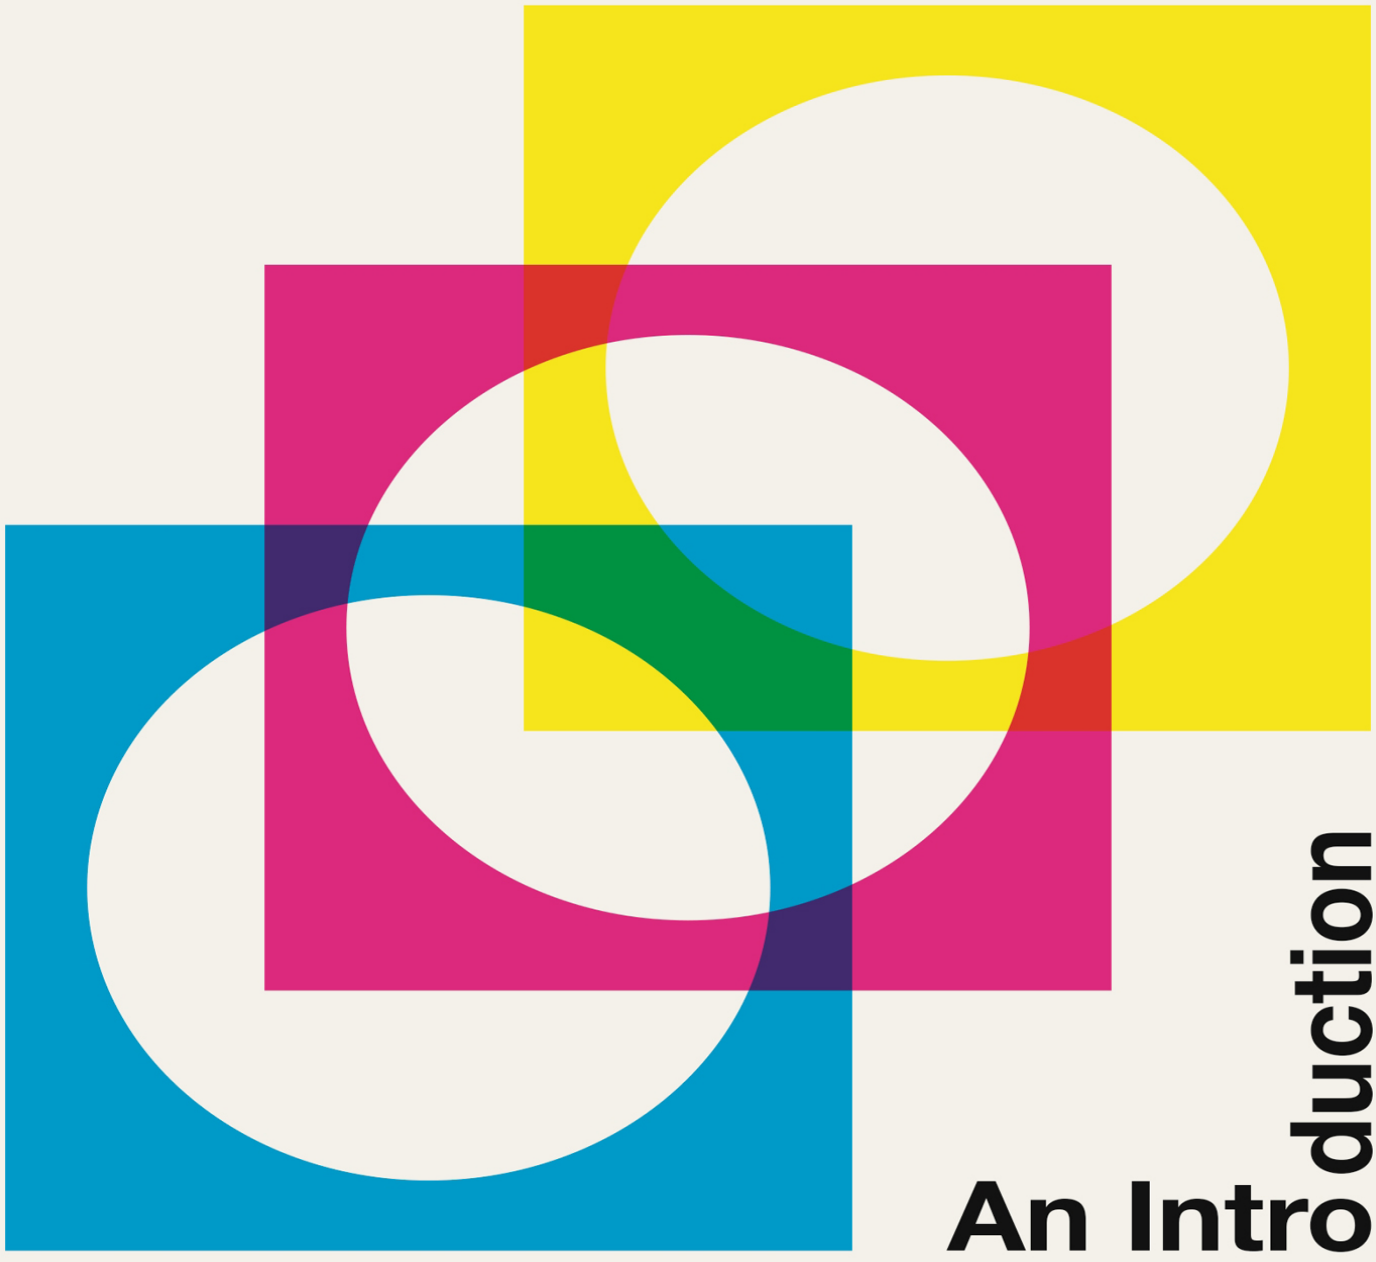

An Introduction

Gerald Gaus and John Thrasher

Philosophy, Politics, and Economics

Philosophy, Politics, and Economics

AN INTRODUCTION

GERALD GAUS

JOHN THRASHER

PRINCETON UNIVERSITY PRESS
PRINCETON *&* OXFORD

8

Social Choice Theory

In the previous chapters, we have focused on how individuals make rational and strategic decisions. Now it's time for us to shift our focus to how individuals with different preferences can (or should) make decisions together. This is an area of mathematics, philosophy, economics, and political science known as social choice theory. At the most general level, there are two main aspects of social choice theory. The first involves theories of preference aggregation, that is, how the preferences that govern the rationality of individual choice can be combined to construct a rational collective or social choice. The second involves theories of collective decision rules, that is, how groups make decisions together. These two topics are related, though distinct, as we will see.

The Problem of Social Choice

Individual rationality, we have seen, can either lead us to efficient outcomes (chapters 2–3) or get us stuck in Pareto-inferior outcomes such as the Prisoner's Dilemma (chapter 5). We also saw that efficiency can sometimes conflict with rights and other values as well as problems of externalities and public goods (chapter 4). All of these considerations have led political theorists to justify government as a coercive organization that can move society to valuable or worthy social states that cannot be brought about by market transactions. Sometimes the problem is that the market has failed to yield efficient results—government, it is said, is necessary to move us to social states in which we enjoy public goods and regulate negative externalities (given non–zero transaction costs). At other times government

pursues values in addition to efficiency; it seeks a fairer distribution of resources. Governments also define and enforce rights, including property rights, which are essential for markets.

This raises the question of what values governments should pursue. If we all agreed that government is to supply exactly the same set of public goods and the precise amounts, there would be no great problem. If we look at a 2 × 2 matrix of the Prisoner's Dilemma, it is obvious what we would want government to do—get us to the cooperate/cooperate payoff, by changing the game so that defection is no longer the dominant strategy. Clearly, though, this is to ignore the very fact on which economic analysis builds: our orderings of outcomes, which is to say our preferences, usually differ. In market transactions we benefit from the diversity of preferences. It is because you order the outcomes differently than I do that we can engage in mutually beneficial trade—Alf prefers a chicken wing to another slice of pizza, and Betty prefers a slice of pizza to a chicken wing, so they trade. But diversity of preferences is much more of a challenge in the political context since, in most cases, when the government acts, we must all do the same thing regardless of our underlying preferences.

There are two very different responses to this problem. We can call them —very broadly—*preference respecting* and *preference transforming* views of the collective choice. On a *preference respecting view*, we suppose that each person has an ordering over possible political outcomes that expresses their overall set of normative criteria. As we have seen, utility functions can accommodate all the normative criteria relevant to choice; we have been at pains to argue that they need not be about what promotes your self-interest, or wealth, or whatever. So long as the preference orderings fit certain requirements, they can be over anything. On this view, the basic data of politics should be the same as the basic data of economics since people's orderings of the outcomes capture all that they see as normatively important for making a choice. The difference between the market and the political forum is that in the market each person makes her own choice based on her preferences and option set, while in the political setting we come together with our individual preference orderings and seek to make a uniform collective choice. If this is true, there is no reason, in principle, why we can't establish criteria of rational collective choice as we did for individual choice.

The *preference transformative* view, however, insists that this is the wrong conception of the political. As Jon Elster understands it, the preference respecting view embodies a confusion between the kind of behavior that is appropriate in the marketplace and that which is appropriate in the forum.

> [In the market] the notion of consumer sovereignty is acceptable because, and to the extent that, the consumer chooses between courses of action that differ only in the ways it affects him. In political choice situations, however, the citizen is asked to express his preference over states that also differ in the way that they affect other people.… [T]he task of politics is not only to eliminate inefficiency but also create justice.[1]

We should clearly distinguish two lines of criticism of the preference respecting approach suggested by Elster's remarks. First, Elster insists that political choice is not simply about efficiency but about justice, which is why aggregating individual preferences into a social choice is inappropriate. While this may be true if we suppose that political agents are simply *Homo economicus*, we have seen that utility functions can include the entire range of normative considerations, including notions of justice. So, the mere idea of starting out with individual preference orderings to determine a social choice about what is just is not, as Elster puts it, "incongruous."[2]

The second criticism, though, has considerably more bite. This is the idea that once we are concerned with collective choice, we can't simply rest, content with the idea that a person's preferences are his own business. Recall Prude and Lewd from chapter 4. Each has a preference about what books the other is to read: should collective political choice take these preferences over what other people do—external preferences—into account? The fact that politics seems to require some people having a say over what other people do, it looks like external preferences of this sort will be a central concern of collective, political choice. This is a clear difference between social choice and market choice, though, since one's choice over what goods to buy or sell doesn't directly concern the actions or choices of other people.

In response to these worries, preference transformative views see politics as a way of coming together in the public forum, arguing, and changing each other's minds. It is not enough to simply assert that one prefers this to that in the public sphere—preferences that were not formed by a properly reflective

process, or that include some sorts of external preferences, or preferences that could not be publicly revealed (say, racist preferences), should be transformed or eliminated in reasoned political discussion.

There is clearly some truth in the transformative view since we do not think that discriminatory external preferences should be the basis of a rational social choice. Still, the contrast between the market and the forum is not quite so sharp as advocates of the transformative view suggest. For one thing, we have seen that all desires or goals can be criticized as not sufficiently reflective or formed by non-autonomous processes. There is no reason intrinsic to market analysis to prevent arguing that some desires, say those for especially dangerous or noxious activities, should not be met by the market. Indeed, Cass Sunstein, an influential legal theorist, has offered a criticism of many market outcomes because they satisfy preferences that ought not to be respected.[3] Even in the market, some sorts of external preferences are a worry.

Another reason that the contrast between preference respecting and preference transformative views of politics is not as clear-cut as one might think is that, even after debate and transformation in the forum, we can still expect an extremely wide range of preference orderings about what is the best social choice. To be sure, some contemporary political theorists, known as “deliberative democrats,” seem to hope for significant consensus on political issues, but this does not look terribly plausible in extensive and diverse societies that make up virtually all democratic nations. Even after deliberation and discussion we are apt to be left with a diverse set of preferences which, presumably, the social choice mechanism then will have to respect, thus leading us back to the preference respecting view of the forum.

