Arrow’s Theorem

Kenneth Arrow’s “impossibility” theorem—or “general possibility” theorem, as he called it—answers a very basic question in the theory of collective decision-making. Say there are some alternatives to choose among. They could be policies, public projects, candidates in an election, distributions of income and labour requirements among the members of a society, or just about anything else. There are some people whose preferences will inform this choice, and the question is: which procedures are there for deriving, from what is known or can be found out about their preferences, a collective or “social” ordering of the alternatives from better to worse? The answer is startling. Arrow’s theorem says there are no such procedures whatsoever—none, anyway, that satisfy certain apparently quite reasonable assumptions concerning the autonomy of the people and the rationality of their preferences. The technical framework in which Arrow gave the question of collective orderings a precise sense and its rigorous answer is now widely used for studying problems in welfare economics. The impossibility theorem itself set much of the agenda for contemporary social choice theory. Arrow accomplished this while still a graduate student. In 1972, he received the Nobel Prize in economics for among others these contributions.

1. The Will of the People?

Some of the trouble with the idea of collective orderings is visible in a simple but important example. Say there are three options *A*, *B* and *C* to choose among. There are three people *1*, *2* and *3* whose preferences among these alternatives will inform this choice, and they are asked to rank them by their own lights from better to worse. The individual orderings happen to be these:

1. *ABC*
2. *BCA*
3. *CAB*

That is, *1* prefers *A* to *B*, prefers *B* to *C,* and prefers *A* to *C*; person *2* prefers *B* to *C*, and so on. Now suppose that, taking the alternatives pair by pair, we derive collective preferences from these individual orderings by voting. We arrive in this way at a collective preference for *A* to *B*, since two people prefer *A* to *B*, but only one prefers *B* to *A*. Similarly, we arrive at a collective preference for *B* to *C*. We might therefore expect to find that there is a collective preference for *A* to *C*, but by this reckoning it is just the other way around since there are two who prefer *C* to *A*.

This is the "paradox of voting." Discovered by the Marquis de Condorcet (1785), it shows how a certain property that facilitates rational choice can fail to pass from individual preferences to the collective preferences derived from them. Voter *1* has *A* at the top of his personal ordering. So this voter's preferences can be *maximized*, by choosing *A*. We can maximize instead the preferences of *2*, or of *3*, by choosing instead their maxima, *B* and *C*. Pairwise majority voting does not result in a collective maximum, though. *A* isn't one because a majority prefers something else, *C*. Likewise, *B* and *C* are not collective maxima. Taken singly, the individuals' preferences lend themselves to maximization; but collectively they do not. This is because pairwise majority voting in this particular instance does not derive an *ordering* of the alternatives from the most preferred, collectively, to the least preferred. It derives a collective preference *cycle* in which each alternative is outdone by some other.

Are other procedures better than pairwise majority voting, or do all of the different ones have shortcomings of their own? Condorcet, his contemporary Jean Charles de Borda (1781), and later Charles Dodgson (1844) and Duncan Black (1948), among many others, all addressed this question by studying various particular aggregation procedures and comparing their properties. Arrow broke new ground by coming at it from the opposite direction. Starting with various properties that aggregation procedures might be expected to have, he asked which of them fill the bill. Arrow built some of the desiderata into the basic technical framework that he developed in order to study the problems of social choice. Among these is *Social Ordering*, which requires that the result of aggregation is always an ordering of the alternatives, never a cycle. Arrow's framework will be discussed in Section 2. Other desiderata are expressed by several "conditions" or constraints that Arrow imposed on aggregation procedures, which are the topic of Sections 3.1 and 4. Briefly, they are: *Unrestricted Domain*, which requires aggregation procedures to be ready for any individual preferences at all; *Weak Pareto*, which requires them to respect unanimous individual preferences; *Non-Dictatorship*, which rules out procedures by which the collective ordering must always agree with the strict preferences of some one person, the dictator, the preferences of others being considered only when the dictator is indifferent; and finally *Independence of Irrelevant Alternatives*, which requires that collective orderings be assembled, so to speak, one pair of alternatives at a time. To determine the collective preference among any two given alternatives, when this condition is satisfied, no further information about individual preferences is needed beyond preferences among these two. Arrow's theorem, stated in Section 3.2, tells us that no procedure for deriving collective orderings satisfies all the desiderata.