In this chapter, we explore political choice on the preference respecting view. If one upholds the pure preference respecting view of politics—Elster’s “market” view—then the following analysis should be of great interest. Insofar as one endorses the transformative view of the political, some of the problems explored in this chapter will be of less concern, as some of the worries about collective choice will stem from allowing citizens to hold any ordering over outcomes that their normative criteria recommend and respecting these orderings in the social choice process. However, even for the transformative theorist, many of the problems we discuss—such as

path dependency and strategic voting—must remain a deep concern. As will questions about voting systems, generally.

Social Choice Theory

According to Thomas Jefferson, self-government requires following a collection of wills expressed by the majority, usually through voting.[4] Thinking of democracy as taking a collection of wills (individual preference orderings) and generating a social will (social preference ordering) is a long-standing and, in many ways, attractive idea. We might think of this "democratic collecting of wills" in two ways.

The most ambitious approach is to aim at a *social preference ordering*. It seems safe to suppose that we would want our social preference ordering to conform to the axioms of ordinal utility such as: asymmetry of strict (social) preference, symmetry of (social) indifference, reflexivity, transitivity, and completeness all seem as important for a social ordering as they were for an individual ordering. Let us call a social outcome with these properties a *social ordering*. We can call such social ordering a social welfare function (SWF). A social welfare function aggregates individual rational ordering into a collected social ordering. We can think of a social ordering as a construction of a social will out of a collection of individual wills.

Social Choice

A *social welfare function* (SWF) takes as *inputs* individual preference orderings and yields as *outputs* a social ordering.

A *collective choice rule* (CCR) takes as inputs individual preferences and yields a best option as an *output*.

Less ambitiously, we might aim only at a *collective choice rule* (CCR) that, from a set of options, selects the best or optimal choice. A CCR takes as inputs people's individual preference orderings and yields a socially best option. This is more akin to a formal conception of voting wherein there is some set of social options, some way of representing preferences over those outcomes (such as voting), and a rule (or set of rules) for determining how those preferences determine the optimal option. So, each individual will be able to order all of the available options according to the traditional

preference binary relations, and some rule will transform those orderings into a collective decision on the "best" option as an output.

In either case, social choice theory has three basic building blocks:

1. Some set of individuals $\{i,j,\ldots\}$
2. Some set of options $\{x,y,z\ldots\}$
3. Binary social ordering relations R,I,P

Individual and Social Orderings Notation

Individual Orderings

For some individual i ($\exists i$)

- $x \succcurlyeq y$ x is at least as good as y
- $x \succ y$ x is strictly preferred to y
- $x \sim y$ x is indifferent to y

Social Orderings

- xRy x is socially at least as good as y
- xPy x is socially strictly preferred to y
- xIy x is socially indifferent to y

$x \rightarrow y$ x implies y

$x \& y$ x and y

$\neg x$ not x

We will use the binary relations ($\succcurlyeq$, $\succ$, $\sim$) for individual orderings of preference, strict preference, and indifference respectively and the binary relations (R,P,I) for the same orderings at the social level.

Collective Rationality

The idea of a social welfare function raises a fundamental question: does it make sense to think of *rational* social choice at all? That is, can we combine individual rational preferences—orderings that meet the conditions of reflexivity, symmetry of indifference, asymmetry of strict preference, and transitivity—and combine them in ways that lead to a rational social preference?

If so, then just as we might be able to identify a rational individual with a consistent chooser who can be represented by a utility function, we can think of a rational society as an agent that makes consistent choices and has a utility function that it maximizes. In that sense we can say that a society can be represented as making rational choices. If not, then we might have to face

the possibility that while we make collective choices, these choices cannot or need not meet the standards of rationality. Either way we go has implications for how much normative weight we should give to the output of these procedures and how we should interpret their meaning.

Some famous and thoughtful economists find the very idea of collective choice and collective rationality a philosophically objectionable notion. According to James Buchanan:

> Rationality or irrationality as an attribute of the social group implies the imputation to that group of an organic existence apart from that of its individual components. If the social group is so considered, questions may be raised relative to the wisdom or unwisdom of this organic being.[5]

According to Buchanan, we should not expect society to display a notion of rationality akin to individuals, so the idea of social choice theory is in many ways misconceived. Buchanan was wary of any conception of social choice for this reason. There is no "social" choice as such. Instead, individuals decide in groups what to do. Democracy is not a representation of a "general will" but a type of group decision that needs to be studied and understood as the representation of individual wills, not a construction of a social one. We will look more carefully at Buchanan's positive views along these lines in the next chapter.

May's Theorem and Majority Rule

One of the enduring questions for political philosophy is whether there is anything special about majority rule. Is choice by a rule "50% + 1 wins" arbitrary? Would 51% + 1 do just as well? How about 55%? Many are deeply attracted to what Brian Barry called "majority principle" as a fundamental aspect of democratic theory with some seeing it as the sine qua non of democratic governance.[6]

An important 1952 theorem by Kenneth May showed that majority rule is indeed special.[7] The power of May's proof, which concerns a collective choice rule (CCR) between two options (x,y), depends on the acceptability or reasonableness of the conditions that he imposes on any acceptable CCR.[8]

A word on this method of axiomatic social choice theory. When we are trying to decide on whether some social choice procedure C is desirable, we

might pursue two different methods. *Output* evaluation would try to determine the sorts of decisions *C* yields, and then evaluate whether we approve of these decisions, and so approve of *C*. So, for example, a group of friends might decide that Alf makes the best decisions about where to go for pizza and so they endorse the rule, "When we are trying to decide where to go for pizza, Alf makes the decision." So here they are reverse engineering their evaluation: if they approve of the outputs, they approve of the procedure.

Another way is to independently evaluate the features of the procedure. Call this a *direct procedural evaluation*. This method asks, "what features would we like our procedure to have?" We then list the characteristics that we think are necessary in a procedure and try to determine whether there are any procedures that meet these criteria, and if there are, what they look like. This is the method of social choice theory: first determine the features that are required in a procedure to aggregate individual preferences into a social preference, and then see what procedures these yield.

May's Conditions

What, then, are some features that would be desirable in a procedure to take individual preferences and yield social preference? Since our interest is in majority rule, we assume that we are designing a voting procedure. May proposes four conditions that must be met for an acceptable voting rule.

(1) **Decisiveness**—The CCR must yield a decision between *x* and *y* for all possible configurations of individual preferences, such that *xPy*, *yPx*, or *xIy*. All possible permutations of *x* and *y* in the individual preference orderings must yield a single definite social choice. We cannot have [(*xPy*) & (*yPx*)] or [(*xPy*) & (*xIy*)] or [(*yPx*) & (*yIx*)].[9]

(2) **Anonymity**—The identity of the voters must have no effect whether *xPy*, *yPx*, or *xIy*. For a given preference profile among all voters, that preference profile will always yield the same decision when the identities of the voters are permutated. If the social choice mechanism is anonymous, it doesn't matter what names are attached to the votes: the CCR counts the votes without paying attention to the voter's identities. So, from the social perspective, when [Alf holds $x \succcurlyeq y$ and Betty holds $y \succcurlyeq x$] it is the same as when [Alf holds *y*

$\succcurlyeq x$ and Betty holds $x \succcurlyeq y$]. In both cases, there is one vote for *x* and one vote for *y*.

(3) **Neutrality**—Here is a somewhat more complicated condition. We'll first state it, and then try to explain it and its rationale. If for a given profile of individual preferences, *xPy* and, keeping that profile, if *x* is relabeled as *w*, and *y* is relabeled as *z*, then it must be the case that *wPz*.

This condition is effectively anonymity for options. The idea is that a collective choice rule should be neutral between alternatives insofar as neither alternative has a built-in advantage, so the labels attached to the proposal don't matter. For example, a rule that a 2/3 decision is required to pass a proposal, otherwise it fails, violates neutrality: it is biased in favor of the status quo. Say that there are two proposals:

x = "Change the constitution: abolish the electoral college and elect the president by popular vote by majority vote"

y = "Do not change the constitution: keep the electoral college and do not elect the president by popular vote by majority vote"

Say *x* gets 51% of the vote; even so, it would not pass. This is because a supermajority is required to change the US Constitution, so *yPx*. Now relabel *x* as *w*, where *w* = "abolish the electoral college and elect the president by majority vote" and relabel *y* as *z*, where *z* = "keep the electoral college and do not elect the president by majority vote." Having eliminated the label in *x* that the proposal changes the Constitution, *w* now wins (51%–49%), and so we have *wPz*, and neutrality is violated.

The rationale behind neutrality is that we do not want our social decision-making procedure to work to the advantage of some options ("no change from the status quo") and to the disadvantage of others ("change the status quo"). Under many voting rules, the decision-making rule favors those people with conservative preferences and places an extra obstacle in front of those with reformist preferences. Neutrality forbids this.

(4) **Strong Monotonicity**—This includes two conditions: *non-negative responsiveness and positive responsiveness*. According to *non-negative responsiveness*, if already *xPy*, and one or more voters change their previous vote for *y* to a vote for *x*, then still *xPy*. A

> social choice mechanism is weakly monotonic if it always is the case that if some individual changes her preferences such that she is *more favorable* to *x* than previously; then if there is any change at all in the social choice, it is an outcome more favorable to *x*. Non-negative responsiveness seems a minimum requirement for an acceptable CCR (or SWF). If we were deciding on what mechanism to adopt to produce a social choice out of individual preference orderings, we would insist that whatever social choice mechanism we adopt, it will never be the case that someone's becoming more favorable to an alternative itself, with no other changes on anyone else's part, causes that alternative to be rejected.

According to *positive responsiveness*, ties can be broken by just one voter changing her mind. If *xIy* and one *y* voter changes to *x*, then *xPy*. That is, the social choice mechanism positively responds to the decision of a single voter in the case of a tie. The rule used by most juries is not positively responsive. Say that a jury is deadlocked six for conviction, six for acquittal. This might be called a tie: 6(*xPy*) and 6(*yPx*). Now say that one person changes his mind such that five jurors are for conviction and seven for acquittal. In a jury this doesn't usually break the tie. If a system fails to be positively responsive, it says that there is nothing to choose between two alternatives even though some individual changes his preferences to support one while everyone else's preferences remain the same. In effect, such a system ignores this person's change of preference.

May's Argument

These, then, are the conditions that May believes any reasonable CCR should meet. May's theorem shows that one—and only one—CCR satisfies all four conditions when deciding between two options: *simple majority voting*.

The reasoning runs like this. The first three conditions imply that if the number of votes for *y* equals the number of votes for *x*, then *xIy*. To see why this is so, suppose that it was false, and so even though *x* and *y* had the same number of votes, *xPy*. We now appeal to Neutrality, and relabel *x* as *w*, and *y* as *z*. We then appeal to Anonymity and change the identity of the voters, so that everyone who voted for *x* becomes a *z* voter, and everyone who was a *y*

voter becomes a *w* voter. Anonymity requires that the identity of the voters doesn't matter, just the number of votes, so it must be the case that if previously *xPy*, now *zPw*. But we have defined *z* as a relabel of *y*; so on the same set of preference profiles, the rule has selected both *x* and *y*, and that violates Decisiveness. Therefore, it is not consistent with our first three axioms to suppose that if *x* and *y* have the same number of voters, *xPy*. We can see that this reasoning will lead us to conclude that if the number of votes for *y* equals the number of votes for *x*, then *xIy*.

The next step is to invoke positive responsiveness. If one voter changes to *x*, then positive responsiveness requires that *xPy*. At this point we have arrived at the simple majority rule. 50% + 1 generates a winner. We also can show that non-negative responsiveness is satisfied by simple majority rule. Suppose that *x* beat *y* by *k* votes; then suppose that one *y* voter changes to *x*, so that *x* now wins by *k* + 1 votes. This change cannot reverse the outcome of *xPy*.