The tenor of Arrow's theorem is deeply antithetical to the political ideals of the Enlightenment. It appears that Condorcet's paradox is indeed not an isolated anomaly, the failure of one specific voting method. Rather, it manifests a much wider problem with the very idea of collecting many individual preferences into one; for, on the face of it anyway, what Arrow's theorem tells us is that there simply cannot be a common will of all the people concerning collective decisions, that assimilates the tastes and values of all the individual men and women who make up a society.

There are some who, following Riker (1982), have taken Arrow's theorem to show that democracy, conceived as government by the will of the people, is an incoherent illusion. Many others, though, have pointed out that certain conditions of the theorem sometimes are unreasonable, and their insights make the prospects for democratic social choice look much brighter. After presenting the theorem itself, this entry will consider some main points of critical discussion. Section 4 considers in detail the meaning and scope of Arrow's conditions, and Section 5 discusses aggregation procedures that are available when not all of them need be satisfied. Section 6 concludes with a brief overview of some proposals to analyze within Arrow's technical framework certain aggregation problems other than that of social preferences.

Amartya Sen once expressed regret that the theory of social choice does not share with poetry the amiable characteristic of communicating before it is understood (Sen 1986). Arrow's theorem is not especially difficult to understand and much about it is readily communicated, if not in poetry, then at least in plain English. Informal presentations go only so far, though, and where they stop sometimes misunderstandings start. This exposition uses a minimum of technical language for the sake of clarity.

2. Arrow's Framework

Arrow's problem of finding an aggregation procedure arises in connection with some given alternatives between which a choice is to be made. The nature of these alternatives depends on the kind of choice problem being studied. In the theory of elections, the alternatives are people who might stand as candidates in an election. In welfare economics, they are different social states such as distributions of income and labour requirements among the members of a society. The different alternatives under consideration are x, y, z,… and the set of them all is X. There are, as well, some people whose tastes and values are to inform the choice among these alternatives. The people are assumed to be finite in number and are enumerated $1,\ldots n$. Arrow's problem of finding an aggregation procedure arises, then, only *after* some alternatives and people have been fixed. It is for them that a procedure is sought. Crucially, though, this problem arises *before* relevant information about the people's preferences among the alternatives has been gathered, whether that is by polling or some other method. The question that the impossibility theorem answers is, more precisely, this: which procedures are there for arriving at a collective ordering of some given alternatives, no matter what some given people's preferences among these alternatives turn out to be?

In practice, meanwhile, we sometimes must select a social decision procedure without knowing exactly for which alternatives and people it will be used. In the recurring elections for some public office, for instance, there is a different slate of candidates each time, and a different population of voters, and we must use the same voting method to determine the winner, no matter who the candidates and voters are and no matter how many of them there happen to be. Such procedures are not directly available for study within Arrow's framework, with its fixed alternatives X and people 1, … n. Arrow's theorem is still relevant to them, though. It tells us that *even* when the alternatives and people are held fixed, then still there is no "good" method for deriving collective orderings. Now, if there is no good method for voting even once, with the particular candidates and voters who are involved on that occasion, then nor, presumably, is there a good method that can be used repeatedly, with different candidates and voters each time.

2.1 Individual Preferences

Arrow assumed that collective orderings will be derived, if at all, from *ordinal* information. This is the kind of information that is implicated in Condorcet's paradox of voting, in Section 1, where each person ranks the alternatives from better to worse but there is nothing beyond this about how strong anybody's preferences are, or how the preferences of one person compare in strength to those of another. In confining aggregation procedures to ordinal information, Arrow argued that "it seems to make no sense to add the utility of one individual, a psychic magnitude in his mind, with the utility of another individual" (Arrow

1963, 11). His point was that even if there are stronger and weaker preferences, and even if the strengths of people's preferences can somehow be measured, nevertheless ordinal information is all that needs to be taken into account because preferences are "interpersonally incomparable." Intuitively, what this means is that there is no saying how much more strongly someone must prefer one thing to another in order to make up for the fact that someone else's preference is just the other way around. Arrow saw no reason to provide aggregation procedures with information about preference strengths because he thought that they cannot put it to any meaningful use.

Accordingly, in Arrow's framework people's preferences are represented by relations R_i among the alternatives: xR_iy means, intuitively, that individual i either prefers alternative x to alternative y or else is indifferent between them. Each individual ordering R_i is assumed to be *connected* (for all alternatives x and y, either xR_iy, or yR_ix, or both) and *transitive* (for all x, y and z, if xR_iy and yR_iz, then xR_iz). That the relations have these structural properties is, for Arrow, a matter of the rationality of the individual preferences they represent; for further discussion, see the entries on Preferences and Philosophy of Economics. Connected, transitive relations are called *weak orderings*. They are "weak" in that they allow ties, or in this connection *indifference* between alternatives. A *preference profile* is a list $<R_1, R_2,\ldots R_n>$ of weak orderings of X, one for each of the people $1, \ldots n$. The list of three individual orderings in the paradox of voting is an example of a profile for the alternatives A, B, and C and people 1, 2, and 3. A profile is a representation of the individual preferences of everybody who will be consulted in the choice among the alternatives. It is in the form of profiles that Arrow's aggregation procedures receive information about individual preferences. Often it is convenient to write $<R_i>$ instead of $<R_1, R_2,\ldots R_n>$. Other profiles are written $<R^*_i>$, and so on.

One important theme in voting theory is the manipulation of outcomes by voters who strategically misrepresent their preferences. Matters relating to strategic voting will come up here (in Sections 4.5 and 5.2) but do not play a part in Arrow's presentation of the impossibility theorem, and were not dealt with seriously in the literature until after its publication. Arrow assumed that people report their preferences truthfully.

2.2 Multiple Profiles

Aggregation procedures in Arrow's framework must reckon with more than just one profile, representing *actual* preferences. They are required to derive collective orderings from many profiles, representing preferences that everybody *could* have.

Variety among preferences is the result, in Arrow's account of the matter, of the different standards by which we assess options. Our preferences depend on our "tastes" in personal consumption. Importantly, in Arrow's account, when it comes to social decisions they depend also on socially directed "values." Now, we are free to have a range of tastes, values, and preferences; and we are free, also, to have them independently of each other. Any one person can have a range of preferences, then, and for any given people there is a range of possible preference profiles. One profile $<R_i>$ represents the preferences of the people $1,\ldots, n$ among the alternatives X in, if you will, one possible world. Another profile $<R^*_i>$ represents preferences of the same people, and among the same alternatives, but in another world where their tastes and values are different.

Arrow's rationale for requiring aggregation procedures to handle many profiles was epistemic. As he framed the question of social choice, a procedure is sought for arriving at a collective ordering of some given alternatives on the basis of some given people's tastes and values. It is sought, though, before it is known just what these tastes and values happen actually to be. The variety among profiles to be reckoned with is, in his account of the matter, a measure of what is known or assumed about everybody's preferences *a priori*, which is to say before they have been elicited. When less is known there are more profiles that might have to be assimilated into a collective ordering; when more is known there are fewer of them. The variety among the profiles that must be reckoned with depends on the nature of the alternatives themselves, and on how preferences among them are formed. See section 4.1 for further discussion of this point.

There are other reasons for working with many profiles, even when the people's actual preferences can be known in advance. For example, we might like to know how the collective ordering derived using some procedure will vary with changes in the input. Serge Kolm (1996) has suggested that counterfactual preferences are relevant when we come to justify the use of one procedure rather than another. Sensitivity analysis, used to manage uncertainty about errors in input, and to determine which information is critical in the sense that the output turns on it, also requires that procedures handle various inputs.

Another advantage of working with many profiles is that it enables us to study "interprofile" conditions. These are constraints that coordinate collective orderings at several profiles. *Independence of Irrelevant Alternatives* is an example. It says that whenever everybody's preferences among two alternatives are the same in one profile as they are in another, the derived collective ordering must also be the same, as far as these alternatives are concerned. There is to be this similarity among derived collective orderings even as people's tastes and values change from one profile to the next. Sections 3.1 and 4.5 discuss in more detail the meaning of this constraint and the extent to which it is reasonable to impose it on aggregation procedures.

Ian Little raised the following objection in an early discussion of (Arrow 1951):

> If tastes change, we may expect a new ordering of all the conceivable states; but we do not require that the difference between the new and the old ordering should bear any particular relation to the changes of taste which have occurred. We have, so to speak, a new world and a new order; and we do not demand correspondence between the change in the world and the change in the order (Little 1952, pp. 423-424).