May's Theorem

Given two options *x, y* majority rule is the only CCR that meets the following four conditions:

1. Decisiveness
2. Anonymity
3. Neutrality
4. Strong Monotonicity

May's theorem, though limited to choices between two options, is still important. At some point we are apt to wonder. "What is so special about making decisions by a vote where over 50% wins? Why over 50% and not over 45% or 60%?" May provides an answer: out of every conceivable CCR, majority voting between two alternatives alone satisfies these four reasonable conditions. Tom Christiano, for instance, argues that majority rule is "intrinsically just," albeit only under "highly simplified circumstances."[10]

It also shows that majority rule is the unique way to settle questions between two options *only* if these four conditions hold. Which is not to say that majoritarianism may not have other virtues, but the claim to be the uniquely acceptable CCR only follows if these conditions hold. This raises a series of questions about how likely these conditions are likely to hold in

most electoral settings. The first problem is that May's theorem only holds when there are two options. This, however, is not always the case. Even in most US presidential elections, there are more than two candidates. This restriction may not be as strong as was once thought, though. Recently, Robert Goodin and Christian List have offered a version of May's theorem that extends it to instances involving three or more options, where it provides a case for plurality voting in single ballot cases.[11]

The second problem is that May's theorem doesn't include any strong rationality condition like transitivity. Nor does it generate a complete social ordering. Instead, it delivers a verdict between two options given a set of individuals and their orderings. This is not, in itself, a drawback but it does show how limited CCRs are in relation to SWF. All they can do is generate a decision given certain input, they don't construct a rational social preference structure.

More seriously, anonymity and neutrality are very strong conditions. Neutrality in particular is a very strong restriction. Most bills introduced into a legislature are attempting to change the status quo in some way and sometimes the status quo has built- in advantages. Whether this will technically violate anonymity depends on the specific case, but we have already noted that any super-majority conditions (i.e., greater than 50% + 1 required to win) will violate neutrality. Anonymity is also violated in some cases where certain voters may have certain privileges, such as the ability to break ties.

Arrow's Theorem

Condorcet Voting and Its Paradox

May's theorem shows that majority rule uniquely satisfies his conditions when choosing among two options. Let's now consider choices among three or more options. Of all the ways of choosing between three or more policies, candidates, etc., the *Condorcet method* seems the most direct extension of simple democracy between two choices (x,y). Simple majority vote handles pairwise choices very nicely; it puts a question to the voters "x or y"? Things get complicated by the introduction of a third alternative, z. We often use plurality, or "first past the post," asking "which do you want, x, y, or z?" But plurality need not give us a majority winner: x can win a plurality contest

with 40% of the vote (with *y* and *z* both getting 30%). The Condorcet method stays truer to the core idea of simple majority rule by taking a vote between each pair of alternatives: now we have three questions rather than just one: *x* or *y*?, *y* or *z*?, *x* or *z*? In each case the winner will be by simple majority: if one option beats every other alternative, then we can say that the majority prefers this option to every other; it is always preferred by a majority when compared against each other alternative, one at a time.

Condorcet voting gives rise to the Condorcet Paradox, as in profile of preferences shown in table 8.1.

TABLE 8.1. Condorcet Paradox Preferences

Alf	**Betty**	**Charlie**
x	*y*	*z*
y	*z*	*x*
z	*x*	*y*

If we take a series of pairwise votes, we get: *xPy*, *yPz*, and *zPx*—an intransitive result. Thus, out of a set of individual preferences that meets all axioms of a rational ordering, we generate an intransitive social "ordering." And what if we insist on transitivity? Then note that Alf alone is decisive over the (*x*,*z*) choice, for he is the only voter who holds that $x \succcurlyeq z$. This line of reasoning will be critical to Arrow's theorem. In some ways Arrow's theorem can be understood as a generalization of the Condorcet Paradox. As we shall see, in one stage of the proof the Condorcet Paradox orderings are invoked.

Arrow's Conditions

Again, we employ the method of *direct procedural evaluation*. We list the characteristics that we think are necessary in a procedure, and then try to determine whether there are any procedures that meet these criteria, and if there are, what they look like. Alas, what Arrow will show is that once we do that, there is no procedure that meets them all!

Kenneth Arrow's argument, fully set out in his 1960 book, *Social Choice and Individual Values*, asked whether any system of social choice that didn't rely on an individual chooser—a dictator as he calls it—can produce a

rational social ordering.[12] Specifically, he sought to determine whether his set of conditions is consistent with the social choice constituting a complete, transitive ordering of the options, that is, whether it will produce a social welfare function, an SWF. So, the requirement that the outcome must be a transitive ordering is entailed by the fact that his concern is whether the conditions he specifies are consistent with yielding a SWF. Some statements of his proof list transitivity as a condition, but that is somewhat misleading. So, Arrow's conditions are:

(1) **Universal Domain (U)**—The SWF must yield a social ordering for a set of options for all logically possible individual orderings of the options. This does *not* mean that there is no limit on the option set (this mistaken view is common). Rather, it means that once we specify an option set (*x*, *y*, *z*) that individuals are ordering, then we allow every possible ordering (permutation) of that option set. We cannot, for example, say that the ordering "$x \succcurlyeq z \succcurlyeq y$" is impermissible.

(2) **Weak Pareto Principle (WP)**—For any pair of options (*x*,*y*), if for all individuals $x \succ y$, then *xPy.* If every individual prefers *x* to *y*, then the social preference is *x* over *y.* We have already encountered this principle in Sen's story of Lewd and Prude in chapter 4.

(3) **Pairwise Independence of Irrelevant Alternatives (IIA)**—The basic idea here is that the social choice between *x* and *y* cannot depend on the presence or absence of a third—irrelevant for this pairwise choice—alternative, *z*. One way to formalize this idea is to suppose that we have a SWF, and we apply it to two different sets of individual preference orderings—by the same people—over the same option set, e.g., {*w*,*x*,*y*,*z*}, if *x* and *y* are in both sets and every individual orders *x* and *y* the same in both sets, then Condition IIA requires that the social preference between *x* and *y* cannot be different in the two sets. An individual might vary all the orderings of other preferences over other options between the two profiles, and the set of options in the two profiles may otherwise differ. Condition IIA requires that whatever social preference our SWF yields between *x* and *y* in the first profile must be the same in the second. Put slightly differently, if an individual prefers *x* to *y* when *z* is available,

adding another option *w* or removing *z* should not change the fact that $x \succ y$. Pairwise choice should be invariant to expansion or contraction of the set of options.

(4) **Non-dictatorship (D)**—There is no individual *i* such that for every option, whatever *i* prefers is the social preference. That is, it cannot be the case that there is some *i*, such that for every option in the option set, if $i{:}x \succ y$, then *xPy*.

The idea of a "Dictator" is easily misunderstood. Arrow will show that if the other conditions are met, there will always be a Dictator. But this does not mean that if we first specify a complete array of preferences—a complete preference profile for a group—we can inspect it at the outset and will find there is one-person *i*'s preference profile such that their profile is equivalent to the entire social preference ordering. In other words, that there is a Dictator in Arrow's sense does not say that we begin with a certain snapshot of all the preferences at time t_1, and that snapshot will always reveal there is some person *i*, whose preferences are always at the outset the same as the social preference.[13] Dictatorship is a *power* such that there is a person *i* who can always change her preferences in such a way that she will have control over the social choice between any *x* and *y* regardless of what others prefer.

The First Stage of the Theorem: From Local Semi-Decisiveness to Global Decisiveness

One way to understand Arrow's proof is as a two-stage argument.[14] First it is shown that *if* an individual has the *power* to decide between any one ordered pair of options, i.e., that individual's preference is the social preference even if no one else agrees, that individual becomes the overall dictator. Second, it shows that there *really is* such a person; hence both steps together establish a dictator. The first stage of the argument, then, is to show how—roughly—if some individual has the sole power to make the social choice between two ordered options, this can give him the power to decide between all other pairs.

We must first introduce two more definitions:

(a) **Semi-decisiveness**—A set of individuals *S* is *semi-decisive* for *x* over *y* if *xPy* when for every individual in *S* $x \succ y$, and $y \succ x$ for every individual not in *S*. Note two things. First, that the set *S* is

semi-decisive for *x* over *y* does not entail that *S* is semi-decisive for *y* over *x*. Second, we can see that semi-decisiveness is rather odd since *S* determines the social preferences only when everyone in it prefers *x* to *y and* everyone not in *S* has the opposite preference, *y* over *x*.

(b) **Decisiveness**—A set of individuals *S* is decisive for *x* over *y* if *xPy* when for every individual in *S*, $x \succ y$. So, if *S* is decisive, when it agrees that $x \succ y$, then *xPy* regardless of what those outside of *S* prefer. Decisiveness is a stronger condition than semi-decisiveness, so if *S* is decisive for *x* over *y*, *S* is *semi-decisive* for *x* over *y*.

We first can show that if there is some individual Alf, who is semi-decisive between any pair (*x*,*y*), then Alf must be a dictator. Starting with Alf's semi-decisiveness over just one pair, he becomes decisive over every pair and is, hence, a dictator.

So, our supposition is that Alf is semi-decisive for *x* over *y*. Let *z* be any other alternative. Suppose that Alf's preference ordering is ($x \succ y$) and ($y \succ z$), and for everyone else in society—call these "The Others"—($y \succ x$) and ($y \succ z$). Remember, given the unlimited domain condition, we can specify any preferences we like. Note that the preference relation between *x* and *z* is only specified (via transitivity) for Alf (for whom $x \succ z$). For Alf, then, we have $x \succ y$, and for The Others, $y \succ x$; Alf's semi-decisive power over the (*x*,*y*) pair implies that *xPy*.

The Weak Pareto principle now can be invoked to determine the social preference between *y* and *z*; since for Alf $y \succ z$, and for The Others $y \succ z$, then by the Weak Pareto principle *yPz*. We also know that a SWF must yield a transitive social preference, thus given (*xPy*) and (*yPz*), it must follow that *xPz*. Again, we know that *xPz* even though we have not supposed anything about any non-decisive individual's preference between *x* and *z*. We have only assumed that for The Others, $y \succ z$ and $y \succ x$. Here the IIA condition comes in: by pairwise independence of irrelevant alternatives, we know that the preferences of The Others over *y* and *z*, and over *y* and *x*, cannot affect their preferences over *x* and *z*.

So, nothing in the assumptions we have made about the preferences of The Others has any implications for their preferences of *x* and *z*; it is only Alf for whom we have specified $x \succ z$. Hence for Alf $x \succ z$ implies that *xPz*

regardless of how The Others order *x* and *z*. Alf's semi-decisive power for *x* over *y* has led to his decisive power for *x* over *z*. So now he is decisive over (*x*,*z*) too. His original semi-decisiveness has spread to decisiveness over a new pair. So, we have our fist conclusion:

(1) If Alf is semi-decisive over (*x*,*y*), he is decisive over (*x*,*z*).

And we can see that his decisiveness will spread even further. Suppose Alf's preferences are ($z \succ x$) and ($x \succ y$); that of The Others is ($z \succ x$) and ($y \succ x$). The Weak Pareto condition implies that, since for everyone $z \succ x$, it must be *zPx*. And since Alf is semi-decisive over (*x*,*y*), his preference $x \succ y$ implies *xPy*. Since we have (*zPx*) and (*xPy*), transitivity entails *zPy*. Again, note that we have no information about what The Others think about the preference relation between *z* and *y;* all we know is that Alf holds $z \succ y$ (by transitivity); so now Alf's semi-decisiveness over (*x*,*y*) has given him decisiveness over *(z,y)*. So:

(2) If Alf is semi-decisive over (*x*,*y*), he is decisive over (*z*,*y*).