Little agreed with Arrow that there would be a different collective ordering were people's tastes different, but unlike Arrow he thought that it wouldn't have to be similar to the actual or current ordering in any special way. Little's objection was taken to support the "single profile" approach to social welfare judgments of Abram Bergson (1938) and Paul Samuelson (1947), and there ensued a debate about which approach was best, theirs or Arrow's. Arguably, what was at issue in this debate was not—or should not have been—whether aggregation procedures must handle more than a single profile, but whether there is to be any coordination among collective orderings at different profiles. This point has been made by among others Sen (1977) and Fleurbaey and Mongin (2005). Then the substance of Little's objection can be accommodated simply by refraining from imposing any interprofile constraints, and the multi-profile framework is the more general of the two. Be this as it may, it is the dominant framework in contemporary social choice theory.

2.3 Domains and Social Welfare Functions

A preference profile is said to be *admissible* if it is compatible with what is assumed or known about everybody's preferences before these have been elicited. The admissible profiles represent preferences that the people could have, and might turn out actually to have, so it is from these that we may expect an aggregation procedure to derive collective orderings. A *domain*, in Arrow's framework, is a set of admissible profiles, each concerning the same alternatives X and people $1, \dots n$. A *social welfare function* f assigns to each profile $<R_i>$ in some domain a weak ordering $f<R_i>$ of X. Intuitively, f is an aggregation procedure. It derives collective ordering $f<R_i>$ from individual preferences $<R_i>$.

Arrow established a convention that is still widely observed of using 'R' to denote the collective ordering derived from $<R_i>$. The social welfare function used to arrive at this ordering is, in his notation, left implicit. One advantage of writing '$f<R_i>$' instead of 'R' is that when we come to state the conditions of the impossibility theorem, in the next section, f will figure explicitly in them. This makes quite clear that what these conditions constrain is the functional relationship between individual orderings and collective orderings.

3. Impossibility

With the conceptual framework now in place, Section 3.1 sets out the "conditions" or constraints that Arrow imposed on social welfare functions, and Section 3.2 states the theorem itself. There is critical discussion of the conditions in Section 4.

Arrow's conditions often are called axioms, and his approach is said to be axiomatic. This might be found misleading. Unlike axioms of logic or geometry Arrow's conditions are not supposed to express more or less indubitable truths, or to constitute an implicit definition of some object of study. Arrow himself took them to be questionable "value judgments" that "express the doctrines of citizens' sovereignty and rationality in a very general form" (Arrow 1951, 31). Indeed, as we will see in Section 4, and as Arrow himself recognized, sometimes it is not even desirable that social welfare functions should satisfy all of them.

Arrow restated the conditions in the second edition of *Social Choice and Individual Value* (Arrow 1963). They appear here in the canonical form into which they have settled since then.

3.1 These Conditions…

A first requirement is that the social welfare function f can handle any combination of any individual preferences at all:

> *Unrestricted Domain* (U): The domain of f includes every list $<R_1, R_2,\dots R_n>$ of n weak orderings of X.

Condition U requires that f is defined for each "logically possible" profile of individual preferences. A second requirement is that, in each case, the collective preference derived by f is to be a weak ordering of the alternatives:

> *Social Ordering* (SO): For any profile $<R_i>$ in the domain, $f<R_i>$ is a complete, transitive ordering of X.

Arrow didn't state this condition separately but incorporated it into the definition of a social welfare function (see Section 2.3). To state the next requirements it is convenient to use some shorthand. For any given individual ordering R_i, let P_i be the *strict* or *asymmetrical* part of R_i: xP_iy if xR_iy but not yR_ix. Intuitively, xP_iy means that i really does prefer x to y, in that i is not indifferent between them. Similarly, let P be the strict part of the collective ordering $f<R_i>$. The next condition of Arrow's theorem is:

> *Weak Pareto* (WP): For any profile $<R_i>$ in the domain of f, and any alternatives x and y, if for all i, xP_iy, then xPy.

WP requires f to respect unanimous strict preferences. That is, whenever everyone strictly prefers one alternative to another, the collective ordering that f derives must agree. Let pairwise majority decision be defined more precisely, now, as follows: x is weakly preferred to y, collectively, if there are as many people who weakly prefer x to y as there are who weakly prefer y to x.[1] Obviously pairwise majority decision satisfies WP, and other familiar voting methods such as Borda counting satisfy it as well (see Section 5.2). WP requires that f is to this extent like them.

[1] Even more precisely, pairwise majority decision is that function f that assigns to a profile $<R_i>$ the relation R such that: xRy if and only if $|\{i: xR_iy\}| \geq |\{i: yR_ix\}|$.