Now let's use our result from step (1). Recall that we discovered: (1) if Alf is decisive over (*x*,*y*), he is decisive over (*x*,*z*). Given definitions (a) and (b), if Alf is decisive over (*x*,*z*) he is also semi-decisive over that ordered pair. Let's now see how the semi-decisiveness over (*x*,*z*) spreads. Suppose next that Alf's preferences are ($y \succ x$) and ($x \succ z$), while for all The Others ($y \succ x$) and ($z \succ x$). Since Alf is semi-decisive over (*x*,*z*), it must be *xPz*. Again, since everyone holds ($y \succ x$), the Pareto principle implies *yPx*. Invoking transitivity, if (*yPx*) and (*xPz*), then *yPz*. We again see that we have no information about the preference of The Others between (*y*,*z*), but we do know that (by transitivity) Alf holds $y \succ z$. So, the mere fact that Alf holds $y \succ z$ implies *yPz*. Thus:

(3) If Alf is semi-decisive over (*x*,*z*), he is decisive over (*y*,*z*).

By the close of the third step, Alf's semi-decisiveness over (*x*,*y*) has spread to decisiveness over the ordered pairs (*x*,*z*), (*z*,*y*), and (*y*,*z*). The same general reasoning can be repeatedly invoked to show that Alf's decisiveness keeps on spreading until he is decisive over every ordered pair among (*x*,*y*,*z*): (*x*,*y*), (*y*,*x*), (*x*,*z*), (*z*,*x*), (*y*,*z*), and (*z*,*y*). Table 8.2 summarizes the necessary steps.

TABLE 8.2. The First Step of the Proof

If Alf Is Semi-Decisive Over	Alf	The Others	Social Preferences	Alf Is Then Decisive Over
1. (x,y); assumption	($x \succ y$) & ($y \succ z$)	($y \succ x$) & ($y \succ z$)	(xPy) & (yPz), so (xPz)	(x,z)
2. (x,y); assumption	($z \succ x$) & ($x \succ y$)	($z \succ x$) & ($y \succ x$)	(zPx) & (xPy), so (zPy)	(z,y)
3. (x,z); from (1)*	($y \succ x$) & ($x \succ z$)	($y \succ x$) & ($z \succ x$)	(yPx) & (xPz), so (yPz)	(y,z)
4. (y,z); from (3)*	($y \succ z$) & ($z \succ x$)	($z \succ y$) & ($z \succ x$)	(yPz) & (zPx), so (yPx)	(y,x)
5. (y,x); from (4)*	($z \succ y$) & ($y \succ x$)	($z \succ y$) & ($x \succ y$)	(zPy) & (yPx), so (zPx)	(z,x)
6. (x,z); from (1)*	($x \succ z$) & ($z \succ y$)	($z \succ x$) & ($z \succ y$)	(xPz) & (zPy), so (xPy)	(x,y)

Once Alf is decisive over all combinations of (x,y,z), we can see that he will be decisive over *all alternatives from the full set of options* (that is, he is a dictator). To see this, take two options out of the larger set at random and call them (v,w). If (v,w) happened to be our original (x,y) alternatives, then we have already shown that Alf is decisive over them. Suppose that one of them was y, but the other was different, so we have a triple {x,y,w}. But we have already shown that if Alf is semi-decisive for x over y, he is decisive for all pairs in the triple containing x and y, so again Alf is decisive overall.

Last, suppose that v and w are different from x and y. Take just v; we have the triple {x,y,v}. We have already shown that Alf is decisive over every pair in a triple containing x and y, so Alf will be decisive for x over v. Now add to the {x,v} set w, giving us the triple {x,v,w}. We have already shown that if Alf is semi-decisive over two options in a triple, he is decisive over all the options; since we have said that he is decisive for x and v—and so he must be semi-decisive between them too—he must be decisive for all ordered pairs of {x,v,w}. But now we have shown that Alf's semi-decisiveness over two options leads to an "epidemic" since if he is semi-decisive over one pair, he is decisive over all pairs.[15] That makes him a dictator.

The Second Stage of the Theorem: Finding the Dictator

The first stage was hypothetical. It holds that *if* there is a person like Alf who has semi-decisive power over one pair (*x*,*y*), he will be a dictator. We now have to show that there is at least one pair over which Alf is semi-decisive; if we do that, given the first part of the proof, we will show he is a dictator. We know that there has to be some decisive set for every ordered pair, even if that *S* is the entire society. It follows by Definition (b), remember, that any set which is decisive is also semi-decisive. So let us examine all the decisive sets over all pairs, and choose the smallest such set (the smallest decisive set over *any* pairwise choice), and call it *V*. Assume that *V* is decisive for (*x*,*y*). We have said that *V* is the smallest decisive set; either it contains one person or more than one. If it contains one, we have found our dictator.

Suppose, then, that it contains more than one person. Divide *V* into two groups, Alf and "*V* minus Alf" (*V-Alf*). Call those outside of *V*, The Others, *O*. Unrestricted domain allows us to suppose any possible preference ordering, so we choose our Condorcet Paradox orderings, as they appear in table 8.3.

TABLE 8.3. Condorcet Paradox Ordering in Arrow's Theorem

Alf	***V-Alf***	***O***
x	*z*	*y*
y	*x*	*z*
z	*y*	*x*

Since for all members of *V* $x \succ y$, it must be that *xPy*: this follows from the definition of *V* being the decisive set over (*x*,*y*). What about *y* and *z*? We see from table 8.3 that only *V-Alf* holds $z \succ y$; if *zPy*, then *V-Alf* would be the decisive set for *z* over *y*. But we have supposed that *V* is the smallest decisive set; since *V-Alf* is a proper subset of *V*, it cannot be a decisive set. So, we know ¬(*zPy*). By completeness, then, it must be that *y* is socially at least as preferred (strictly preferred, or indifferent) to *z*, so *yRz*. If, though, *xPy* and *yRz*, transitivity implies *xPz*. If *x* is preferred to *y*, and *y* is at least as good as *z*, then *x* must be preferred to *z*.

However, Alf is the only person who holds $x \succ z$, so Alf is semi-decisive for *x* over *z*. Therefore, we have located a single person who is semi-decisive over one pair. But, by the first part of the proof, that shows him to be

decisive over all options, so he is a dictator. But we have assumed Non-dictatorship: hence we have contradicted our original assumptions. All the conditions cannot be met.

Arrow's Theorem

Given more than two options and more than two individuals, there is no SWF that can generate a social ordering that meets the following conditions:

1. Unlimited Domain
2. Weak Pareto
3. Independence of Irrelevant Alternatives
4. Non-Dictatorship

Given conditions 1–3, the social outcome is either an ordering (transitive) or non-dictatorial, but not both.

Representation and Coherence

Arrow's proof is sometimes said to reveal a conflict between "representation" and "coherence." We can obtain coherent social preferences if we are willing to abandon the requirement that the social preference reflects the preferences of the body of citizens—we can accept a dictator or abandon the Weak Pareto principle—or we can ensure that the social preference reflects the preferences of the citizens, but then, as the Condorcet Paradox shows, we can end up with an incoherent social preference. One can easily see why Arrow's theorem is seen as a challenge to the rationality of democracy. If the aim of democracy is to generate a social decision that (1) represents the preferences of the citizens no matter what their preferences, and yet (2) is coherent, it seems that democracy aims at the impossible. Not all the conditions can be met; the set is contradictory.

The Importance of Arrow's Theorem

Does Arrow's Theorem Challenge Democracy?

Arrow's theorem shows that there is no way to construct a Social Welfare Function that is guaranteed to meet his conditions. Democracy can be seen as a way of aggregative preferences (notions of betterness) into a social decision. So then does Arrow's theorem undermine the rationality of democracy? Interestingly, some insist that it must cause us to question whether democracy can be said to be a way to generate a reasonable social

choice, while others dismiss the theorem as interesting, but not crucial. There are four important ways to challenge it.

First, we might dispute whether Arrow's conditions are really intuitively compelling; to the extent that we do not mind dropping one of the conditions, the proof should not cause concern. The pairwise independence and unrestricted domain conditions both have been subject to considerable debate.

Second, it is sometimes argued that Arrow's theorem is concerned with mere "preferences," but democratic decision-making pertains to rational judgments about what is in the common good; so, it is said, Arrow's problem of how to aggregate individual preferences into a social preference is irrelevant to democratic decision- making. This challenge is, we think, misguided, for at least two reasons.

(a) As we have stressed throughout, a "preference" is simply a ranking of one option over another—it does not necessarily involve a liking, any sort of selfishness, etc. If democratic politics is about asking people to choose among candidates or policies (for whatever reasons), the idea of a preference is entirely appropriate.

(b) Christian List and Philip Pettit have presented an Arrow-like impossibility theorem that concerns the aggregation of judgments based on several reasons. So even if we understand democratic politics as a "forum" in which people debate (and vote on) the reasons for their political judgments rather than simply reporting their conclusions (rankings), a result very much like Arrow's obtains.[16]

Third, some insist that in democratic politics we do not aim at a social ordering, but only a social choice, that is, we should see democracy as a CCR (Collective Choice Rule), not a SWF, and thus, again, we need not worry about the theorem. We consider that reply in the next section.

Fourth, it can be argued that Arrow's theorem relies on the Condorcet Paradox preference orderings in finding the dictator, but if this pattern of orderings does not, as a matter of fact, often arise, then we won't actually confront the inconsistency at the heart of the proof. Indeed, it is important to realize that unrestricted domain does a lot of work in the theorem.

So, are Condorcet Paradox orderings a problem in politics? There is a nice answer to when we can expect Condorcet Paradox sorts of orderings to arise. Let us call a *dimension* an option space along which each person has some ideal point, and in each direction, the further you get from the person's ideal point, the less preferred is the option. Right to Left in politics; Dove to Hawk on war issues; pro to con on civil rights—all these are examples of an option dimension.

Consider a voter, Alf. Say he sees the options (*x*, *y*, *z*) as arranged along the Left/Right dimension in politics. Alf is a middle-of-the-road independent, so he prefers *y* as the best; *x* and *z* are worse than *y*. But let us say he thinks that *z* is better than *x*. Betty too sees the options in terms of a Left/Right dimension, but being a Leftist, she ranks *x* as the best, Alf's middle-of-the-road *y* as the second best, and *z* as the worst. Charlie also sees the options along this dimension, but as a conservative he orders them $z \succ y \succ x$. We thus obtain figure 8.1.

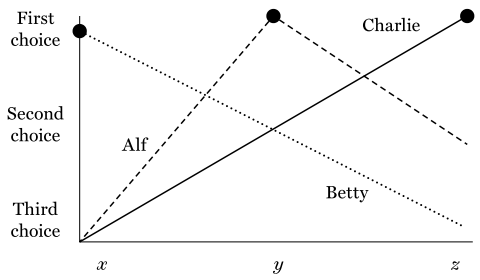

FIGURE 8.1. Single-Peaked Preferences

In this case, although Alf, Betty, and Charlie disagree on the best option, because they order their options along the same dimension, all their preference curves are *single-peaked*. If all the preferences are single-peaked, Condorcet Paradox preference orderings cannot arise.

Assume, though, that Betty no longer sees the options in terms of Left/Right, but in terms of for or against free markets, and she rejects free markets. And suppose that she sees *x* as the best option (the most left and the most restrictive in terms of markets) but her second choice is *z*, the proposal of the right, since she thinks that conservatives at least appreciate the limits of the market—the conservatives too have doubts about "liberal economics" characteristic of *y*. She thus changes her ordering, and we get figure 8.2.