The next condition ensures that the collective ordering is not based on the preferences of any one person. Let us call person d a *dictator* of f if for any alternatives x and y, and for any profile $<\ldots, R_d,\ldots>$ in the domain of f: if xP_dy, then xPy. Whenever the dictator strictly prefers one thing to another the collective ordering always agrees, no matter what anyone else prefers. The relevant condition is now:

Nondictatorship (D): f has no dictator.

To illustrate, pick some person, d, and from each profile $<R_i>$ in the domain take the ordering R_d representing the preferences of d. Now, in each case, simply let the collective ordering be that. In other words, for each profile $<R_i>$, let $f<R_i>$ be R_d. This social welfare function f bases the collective ordering entirely on the preferences of d, its dictator. It is intuitively undemocratic and D rules it out.

The remaining condition of Arrow' theorem is *Independence of Irrelevant Alternatives.* It requires that the collective ranking of a pair of alternatives is independent of anybody's preferences concerning alternatives other than them. Figuratively speaking, what this means is that a social welfare function must assemble collective orderings from individual orderings one pair of alternatives at a time. To state the Independence condition precisely, another piece of shorthand comes in handy. For any given relation R, and any set S, let $R|S$ be the *restriction of R to S*. It is that part of R concerning just the members of S.[2] A list of relations is restricted by restricting, in turn, each relation in the list. In the case of a profile $<R_1, R_2,\ldots R_n>$, the restriction to S, written $<R_1, R_2,\ldots R_n>|S$, is just $<R_1|S, R_2|S,\ldots R_n|S>$. For instance, the profile from the paradox of voting in Section 1 is, once again:

1. *ABC*
2. *BCA*
3. *CAB*

Its restriction to the set $\{A,C\}$ of alternatives is:

1. *AC*
2. *CA*
3. *CA*

The relevant condition is now:

Independence of Irrelevant Alternatives (I): For all alternatives x and y in X, and all profiles $<R_i>$ and $<R^*_i>$ in the domain of f, if $<R_i>|\{x,y\} = <R^*_i>|\{x,y\}$, then $f<R_i>|\{x,y\} = f<R^*_i>|\{x,y\}$.

I says that whenever two profiles $<R_i>$ and $<R^*_i>$ are identical, as far as some alternatives x and y are concerned, so too must the derived collective orderings $f<R_i>$ and $f<R^*_i>$ be identical, as far as x and y are concerned. For example, consider the profile:

1. *BAC*
2. *CAB*
3. *BCA*

Its restriction to the pair $\{A,C\}$ is identical to that of the profile of the paradox of voting. Suppose the domain of a social welfare function includes both of these profiles. To satisfy I, this function must derive from each one the same collective ordering of A with respect to C. The collective ordering of A and C is in this sense to be "independent" of anybody's preferences among either of them and the remaining "irrelevant" alternative B. The same is to hold for any two profiles in the domain, and for any other pair taken from the set $X = \{A, B, C\}$ of all alternatives. Some voting methods do not satisfy I (see Section 5.2),

[2] That is, $x\ R|S\ y$ holds if and only if xRy and both x and y are members of S.

but pairwise majority decision does. To determine the collective preference among *x* and *y*, by this method, you need no more than the individual preferences among *x* and *y*.

3.2 … are Incompatible

Arrow discovered that, except in the very simplest of cases, the four conditions on social welfare functions are incompatible.

> *Arrow's Theorem*: Suppose there are more than two alternatives. Then no social welfare function *f* satisfies *U*, *WP*, *D*, and *I*.

Arrow (1951, 1963) has the original proof. See among many other works (Kelly 1978), (Campbell and Kelly 2002), (Geanakoplos 2005) and (Gaertner 2009) for variants and different proofs.

4. The Conditions, again

Taken separately, Arrow's conditions do not seem so severe. Apparently, they ask of an aggregation procedure only that it will arrive at a collective ordering no matter what everybody prefers (*U*), will resemble certain democratic arrangements in some ways (*WP* and *I*), but will not resemble certain undemocratic arrangements in another way (*D*). Taken together, though, these conditions exclude all possibility of deriving collective orderings. It is time to consider them more closely.

4.1 Unrestricted Domain

Arrow's domain condition *U* says that the domain of the social welfare function includes every list of *n* weak orderings of *X*. For example, suppose the alternatives are *A*, *B*, and *C*, and that the people are *1*, 2, and 3. There are *13* weak orderings of three alternatives so the unrestricted domain contains *2197* (that is, 13^3) lists of weak orderings of *A*, *B*, and *C*. A social welfare function *f* for these alternatives and people, if it satisfies *U,* maps each one of these "logically possible" preference profiles onto a collective ordering of *A*, *B*, and *C*.

In Arrow's account, the different profiles in a domain represent preferences that the people might turn out to have. To impose *U*, on his epistemic rationale, amounts to assuming that they might have any preferences at all: it is only when their preferences could be anything that it makes sense to require the social welfare function to be ready for everything. Arrow wrote in support of *U*: "If we do not wish to require any prior knowledge of the tastes of individuals before specifying our social welfare function, that function will have to be defined for every logically possible set of individual orderings" (Arrow 1963, p. 24).

There have been misunderstandings. Some think *U* requires that the social welfare function can handle "any old" options. It does nothing of the sort. What it requires is that the social welfare function can handle the widest possible range of *preferences* among the options; and whether there happen to be many of these or only a few is beside the point: the domain of a social welfare function can be completely unrestricted even if there are in *X* just two options to choose among. One way to sustain this unorthodox understanding of *U* is, perhaps, to think of Arrow's *x*, *y*, *z*… not as alternatives properly speaking—not as candidates in elections, social states, or what have you—but as names or labels that *represent* these on different occasions for choosing. Then, it might be thought, variety among the options to which the labels can be attached will generate variety among the profiles that an aggregation procedure might be expected to handle. Blackorby, Donaldson and Bossert (2006) toy with this idea at one point, but they quickly set it aside. It does not seem to have been explored in the literature.

There is, of course, nothing to keep anyone from reinterpreting Arrow's basic notions including X in any way they like; a theorem is a theorem, and Arrow's theorem will still hold true no matter what interpretation is given to it. It is important to realize, though, that to interpret Arrow's x, y, z... as labels is not standard, and it can only make nonsense of much theory of social choice.[3]

Arrow already knew that U is a stronger domain condition than is needed for an impossibility result. The *free triple property* and the *chain property* are examples of weaker conditions that jointly replace U in some versions of Arrow's theorem (Campbell and Kelly 2002). These versions, being more informative, are, from a logical point of view, better. U is simpler to state than their domain conditions, though, and might be found more intuitive. Notice that the weaker domain conditions still require a lot of variety among profiles. A typical proof of an Arrow-style impossibility theorem requires that the domain is unrestricted with respect to some three alternatives. In this case there is always a preference profile like the one implicated in the paradox of voting in Section 1, from which pairwise majority decision derives a cycle.

Whether it is sensible to impose U or any other domain condition on a social welfare function depends very much on the particulars of the choice problem being studied. Sometimes, in the nature of the alternatives under consideration, and the way in which individual preferences among them are determined, imposing U certainly is not appropriate. If for instance the alternatives are different ways of dividing up a pie among some people, and it is known prior to selecting a social welfare function that they are selfish, each caring only about the size of his own piece, then it makes little sense to require of a suitable function that it can handle cases in which some prefer to have less for themselves than to have more. Arrow made this point as follows:

> [I]t has frequently been assumed or implied in welfare economics that each individual values different social states solely according to his consumption under them. If this be the case, we should only require that our social welfare function be defined for those sets of individual orderings which are of the type described; only such should be admissible (Arrow 1963, p. 24).

Section 5.1 considers some of the possibilities that open up when there is no need to reckon with all "logically possible" preference profiles.

4.2 Social Ordering

Condition *SO* requires that collective preferences be weak orderings, connective and transitive relations. Arrow did not state *SO* as a separate condition. As we saw in Section 2.3 he built it into the very notion of a social welfare function. In justification he offered only that *SO* is needed if collective preferences are to "reflect rational choice-making" (Arrow 1951, p. 19). Criticized by James Buchanan (1954) for transferring rationality properties of individual choice to social choice, he later claimed that transitivity of collective preferences is necessary for social choices to be independent of the path taken to them (Arrow 1963, p. 120). He did not develop this idea further.