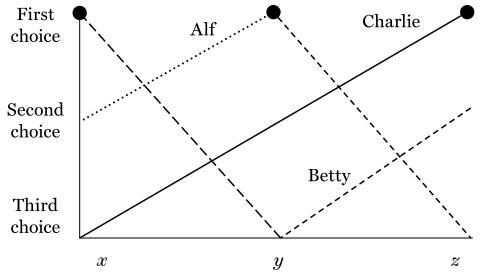

FIGURE 8.2. Double-Peaked Preferences

Betty no longer orders her preferences along the *x*-*y*-*z* dimension; we can see that her preference curve is *double-peaked*: this is not the relevant dimension along which she orders her options. Even though *y* is closer to her ideal point than is *z*, she prefers *z* to *y*. And now we have a Condorcet Paradox ordering. Thus, when voters do not agree on the relevant dimension of the option space, Condorcet Paradox orderings can arise.

How often do cycles actually arise? There is spirited debate about this in the literature on democracy: some think that uncontrived cycles are rare, while others believe they are more common. As we will see, however, the main importance of Condorcet Paradox type cycles may be the way they lend themselves to various sorts of contrived manipulation.

Collective Choice Rules

It is hard to have faith in any SWF in the light of Arrow's theorem. To be sure, as we have stressed, SWFs can, under some profile of preferences, satisfy the Pareto principle, non-dictatorship, and the independence of irrelevant alternatives. But when the preference profiles of the citizens display considerable multidimensionality, Arrow-type problems come to the fore. Moreover, as we shall see, minorities can generate voting cycles by misrepresenting their true preferences, thus magnifying the problems to which Arrow points.

At this juncture, though, students of democracy may insist that social welfare functions are not really of much interest. After all, we do not want a social ordering, just a social *choice*: for any set of options, we only want to pick out the *best*. Our interest should be in collective choice rules, not social welfare functions. At first blush this looks inviting since some CCRs meet all of Arrow's conditions, so the proof does not preclude all Collective Choice Rules. The difference between a SWF and a CCR is, of course, the

transitivity requirement—Arrow requires strict transitivity of the social ordering of all options.

What does a CCR require? Let us follow Sen in requiring of any collective choice rule that, for any set of options {*w,x,y,z*}, the CCR can select one option as best. Now Sen shows that, so understood, such a CCR requires social preferences that are reflexive, complete, and *acyclical* over the entire set of preferences.[17] Social preferences are acyclical if {*wPx*, *xPy*, *yPz*} then *wRz*, that is, *w* must be, socially, at least as good as *z*. We can easily see how cyclical preferences mean that there may be no best choice from some sets. Consider the Condorcet Paradox preferences and consider what is the best choice from the set {*x,y,z*}. There is, of course, no best choice from the entire set: we get *xPy*, *yPz*, *zPx*, so we cannot choose. So, if we are going to guarantee that a choice can always be made, we must ensure that we do not have cycles; thus, we require that the results be *acyclical*. To get a better idea of a Collective Choice Rule, consider an example of a rule that Sen discusses:

> *Pareto Choice or Indifference*—if for everyone $x \succ y$, then *xPy*; if it is not the case that for everyone $y \succ x$, then *xRy* (that is, *x* is socially at least as good as *y*).

So, on this rule, *x* is socially preferred to *y* only if everyone prefers it; and, of course, the same goes for *y* vis-à-vis *x*. If there is any disagreement in people's strict preference profiles about *x* and *y* (for some $x \succ y$, for others $y \succ x$), then (*xRy*) and (*yRx*)—*x* is socially at least as good as *y*, and *y* is socially at least as good as *x*. And that means that socially *x* and *y* are indifferent (*xIy*).

Using this Collective Choice Rule, our Condorcet Paradox orderings yield a "choice" of *xIy, yIz, xIz.* But in this case the rule satisfies all of Arrow's requirements. Suppose instead we just consider Alf's and Betty's preferences. For Alf ($x \succ y \succ z$) and for Betty ($y \succ z \succ x$). Since for both $y \succ z$, the CCR yields *yPz*; since they disagree on the other pairings, we get social indifference—*xIy*, *yIz*. Now, however, transitivity is violated. Given (*xIy*) and (*yPz*), we should get (*xPz*): if society is indifferent between *x* and *y*, but *y* is better than *z*, transitivity would hold that *x* is better than *z*. Instead, our CCR says *xIz*.

However, acyclicity does still hold: we get a “best” (even if it is a tie) from every possible set. We also can see that Pareto choice or indifference satisfies all the rest of Arrow’s conditions. It obviously satisfies the Pareto condition, since it is defined never to yield a choice against it. It does not violate the independence of irrelevant alternatives, since we can see that in deciding between *x* and *y*, it is only concerned with people’s preferences over *x* and *y*; we have a good inkling that unrestricted domain will not cause trouble, since we have used Condorcet Paradox orderings; and we can never have a dictator, since *xPy* only if everyone agrees.

Problems of Collective Choice Rules

However, while Arrow’s theorem does not apply to CCRs (since they do not require transitivity), they have their own Arrow-like problems. Arrow, it will be recalled, showed that a SWF meeting the pairwise independence condition (IIA), unrestricted domain (U), and the Pareto principle (P) always has a dictator; it has been shown that any CCR that meets *IIA*, *U*, and *P* will always have something very much like an *oligarchy* (a subgroup) that can always veto a social choice.[18] That is, the dynamic of Arrow’s theorem, in which decision-making power becomes concentrated in one person, also obtains for CCR, though in a somewhat mitigated form. Instead of a dictator, we find subgroups that are able to impose their will by blocking whatever others want or (given slightly different conditions) actually imposing their will.

Transitivity can be seen as a condition of consistent choice in Social Welfare Functions. When we consider Collective Choice Rules, we still need some conditions of consistent choice. Two properties seem attractive: the *contraction* and *expansion* properties. According to the *contraction* property, if *x* is chosen from the entire set S, it must be chosen from all subsets of S in which *x* is included. Our polite mango refuser from chapter 2 violates this by selecting a mango from the set {mango, mango, apple} but an apple from the subset {mango, apple}. Our chooser will also violate the weak expansion principle: if an option is chosen from each of two subsets, it must still be chosen when the sets are combined. Suppose our person is confronted with two sets {apple, apple, mango} and {apple, mango}. Because she will not take the last mango, she will choose {apple} from the first set and {apple} from the second. But if we combine the two sets to get {apple, apple, apple,

mango, mango} she will choose a mango, thus violating the weak expansion property. So according to the *contraction principle*, if from a set {*x*,*y*,*z*} a person chooses *x*, he must choose *x* from the other subsets in which it enters {*x*,*y*}, {*x*,*z*}. If it is the best in the entire set, it must be best in every subset it is in. It would be odd indeed if, thinking that *x* is the best choice when confronted by three options, one no longer thinks that *x* is best if one considers it in relation to just one of those options.

According to the *expansion principle* (roughly), if *x* is chosen as best in the set {*x*,*y*,*z*}, and if there is a set {*w*,*x*,*y*,*z*}, *y* or *z* cannot be the social choice from the larger set; if *x* is better than *y* and *z* in the small set, it must be better than they are in the larger set. This expansion principle implies the *weak expansion principle:* if *x* is chosen from the set {*x*,*y*,*z*} and *x* is chosen from the set {*v*,*w*,*x*}, it must be chosen from the union of those two sets.

These seem to be reasonable conditions of consistent choice since collective choice rules that violate them yield inconsistent choices about what is the best depending on how we divide the options. If we violate these principles, then an option that is best in a large set may not be best in a smaller set; an option that beats another in a small set may lose to that same option in a larger set. However, it can be shown that the expansion property applied to CCRs essentially turns them into SWFs, and so Arrow's theorem then applies to such CCRs. Thus, if we wish to avoid Arrow's theorem by adopting a CCR, we must abandon the expansion principle as a principle of rational choice. And that involves a considerable cost. But if we do not abandon it, Arrow's theorem applies to our collective choice rule.

Extending Social Choice Rules

We know from May's theorem that simple majority decision-making over two alternatives is uniquely characterized by universal domain, anonymity, neutrality, and strong monotonicity. The problem—and this is at the heart of Arrow's theorem—is that we seldom have only two alternatives since usually there are a number of possibilities among which we are trying to decide. Is there any way to extend the wonders of majority decision-making to these cases with three or more alternatives? Although there are indeed methods—some of which are widely used—that seek to extend majority decision-making to three or more alternatives, each of them has flaws.

Remember that a social choice mechanism is some way of taking all the preference orderings of individuals and generating out of these a Collective Choice. If a full social ordering is produced, we have a Social Welfare Function; if only a best choice is made, we have a Collective Choice Rule.

Condorcet Method

We have considered the Condorcet method in our explication of the Condorcet Paradox. It seems the most direct extension to majority vote between two alternatives to a vote among more than two alternatives: it splits the choices into pairs, and then makes pairwise majority decisions. We have seen from the Condorcet Paradox that the Condorcet method violates transitivity. This also means that it violates May's Decisiveness condition, since it cannot identify a best from the set.

A Weakly Positional Method: Plurality

Let's distinguish two sorts of positional methods, Weakly Positional and Strongly Positional methods.

A *weakly positional* method uses a categorical (yes/no; best choice) ballot, and then determines the outcome based upon tallying the categorical ballots. A categorical ballot doesn't ask you to order the candidates. Instead, it simply asks you which one do you like the best? You ignore all the other options and vote for the one that you think is the best.

Plurality is employed in most elections in the UK and the United States. These systems are based on a categorical ballot. When a Republican, a Democrat, and an Independent are running against each other, voters are not asked to rank them in order of preference. All they are asked to do is report their first preference. The plurality method then allows society to order the candidates by the number of votes they get: the candidate that is at the top of the list wins. Plurality is probably best conceived of as a Collective Choice Rule, since it seldom cares about anything but the candidate who gets the most votes. Yet it can also function as a SWF, when it makes a difference which party came in second (perhaps they will form the official opposition).

We call this a *weakly* positional method because (1) the voters do not order the candidates but (2) all the candidates run against each other at the

same time, and (3) the candidate who comes out on top—has the top position—wins.

Plurality is a simple method. Unlike the Condorcet method or the Borda method (see below), we do not have to do a great deal of work in order to find the plurality winner. The problem is that plurality is able to select what is, in a fairly clear sense, the most unpopular option, as shown in table 8.4.[19]

TABLE 8.4. Plurality Voting Example

Voters	1st	2nd	3rd	4th
9	*w*	*z*	*x*	*y*
6	*x*	*y*	*z*	*w*
2	*y*	*x*	*z*	*w*
4	*y*	*z*	*x*	*w*
5	*z*	*x*	*y*	*w*

Note that *w* gets the most first-place votes (9) but everyone else (17) places it last. William Riker argued that this has occurred in actual American elections, such as the US 1912 presidential election, and perhaps the critical 1860 election that elected Abraham Lincoln.

Viewed as a Collective Choice Rule, plurality violates the expansion condition. It can beat an option in a small set but loses to that same option in a larger set. Suppose that we have a five-member electorate, choosing between only *x* and *y*, and the results are as presented in table 8.5, where *x* gets three first-place votes and *y* gets two, so *xPy*. Now *z* enters, and two voters (Betty and Charlie) prefer it to *x*, giving us the results in table 8.6.

TABLE 8.5. Plurality Voting with Two Options

Alf	Betty	Charlie	Doris	Eugene
x	*x*	*x*	*y*	*y*
y	*y*	*y*	*x*	*x*

TABLE 8.6. Expanded Plurality Voting with Three Options

Alf	Betty	Charlie	Doris	Eugene
x	*z*	*z*	*y*	*y*
y	*x*	*x*	*x*	*x*

Alf	Betty	Charlie	Doris	Eugene
z	*y*	*y*	*z*	*z*

Notice that no one changes their relative orderings of *x* and *y* in the two profiles; all that changes is that they now must insert *z* into their orderings. But we see with the addition of *z*, which is not relevant to the social preference between *x* and *y*, the social ordering goes from *xPy* to *yPx* (we have a tie for first place, *zIy*). This violates the expansion consistency condition: you can defeat a victor, *x*, without changing *anyone's* preference, simply by adding another option. Thus, "splitting the opposition" by introducing a third alternative is a typical strategy in systems that use plurality.

A Strongly Positional Method: Borda Count

A strongly positional method presents the voters with an ordinal (rather than categorical) ballot, and then compares all the candidates at the same time to get a social ordering. A strongly positional method is the Borda count. The Borda count tallies the ranks that each voter assigns to each candidate. The candidate with the lowest score wins. Table 8.7 calculates the Borda winner (and ordering) from the plurality profile from above.

TABLE 8.7. Borda Count Example

	1st	2nd	3rd	4th	Total
w	9 (× 1)	0	0	17 (× 4)	77
x	6 (×1)	7 (× 2)	13 (× 3)	0	59
y	6 (× 1)	6 (× 2)	5 (× 3)	1 (× 4)	37
z	5 (× 1)	13 (× 2)	8 (× 3)	0	55

So, we get *yPzPxPw*—a different result from the plurality rule. Borda has many attractive qualities. It has, for example, a high tendency to track the Condorcet winner when there is one: when Condorcet yields cycles, Borda will usually yield a tie.

Yet, it has drawbacks. As a positional method it violates Independence (IIA). We can actually identify two Independence conditions. Recall the characterization that I employed in Arrow's theorem: suppose we have an SWF, and we apply it to two different sets of individual preferences orderings over the same option set. Suppose that *x* and *y* are in both sets, and every individual orders *x* and *y* the same in both sets. That is, whatever individual *i*'s preference is between *x* and *y* in the first profile, he has the same pairwise preference in the second profile.

To see how Borda violates this condition, compare tables 8.8 and 8.9.

TABLE 8.8. Borda Count Basic Case

Voters	**1st**	**2nd**	**3rd**	**4th**
9	*w*	*z*	*x*	*y*
6	*x*	*y*	*z*	*w*
2	*y*	*x*	*z*	*w*
4	*y*	*z*	*x*	*w*
5	*z*	*x*	*y*	*w*

TABLE 8.9. Borda Count with Preferences that Violate IIA

Voters	**1st**	**2nd**	**3rd**	**4th**
9	*w*	*z*	*x*	*y*
6	*x*	*y*	*w*	*z*
2	*y*	*x*	*w*	*z*
4	*y*	*z*	*x*	*w*
5	*z*	*x*	*y*	*w*

Two groups of voters have changed their relative evaluation of *z* and *w*—they have decided that *w* is better than *z*—but no one has changed their relative evaluations of *x* and *z*. Yet, because *z* now gets additional points by those voters who have relegated it to their fourth position, society goes from *zPx* to *xPz*, a violation of pairwise independence.

The Independence condition can also be understood as an *expansion or contraction consistency* condition. On this interpretation, we take option set {*x*,*y*,*z*} and we either expand it to include *w*, or contract it to omit *z*. According to this interpretation, if everyone's ordering between *x* and *y* is

the same in both profiles, the social ordering of *x* and *y* should be the same in both profiles. Consider table 8.10, which is simply a contraction of our initial profile with *y* omitted, and no one has changed *any* orderings of the remaining *w*, *x*, or *z*.

TABLE 8.10. Borda Count Contraction Inconsistency

Voters	1st	2nd	3rd
9	*w*	*z*	*x*
6	*x*	*z*	*w*
2	*x*	*z*	*w*
4	*z*	*x*	*w*
5	*z*	*x*	*w*

Again, society goes from *zPx* to *xPz*,[20] violation of this understanding of pairwise independence.

Cardinal Methods

A positional method such as Borda counts "aheadness"—it counts how many places ahead one option is to another. To say that *w* ended up three places behind *y* in the social ordering is not to say that anyone, or society, preferred *y* three times as much as *w*, or even that it far more intensely prefers *y* to *w*. However, a number of democratic theorists—especially, we think, economists—have felt that we should measure how *much* you prefer an option. We thus may seek to measure the utility an option gives a voter—say, using von Neumann–Morgenstern utilities—and then aggregate them.[21] But wait! We also need full interpersonal comparisons, so that we can say that if for Alf $\mu(y) = 1$, that is the same utility as Betty = $\mu(y) = 1$. We have already looked at a bit of the debate surrounding interpersonal comparisons of utility in chapter 2. To use the cardinal method of voting, there must be a good argument for comparing people's utilities in a certain way.

A simpler method—one that does not seek to account for attitudes toward risk—is just to ask each person to place each option on a cardinal scale, and again we assume that the top and bottom of these scales are in some sense "equivalent" for each voter. At this point we could either simply add the utilities (this is sometimes called the "Bentham method") or, alternatively,

we perform some other aggregation method such as multiplying the utilities (this is sometimes called the “Nash method”). Cardinal methods violate Independence for much the same reason as Borda does.

Elimination Methods

Given that majority vote over two options has the wonderful qualities May demonstrated, perhaps what we should be doing is trying to get down to a binary choice. Elimination procedures seek to do this through a series of votes (or stages) in which a large set of options is reduced to two. Runoff elections and primary systems are examples of elimination procedures.

Some elimination procedures are built into the way the voters are counted. Consider for example, a system along the lines of that used in elections to the Senate in Australia—the single transferable vote (STV). Elections to Australia’s Senate (the upper chamber of the legislature) combine the single transferable vote with proportional representation. Voters have to order all the candidates running from first to last (or at least some minimum number).[22]

TABLE 8.11. Single Transferable Vote (STV) Example

Voters	**1st**	**2nd**	**3rd**	**4th**
9	*w*	*z*	*x*	*y*
6	*x*	*y*	*z*	*w*
2	*y*	*x*	*z*	*w*
4	*y*	*z*	*x*	*w*
5	*z*	*x*	*y*	*w*

The single transferable vote works by rounds. If a round does not produce an option that has more than 50% of the vote, the option with the least number of votes in that round is eliminated, and its votes are *transferred to* the voter’s next preference. If no majoritarian winner yet appears, the procedure is repeated until one option gets a majority. So, at the end of the process, some option does get a majority of support.

In table 8.11, no option receives a majority in the first round, so the option with the least votes in that round is eliminated: in this case *z*, which

has five first-place votes. So, we now have the profile presented in table 8.12.

TABLE 8.12. Second Stage of the STV

Voters	1st	2nd	3rd
9	*w*	*x*	*y*
6	*x*	*y*	*w*
2	*y*	*x*	*w*
4	*y*	*x*	*w*
5	*x*	*y*	*w*

There is still no majority winner, so *y* (the remaining option with the lowest first-place votes), and its votes are transferred to the voter's next choice, as in table 8.13.

TABLE 8.13. Third Stage of the STV

Voters	1st	2nd
9	*w*	*x*
6	*x*	*w*
2	*x*	*w*
4	*x*	*w*
5	*x*	*w*

At this point *x* emerges as the majoritarian winner. Alas, elimination procedures violate May's condition of non-negative responsiveness (which is implied by Arrow's Weak Pareto Condition). To see this, let's go back to the original figure, shown in table 8.14.

TABLE 8.14. STV with Monotonicity Violation

Voters	1st	2nd	3rd	4th
9	w	z	x	y
6	x	y	z	w
2	y	x	z	w
4	y	z	x	w
5	z	x	y	w

Recall that on the single transferable vote system, x is the winner. Now non-negative responsiveness (monotonicity) requires that it is impossible for x to be hurt by being *more* preferred: someone voting for x cannot cause x to lose!

To see how the single-transferable vote violates this condition, suppose that the two voters who rank the options $y \succ x \succ z \succ w$, decide they prefer x to y, so they switch their order to $x \succ y \succ z \succ w$. So, let's go through the elimination procedure again. Note that y is now the weakest option in the first round, so it is eliminated, and its votes transferred to the voter's next choice, producing the second-round tally in table 8.15.

TABLE 8.15. STV Winner after Vote Switching

Voters	1st	2nd	3rd
9	w	z	x

Voters	1st	2nd	3rd
6	*x*	*z*	*w*
2	*x*	*z*	*w*
4	*z*	*x*	*w*
5	*z*	*x*	*w*

Now *x* is eliminated (having 8 votes to *z*'s 9), and *z* goes on to win. Thus, because two voters became more favorable to *x*, it lost. Many, including Riker, believe this is a serious flaw in a voting system. It also, we shall see, makes elimination systems subject to manipulation.

What to Choose?

So where are we in deciding how to extend majority voting to three or more options? In our core example, Plurality chooses *w*, the Condorcet method chooses *z* (check this for yourself), Borda selects *y*, and the Single Transferable vote gets us *x*. We have managed to select all four possibilities through commonly used voting systems. In this case, pick your favorite way of extending majoritarianism to three options, and you can also pick your favorite option. As we will see, this is a generalizable feature of voting systems.

Path Dependency and Agenda Manipulation

Arrow stressed the importance of a final choice being *independent of the path to it.*[23] The crux of *path independency* is that the best choice from the set {*w,x,y,z*} should not vary depending on the order in which we take up the options, or the subsets we first consider in our path to a final choice. Path *in*dependency assures us that when our CCR (Collective Choice Rule) selects a "best option" from a set, this selection is not an artifact of the order in which we considered the options—the path by which the final choice was reached. If we considered the same options in a different order, we would have arrived at the same choice from the same set.

If we aim at a CCR that really uncovers what is the best social choice given a profile of individual preferences, we will insist on a path-

independent choice rule. When we employ a path-*de*pendent choice rule, we must accept that while our CCR selected *y* as best from the set, it could have selected some other option from the same set if we had considered the options in another order. In that case our choice is an artifact of the path. However, it turns out that path independence requires the *contraction* and *weak expansion principles*; thus, any path-independent CCR is subject to Arrow-like problems of concentration of decision-making power in some oligarchic-like group.[24]

If we want to avoid Arrow-like concentration of power problems, we must embrace path-*de*pendent CCRs. Well, we might ask, what is so bad about path-dependent Collective Choice Rules? Why not avoid Arrow-like problems by simply rejecting the contraction and weak expansion principles?

Consider more closely what follows from path dependency, such as the procedure used to decide on bills in the US Congress. First a proposal (call it *x*) is brought before Congress; amendments (call one *y*) are then proposed, and a vote is taken between *x* and *y*. Finally, there is a vote between no bill at all (call it *z*) and whatever bill won the *x* versus *y* contest. This means that one of the pairs—either (*x*,*z*) or (*y*,*z*)—is never voted on. Since the outcome depends on the order in which the options are considered—or, we might say, the way the total option set is divided up in decision-making), those who control the agenda—the order in which the options are considered—have great power in determining the outcome. A committee chair who sets the agenda may well be able to determine the outcome! Path-dependent CCRs thus allow for *manipulation of the agenda.*

Path Dependency in Legislation

Again, there is lively dispute about how pervasive path dependency and agenda manipulation are. William Riker famously argued that they were the stuff of politics. Consider, for example, Riker's account of the aid to education bill in the late 1950s and early 1960s.[25] The Democratic leadership in the US House of Representatives put forward a bill that would provide the states with federal funds for building schools. The Democrats were generally in favor of such aid while the Republicans opposed it. However, each time the bill was proposed, Adam Clayton Powell, an African American member from Harlem, offered an amendment, requiring that aid would only be given

to states whose schools were “open to all children without regard to race.” This was called the Powell Amendment.

Southern Democrats—who supported federal aid to education, since they would be net recipients—opposed the amendment, since under the amendment the South would not receive aid, having segregated school systems. In contrast, urban Democrats from the North supported the amendment; but the Democratic leadership opposed the amendment, primarily because it would make the school aid bill harder to pass. The Republicans, who were generally against the expansion of the federal government, opposed federal aid to education. To analyze the voting, let:

x = the bill with the Powell amendment
y = the original, unamended bill
z = the status quo (no federal aid to education)

Let us distinguish four groups in Congress:

The Powellians ($x \succ y \succ z$)—These members preferred the amended bill to the original bill and preferred both to no bill at all. Northern Democrats would form the core of this group.