Charles Plott (1973) later elaborated a suitable notion of path independence. Suppose we arrive at our choice by "divide and conquer:" first we divide the alternatives into some smaller sets—perhaps because these are more manageable—and choose from each one. Then we gather together the choices made from

[3] Of course, the *terms* 'x', 'y' and 'z' that Arrow used to refer to alternatives are representations, and they might be called labels. Similarly the numerals '1',… 'n' used to refer to the people might be called labels, but that is another thing. For example, interpreting X as a set of labels blurs the standard distinction between the Independence condition and Strong Neutrality (Section 4.5). Independence requires, for any given pair of alternatives, that the social welfare function maintains a certain consistency as we go from one profile to the next within its domain. Strong Neutrality requires this, and also consistency as we go from one pair of alternatives to the next. Now, if we think of x and y are labels that pick out different alternatives—properly speaking—in different profiles, then Independence already demands consistency from one pair of alternatives to the next.

the different sets, and choose again from among these. There are many ways of making the initial division, and a choice procedure is *path independent* if the final choice is independent of which division we start with (Plott 1973, p. 1080). In Arrow's account, social choices are made from some given "environment" S of feasible alternatives by maximizing a social preference relation R: the choice C(S) from among S is the set of those x within S such that for any y within S, $x\ R\ y$. It is not difficult to see how intransitivity of R can result in path dependence. Consider again the paradox of voting of Section 1. There is a strict social preference for A above B, and for B above C; but, contrary to transitivity, there is also a strict social preference for C above A. Starting with the division $\{\{A,B\}, \{B,C\}\}$, our choice from among $\{A,B,C\}$ will be $\{A\}$; but starting instead with $\{\{A,C\}, \{B,C\}\}$ we will arrive at $\{B\}$.

Plott's analysis reveals an important subtlety: it is only *strict* preferences that must be transitive to secure path independence. As far as Arrow's rationale is concerned, then, we can as Plott observes weaken *SO* by requiring only that social preferences determine a complete *quasi-ordering* of the alternatives—a relation whose strict component is transitive, but whose indifference component need not be. Sen (1969, p. 387) had already demonstrated the compatibility of this weaker requirement with all of Arrow's other conditions. Any hopes that this might by itself resolve Arrow's problem were dampened by Allan Gibbard's demonstration that the only "constitutions" that this weakening of *SO* makes available are what he called *liberum veto oligarchies*. There has in every case to be some set of individuals such that the society always strictly prefers one thing to another when everybody in this set does, and never strictly prefers one thing to another if someone in this set has the opposite strict preference (Gibbard 1969, 2014).

4.3 Weak Pareto

Condition *WP* requires that whenever everybody ranks one alternative strictly above another the collective ordering agrees. This has long been a basic assumption in welfare economics and might seem completely uncontroversial. For Arrow, some of its demands are "not debatable except perhaps on a philosophy of systematically denying people whatever they want" (Arrow 1951, p. 34).

However this may be, *WP* in combination with *U* tightly constrains the possibilities for social choice. This is evident not only from Arrow's theorem, but especially also from Sen's (1970) demonstration that these conditions conflict with the idea that for each person there is a personal domain of states of affairs within which his preferences are decisive, and must prevail whenever they conflict with the preferences of others. This important problem of the "Paretian libertarian" meanwhile has its own extensive literature.

We may think of *WP* as a vestige of what Sen called:

> *Welfarism*: The judgement of the relative goodness of alternative states of affairs must be based exclusively on, and taken as an increasing function of, the respective collections of individual utilities in these states (Sen 1979, p. 468).

WP asserts the demands of welfarism in the special case in which everybody's strict preferences coincide. Sen argued that even these apparently modest demands might in some instances be found excessive on moral grounds, in light of non-welfare information concerning the nature of the choices that people make, and the motivations underlying their individual rankings (Sen 1979, Section IV).

4.4 Non-Dictatorship

Arrow's Condition D requires that there is no one person d whose preferences are *decisive* in the sense that whenever one alternative strictly outranks another in d's ordering, it always does so in the collective ordering as well. Arrow's non-dictatorship condition is very suggestive and has attracted little comment in the literature.

Certainly D rules out some genuinely dictatorial arrangements, such as choosing some one person and identifying the social ordering with his preferences, no matter what anyone else prefers. But this condition overshoots. In addition, it rules out some arrangements that are not intuitively dictatorial at all. Technically, an Arrovian dictator is just someone whose preferences invariably are a subset of the society's preferences, and this by itself doesn't entail that they form a "basis" for these social preferences, or that the dictator has any "power" or "control" over social preferences. For example, someone will count as a dictator in this technical sense who, though having no influence on the matter, does have some way of knowing what is socially preferred, and in every case makes sure that his own preferences agree. Aanund Hylland made a related point while objecting to the unreflective imposition of D in single profile analyses of social choice: "In the single-profile model, a dictator is a person whose individual preferences coincide with the social ones in the one and only profile under consideration. Nothing is necessarily wrong with that; the decision process can be perfectly democratic, and one person simply turns out to be on the winning side on all issues" (Hylland 1986, p. 51, footnote 10).