The School Aiders ($y \succ x \succ z$)—This group preferred school aid to anything else. This included Democrats following the Democratic party leadership—who saw the Powell amendment as killing school aid—and some Republicans.

The Southerners ($y \succ z \succ x$)—As the story indicates, they were segregationist Democrats who supported school aid, but strongly opposed the Powell amendment.

Republicans Against Aid ($z \succ x \succ y$)—This probably is not quite right; Riker argues that the Republicans split into two groups, as we will see later on. But for now, we will treat them as one group with this preference order.

Riker advances the following estimate of the relative strengths of the groups:

The Powellians ($x \succ y \succ z$): 132
The School Aiders ($y \succ x \succ z$): 67
The Southerners ($y \succ z \succ x$): 130

Republicans against Aid ($z \succ x \succ y$): 97

Given this estimate, the result of Condorcet voting would be *xPy* (229 to 197), *yPz* (329 to 97), *zPx* (227 to 199)—a cycle. Note that the original bill (*y*) easily defeats no bill (*z*) by a large margin of 329 to 97. But because of path dependency, *z* was actually selected. Recall that given the path-dependent process followed by Congress, once *y* was defeated by *x* it was eliminated from consideration, leaving only *x* to survive and confront *z*, with *x* losing to *z*. We see, though, that *y* would have defeated *z*, since we are in a cycle.

Gerry Mackie has investigated Riker's examples (including the case just presented) and has argued that careful investigation does not support Riker's conclusions. He makes two distinct claims.[26]

First, he challenges Riker's estimations of the voting strength of these groups. If *y* (the original bill) met *z* (no bill), Mackie argues that, based on another similar vote, *z* would have defeated *y*, hence there would have been no cycle.[27] Whereas Riker holds that *y* would have beaten *z* by a large margin, Mackie argues that *z* would have beaten *y* by a smaller one. So, there is dispute about the true strengths of the groups, and whether there was a cycle—Riker himself suggests that there may not have been a "real" cycle, but that some of the Republicans voted "strategically."

Second, Mackie argues that even if Riker's estimates of the strength of the groups are correct, Riker's conclusion that there is no coherent way to combine preferences still does not follow. "Does an alleged cycle," Mackie asks, "show once again that the aggregation of preferences is incoherent?" "No," he answers, because there would be other ways to choose among these preferences that would not result in a cycle. "If there was a cycle, then the problem is not with the preference rankings, the problem is with the voting procedure."[28] We are not sure what Mackie means by this. Consider again our Condorcet preference orderings.

TABLE 8.16. Condorcet Preference Ordering

Alf	**Betty**	**Charlie**
x	*y*	*z*
y	*z*	*x*
z	*x*	*y*

We know that pairwise voting yields a cycle, *xPy*, *yPz*, *zPx*. But, as Mackie suggests, we might employ another procedure such as the Borda count that will avoid a cycle. In this case, the result is a three-way tie. In some cases, the Borda method can yield a strict social preference even when pairwise voting yields a cycle.

This seems the basis of Mackie's claim that the problem is pairwise voting, not the preference rankings—some other system of voting would yield an answer. Yes, but recall Arrow's theorem: the problem is that every system of voting violates one of the conditions. Borda does have the nice advantage that it gives us a social preference in these cases that is path-independent: there is no ordering of the voting process. Pairwise voting violates transitivity in the case above—to avoid that result, use of Condorcet voting typically imposes an order on the way the pairs are taken up, and that leads to path dependence. But if we use Borda, as we have seen in the previous sections, we then violate pairwise independence of irrelevant alternatives (IIA).

The lesson of Arrow's theorem is that, as it were, we must choose our poison—will it be intransitivity, violation of the pairwise independence condition, the Pareto principle, unrestricted domain, or non-dictatorship? We can always avoid cycles, but then we must violate some other condition. The "incoherence" of aggregation systems is not that they must yield cycles—that is clearly not the case, and Riker never thought it was—but that they must violate some basic axiom.

Agenda Control in the Flying Club

Michael Levine, a professor of law at USC, went to Charles Plott, a professor of economics at Cal Tech, and said, "Charlie, I'm chairman of an agenda committee of my flying club.[29] We're going to decide on a new fleet. They made me chairman because I'm the only lawyer in the club. But you know more about voting than I do. What's the fair way to do this?" Plott replied that there is no true aggregation of individual preferences into a social preference; lots of reasonably fair rules will yield lots of different results. So, we might as well get the one you want. Let's arrange things that way.

The flying club had sixty-five members who, in 1973, paid $1800 each to join, plus $30 a month, and $25 per hour for the plane rental. They had a

fleet of six single-engine planes: three five-year-old Beechcraft E33As, one three-year-old Beechcraft F33A, one three-year-old Beechcraft V35, and one new Cessna 210. The members preferred roomy, fast aircraft, which tend to be expensive. They were dissatisfied with the current fleet because:

- It was not uniform (the V35 and the Cessna 210 had different operating characteristics from the rest of the fleet, a safety concern).
- The older planes were shabby and disliked. They wanted better avionics equipment.
- The planes were not always available when people wanted them.

Members had a clear preference for eliminating the one V35. But, given the club's resources, they were limited to a fleet of four Beechcraft E33A and F33A and possibly one or two six-seated airplanes, either a Beechcraft A36 or Cessna 210. Table 8.17 provides the options.

TABLE 8.17. Flight Club Options

Make and Model	Abbreviation Used	Number of Seats	Estimated Rental Per Hour
Beechcraft A36	A	6	$31.50
Beechcraft F33A	F	4	$30.00
Beechcraft E33A (refurbished)	E	4	$24.00
Cessna 210	C	6	$27.00

Most members prefer Beechcraft, while members of the governing board preferred Cessnas because of their new Cessna dealership. A few highly influential members wanted Beechcraft A36s, the most luxurious and expensive single-engine Beechcraft. On the other hand, many members seemed concerned about expenses; they expected to refurbish three E33As but didn't want to buy additional E33As.

Levine's Preferences:

(1) Include in the fleet some six-seated aircraft for family trips. He preferred the cheaper Cessna 210s to the more expensive Beechcraft.
(2) If the 210 was bought, he would like two.
(3) If the A36 was bought, he would want one.
(4) He was indifferent about buying new F33As or E33As to refurbish.

So, his overall fleet preferences were:

1. (5E, 2C) or (3E, 2F, 2C)
2. (5E,1C) or (3E, 2F, 1C)
3. (6E, 1A) or (3E, 3F, 1A)
4. (5E, 2A) or (3E, 2F, 2A)
5. and lower (all E and F)

Levine believed this differed from the rest of the club, who might have preferred all Es, and if the rest of the club did accept a secondary fleet of six-seat planes, they probably preferred As to Cs.

The problem he and Plott confronted was how to devise an agenda that would produce a fleet with one or two Cs against a possible majority for Es and maybe As. The way they would do this is to devise an *agenda* to decide the following questions sequentially:

1. What type of aircraft should be in the primary fleet?

This question was to be decided by Borda count, not because they had any doubt about the outcome, but to ease the concerns of those members who feared being stuck with an expensive fleet, i.e., one that was foisted on them by the influential members who wanted As, and make them agreeable to a seven-plane fleet, which would be more likely to include Cs than a six-plane fleet.

2. How many planes do you want?

This was the crucial question. They believed a plurality wanted a fleet of six four-seat planes. It was important to pose this question in a way to isolate this plurality. They wanted to produce a majority coalition for seven planes, consisting of those who wanted seven four-seaters and those who wanted a fleet with some six-seaters, either As or Cs—the odds of getting a six-seater were better in a seven-plane fleet.

3a. Do we want a mixed fleet?

The third question was intended to pit all those who wanted any kind of six-seat plane—either A or C—against those who wanted a fleet of all Es and Fs. Note here that Plott and Levine were seeking to use the supporters of A to reach a point on the agenda where C could defeat A.

3b. Should the secondary fleet be one or two planes?

In the course of the meeting this question was added to the agenda.

4. **What type of aircraft should be in our secondary fleet?**

   This was intended to finally pit the As against the Cs. Given their expectation that a seven-plane fleet would be adopted, those who were concerned with expenses would side with those who favor the cheaper planes, the Cs.

5. **How elaborately should we equip our aircraft?**

   Levine and Plott saved this for last to make sure that expensive outfitting wouldn't preclude a seven-plane fleet.

The results of this elaborate agenda went in line with their expectations. First, the club tied on the fleets of (3E, 2F, 2C) and (3E, 2F, 2A), but since the governing board preferred Cessnas (they owned the dealership!), this ensured Levine's s favorite outcome of (3E, 2F, 2C). Second, a questionnaire was sent to all members after the meeting asking them to rank the alternatives. If all the main alternatives had confronted each other in a round robin tournament, the winner would have been (3E, 3F, 1A). That would be the Condorcet winner. Levine's preference would have come in second.

The chairman of the club, however, sought to impose his own agenda; he tried to change the opening question to "should we have a fleet of all Beechcraft?" This probably would have been passed and would have precluded Levine's favored fleet. He also tried to introduce before Levine's second question: "If we have a secondary fleet, do we want Beechcraft or Cessena?"

Finally, after question 2 but before 3a, he tried to introduce the question, "Do we want at least one A?" An affirmative answer to any of these questions would have ruled out a fleet with C.

On all these points he was ruled out of order, preserving Levin's agenda.

The result is that Plott and Levine were effectively able to manipulate the Flying Club into selecting their preferred alternatives and to outmaneuver the attempts by the chairman to do something similar. The lesson, according to Riker, is that this type of manipulation is the essence of politics, what he calls "Heresthetics." This is the art of structuring the choice situation in order to win. Agenda manipulation is at the heart of this skill, which Riker argues that all genuinely good political actors must possess.

Riker's lesson goes deeper, though, since the possibility of Heresthetical techniques arises because of the underlying instability of results that Arrow's theorem and similar results show. Agenda manipulation is the result of disequilibrium, such as we find in a Condorcet cycle. In that there is no equilibrium, we can expect a wide variety of results from the same set of preferences; we can go down a wide array of paths, each one yielding a different outcome. And disequilibrium occurs when voters employ different dimensions of evaluation.

Strategic Voting

The problems we have thus far been exploring all have assumed that each individual is reporting their "true" preferences. Of course, on revealed preference theory it does not make any sense to distinguish what a person truly prefers from their behavior; but even so, there is a sense in which a person might order options not because they want the outcome as such, but in order to induce a further outcome. So, we suppose that people can indeed misreport their preferences; by doing so they often can obtain their favored outcome. The most obvious case of strategic voting involves multi-stage votes—such as votes on amendments and bills in Congress—and runoff elections.

For example, in the Aid to Education bill discussed previously, Riker notes that the Democratic leadership in the House was convinced that the Republicans were not really in favor of the Powell amendment, but were voting strategically just to help kill the aid to education bill.[30] That is, the Democratic leadership was convinced that the Republicans' true preference ordering was: $z \succ y \succ x$ (no bill, original bill, amended bill). If that was the true preference ordering of even half the Republicans, the amendment would have failed, and Congress would have voted on the unamended bill, which would have passed. By misreporting their true preferences—voting for their third choice over the second choice when the amendment was initially proposed—these Republicans helped to secure their first choice.

Now, what amounts to "true" preference here is a little misleading. The Republicans, for instance, may have preferred $y \succ x$ if they alone were able to select from that pair, but of course this was not the situation in which they found themselves. The objection or even identification of strategic voting

seems to presume that voting is more akin to a parametric decision problem, where the voter acts to bring about an outcome directly, rather than a strategic problem, where the voter acts—in conjunction with other voters—to bring about an outcome. To ignore the strategic nature of any voting situation is to misrepresent the choice problem, though. Unless we assume that voters are primarily concerned with the act of voting rather than with the outcome voting is meant to bring about, we shouldn't be surprised that they will respond to the strategic situation they find themselves in.

Other Examples of Strategic Voting

The Borda count discussed above is clearly subject to strategic voting. Consider table 8.18.

TABLE 8.18. Example of Strategic Voting

Alf	**Betty**	**Charlie**	**Doris**	**Eugene**
x	*x*	*x*	*y*	*y*
y	*y*	*y*	*x*	*x*
z	*z*	*z*	*z*	*z*

If we use the Borda count, we get $x = 7$ points; $y = 8$ points, $z = 15$ points, so *xPy.* Suppose that Doris, who wants *y* to win and knowing that *z* will win otherwise, strategically reports that *x* is her third, not her second choice. The outcome would then be 8 points for both *x* and *y*, producing a tie (*xIy*).

Elimination procedures and runoffs are also clearly subject to strategic voting. Think again of our example of a single-transferable vote.

TABLE 8.19. Strategic Voting with the STV

Voters	**1st**	**2nd**	**3rd**
9	*w*	*x*	*y*
6	*x*	*y*	*w*
2	*y*	*x*	*w*
4	*y*	*x*	*w*
5	*x*	*y*	*w*

Suppose the *w* voters realize that *w* simply can't win and that if they report their true preferences, the result, we have seen, is *x* winning, their third choice. If they vote for *z* rather than *w* in the first round, however, we have seen that *z* will go on to beat *x* if it can survive the first round, so they get their second rather than their third choice by voting against their first choice. In elimination procedures it is often the case that, if one opponent who would have defeated your candidate in the final round can be eliminated early on, your candidate can go on to win.

Logrolling: Vote Trading in Legislatures

Another case of voting strategically occurs when legislators vote against their preferences on Issue 1 (which they care less about) to achieve their preferred outcome on Issue 2 (which they care more about). Sometimes this is called "logrolling" or "vote trading."[31] Assume that we have two issues to be decided in a legislature split into three factions.

x: whether to provide increased federal funds for urban renewal projects, which will mostly benefit the Rustbelt Areas. Someone *against* this we can say is for not-*x*, which will be designated ¬*x*.
y: whether to increase expenditure for defense. Those in the Sunbelt Areas will primarily benefit. Someone against this is for ¬*y*.

Our three groups are, then: Rustbelters, Sunbelters, and Others. Each has 100 votes in our imaginary legislature. Their preferences are ordered in table 8.20.

TABLE 8.20. Logrolling Preferences

Rustbelters	**Sunbelters**	**Others**
x & ¬*y*	¬*x* & *y*	¬*x* & ¬*y*
x & *y*	*x* & *y*	¬*x* & *y*
¬*x* & ¬*y*	¬*x* & ¬*y*	*x* & ¬*y*
¬*x* & *y*	*x* &¬*y*	*x* & *y*

According to the *Rustbelters'* preferences, they want urban aid (*x*), but they don't want to also incur the costs of increased defense spending (*y*). Perhaps they hope to fund the aid to urban areas with defense cuts. But they would

prefer both the urban aid and increased defense spending (x & y) to no urban aid at all ($\neg x$). If they can't get their urban aid, they want to at least hold the line on federal spending so that their economies might improve. The last thing they want is increased federal spending and no urban aid ($\neg x$ & y).

According to the *Sunbelters'* preferences, they not only want to resist defense cuts, but also want some increases in defense spending to spur growth in their states, many of which rely on defense contractors. They don't want to increase the federal deficit with urban renewal expenditure (x); however, they would prefer to have both urban renewal and increased defense spending to no increase in defense spending (x & y). But if they can't get the increase they want in defense spending, they certainly want to at least hold the line on other new expenditures, so they can go back to their districts arguing that they were in favor of keeping the lid on federal expenditure ($\neg x$ & $\neg y$). The worst thing that could happen for them is for the Rustbelters to get their urban renewal aid, thus increasing the deficit but leaving them empty-handed (x &$\neg y$).

The *Others* are not enamored of either project; they are worried about the deficit and see no pressing need for either expenditure ($\neg x$ & $\neg y$). But if there has to be an expenditure, it always sits better with their constituents if it is for defense ($\neg x$ & y). They don't, however, want to fund both projects at the same time—that would lead to terrible strains on the budget; so they prefer to fund the Rustbelters alone rather than to pay for both projects.

If we suppose sincere voting on the issues separately, both x and y are defeated, leaving us with ($\neg x$ & $\neg y$). However, there is a clear possibility here for vote trading since the Rustbelters will agree to vote for the defense spending when that vote comes up (even though they prefer $\neg y$ to y), and the Sunbelters will agree to vote for urban aid when that vote comes up (even though they prefer $\neg x$ to x) in order to get the conjunction of both bills rather than neither. If they do so, the decision in favor of neither program ($\neg x$ & $\neg y$) is changed to (x & y) with both succeeding. They each get their second preference instead of their third. Here, as in the case of the w voters in the STV case, the strategic vote against their first preference gained them their second choice. Yet in one sense, society ends up with a result that would be clearly rejected if the votes were not linked. As we shall see next, this dynamic might lead to the society making choices that are worse for everyone.

These, of course, are just examples. However, Allan Gibbard and Mark Allen Satterthwaite, using versions of Arrow's theorem, have shown elsewhere that, roughly, all nonrandom, nondictatorial voting procedures that are sensitive to the voter's full statement of preferences are subject to strategic manipulation.[32] Thus, to vote over three options is to open the door to strategic manipulation.

Conclusion

In this chapter we looked at the logic of collective choice. We saw that there are important differences between social choice that aims at generating a complete representation of a social welfare function compared to one that aims only at generating a determinate social choice. The first approach runs afoul of Arrow's theorem, which shows that a social welfare function can only be socially rational if it is dictatorial. This poses problems for certain conceptions of collective choice and for populist theories of democracy. Collective choice rules have their own set of challenges. We saw that there is no generally acceptable voting rule that satisfies all of our notions of what collective choice should be. Further, collective choice rules are manipulatable, susceptible to strategic voting and agenda manipulation. Although many have argued that the results of social choice theory challenge the very idea of democracy, the pure logic of collective choice is too formal and lacking in substance to allow us to draw such sweeping conclusions. In the next chapter, we add substance to this model and show some of the implications for political theory.

Discussion Questions

(1) What is most important to democracy: that it is preference transforming or preference respecting?
(2) Does it make sense to talk of "collective rationality"? Was Buchanan right?
(3) Is May's theorem important for democracy? Do you have doubts about any of his conditions?
(4) What do you find confusing about Arrow's Theorem?

(5) Provide three examples of political disputes which you think are critically multi-dimensional.
(6) Do you think voters ought to vote strategically?
(7) Is vote trading in a legislature ethical? Is it efficient?
(8) Defend what you think is the best rule for deciding among three or more options against what you think are the three most serious objections.
(9) Did Levine and Plott do anything wrong? Is there an ethical way to set an agenda that is not biased?

---

1. Elster, "The Market and the Forum: Three Varieties of Political Theory," 111.
2. Elster, "The Market and the Forum: Three Varieties of Political Theory," 111.
3. Sunstein, Free Markets and Social Justice, 18ff.
4. Padover, Thomas Jefferson on Democracy, 15.
5. Buchanan, "Social Choice, Democracy, and Free Markets."
6. Barry, "Is Democracy Special."
7. May, "A Set of Independent Necessary and Sufficient Conditions."
8. We could also see this as a SWF over two options. The distinction between a social welfare function and a collective choice rule only becomes important when we have three or more options in our option set.
9. This, of course, follows if we are aiming at a social ordering.
10. Christiano, The Constitution of Equality, 290.
11. Goodin and List, "A Conditional Defense of Plurality Rule."
12. Arrow, Social Choice and Individual Values.
13. There are versions of Arrow's proof that provide a similar result for this "snapshot" case, where preferences are held fixed. The conditions, however, are different. In these versions of the proof, it is indeed the case that, looking at the preference profile, there is one person whose ordering is equivalent to the social ordering. It is interesting that there can be trivial dictators in this sense: if there is a person who is indifferent among all alternatives, then the social preference will imply her preference. She holds $x \succcurlyeq y$ and $y \succcurlyeq x$. If the social preference is xPy, that is consistent with her preference $x \succcurlyeq y$. If the social preference is yPx, that is consistent with her preference $y \succcurlyeq x$. And of course, if the social preference is for indifference, that is consistent with her preferences too. So, she always gets her preference satisfied. Even more obviously, on this "snapshot" view, if everyone happens to have the same preferences, then everyone is a Dictator (by virtue of Weak Pareto)! For an accessible fixed-profile proof, see Felman and Serrano, "Arrow's Impossibility Theorem."
14. The explication here follows Sen's in Collective Choice and Social Welfare, chap. 3.
15. See Craven, Social Choice, 36ff.
16. See List and Pettit, "Aggregating Sets of Judgments." The List-Pettit result is relevant to contexts in which the several reasons are relevant to a conclusion. Suppose a legal doctrine says that a person is responsible for his action ϕ at law if and only if conditions *A*, *B*, and *C* all hold. Instead of taking a vote among jurors and asking each for their final decision ("Do conditions *A*, *B*, and *C* all hold?"), we might take a series of votes: "Does condition *A* hold?," "Does condition *B* hold?," and "Does Condition *C* hold?" After taking a series of votes on these three questions, we might say that

the person is liable if and only if a majority of jurors have answered "yes" to all three questions. The answer generated by this procedure can differ radically from the answer that would be generated by asking each juror, "Do all three conditions hold?"

17. Sen, Collective Choice and Social Welfare, 14–18.

18. We are simplifying a bit here. If one adopts a condition slightly stronger than acyclicity, called "quasi-transitivity" (transitivity of strict preference only), then there will always be an oligarchic group who can impose their preferences; if one adds positive responsiveness to the conditions, then such a CCR will have a dictator. The results for a CCR that only meets acyclicity are weaker, uncovering a vetoer, who is very similar to an oligarch group. See Riker, Liberalism Against Populism, 131.

19. Table take from Riker, Liberalism Against Populism, 50.

20. *x* now gets 30 points and *z* gets 43.

21. In order to say society preferred *y* three times as much, we would need a ratio scale with a non-arbitrary zero point, but we can set that aside for now.

22. Technically, in Senate elections in Australia, voters can vote "above the line" or "below the line." The second option involves ranking at least twelve candidates and as many as all of them. Voting "above" the line involves ranking at least six parties. The preference rankings (not necessarily the candidates) of each party are then used instead of the direct ranking of candidates. Before the election, the parties make deals with one another, "trading preferences" to include other parties' candidates higher in their rankings.

23. See Sen, Collective Choice and Social Welfare, 48n.

24. See Riker, Liberalism Against Populism, 135.

25. This account comes from Riker, The Art of Political Manipulation, chap. 11. All quotes are from this chapter.

26. Mackie, Democracy Defended.

27. Mackie, Democracy Defended, 201.

28. Riker, The Art of Political Manipulation, 200. Mackie also disputes this claim of Riker's (and just about every other claim Riker ever made!). See Mackie, Democracy Defended, 201ff.

29. We follow the presentation of this classic example from Riker, The Art of Political Manipulation. You can get the example firsthand from the principles involved, however, in Levine and Plott, "Agenda Influence and Its Implications."

30. Riker, Liberalism Against Populism, 152–53.

31. On the ethics of this practice, see Thrasher, "The Ethics of Legislative Vote Trading."

32. Gibbard, "Manipulation of Voting Schemes: A General Result"; Satterthwaite, "Strategy-Proofness and Arrow's Conditions."