Even pairwise majority voting, that paradigm of a democratic procedure, is in Arrow's technical sense sometimes a dictatorship. Consider Zelig. He has no tastes, values or preferences of his own but temporarily takes on those of another, whoever is close at hand. He is a human chameleon, the ultimate conformist.[4] Now, Zelig finds himself on a committee of three that will choose among several options using the method of pairwise majority voting and, given his peculiar character, the range of individual orderings that can arise is somewhat restricted. In each admissible profile, two of the three individual orderings are identical: Zelig's and that of whoever is seated closest to him at the committee meeting.[5] Suppose it turns out that Zelig strictly prefers one option x to another, y. Then someone else does too; that makes two of the three and so, when they vote, the result is a strict collective preference for x above y. The committee's decision procedure is, in Arrow's sense, a dictatorship, and Zelig is the dictator. But of course really Zelig is a follower, not a leader, and majority voting is as democratic as can be. It's just that this one mad little fellow has a way of always ending up on the winning side.

Arrow imposed D in conjunction with the requirement U that the domain is completely unrestricted. Perhaps D expresses something closer to its intended meaning then. With an unrestricted domain, a dictator, unlike Zelig, is someone whose preferences conflict with everybody else's in a range of cases, and it is in each instance his preferences that agree with the collective ordering, not theirs. In any case, whether D is appropriate depends on details of the choice problem at hand. Sometimes there is nothing undemocratic about having an Arrovian dictator.

4.5 Independence of Irrelevant Alternatives

Condition I is not Arrow's own formulation. It is a simpler one that is now the standard in expositions of the impossibility theorem. Arrow's version concerns choices made from within various "environments" S of feasible options by maximizing collective orderings:

> *Independence of Irrelevant Alternatives* (*choice version*): For all environments $S \subseteq X$, and all profiles $\langle R_i \rangle$ and $\langle R^*_i \rangle$ in the domain of f, if $\langle R_i \rangle|S = \langle R^*_i \rangle|S$, then $C(S) = C^*(S)$.

Here $C(S)$ is the set of those options within the set S that are, in the sense of the ordering $f\langle R_i \rangle$ derived from $\langle R_i \rangle$, as good as any other; and $C^*(S)$ stands for the maxima by $f\langle R^*_i \rangle$.

This is Arrow's Condition 3 (Arrow 1951, p. 27). Notice that Arrow didn't actually write 'all environments S'. He wrote 'a given environment S', which is ambiguous: in Arrow's formulation, condition 3 can be read

[4] This example derives from Woody Allen's spoof documentary *Zelig* (1983).
[5] The domain is restricted but it still contains many profiles. Suppose there are three alternatives to choose between, that the other two members of the committee might have any preferences at all, independently of one another, and that Zelig might end up with the preferences of either one. Then the decision procedure might be called on to handle any of *338* (or $13^2 \times 2$) preference profiles.

either as concerning all environments, or just some particular one. The universal reading is the one that is needed for an impossibility theorem, though, and it is this reading that secures logical equivalence with I. Crucially, every pair $\{x,y\}$ of alternatives is an environment.

Iain McClean (2003) finds a first statement of Independence, and appreciation of its significance, already in Condorcet (1785). Meanwhile, this condition has created much controversy and not a little confusion. Some of each can be traced to an example with which Arrow sought to motivate Independence. When one candidate in an election dies after polling, he wrote,

> …the choice to be made among the set S of surviving candidates should be independent of the preferences of individuals for candidates not in S… Therefore, we may require of our social welfare function that the choice made by society from a given environment depend only on the orderings of individuals in that environment (Arrow 1963, 27).

Arrow took this to amount to the choice version of Independence, for he continued:

> Alternatively stated, if we consider two sets of individual orderings such that, for each individual, his ordering of those particular alternatives in a given environment is the same each time, then we require that the choice made by society from that environment be the same when individual values are given by the first set of orderings as they are when given by the second (Arrow 1963, 27).

It is not clear why Arrow thought the case of the dead candidate has anything to do with different values and preference profiles. As he set the example up, it is natural to suppose that everybody's values and preferences stay the same while one candidate becomes unfeasible ("we'd all still prefer A, but sadly he's not with us any more"). Apparently, then, this example misses its mark. There has been much discussion of this point in the literature. See for example (Hansson 1973) for the widely accepted view that Arrow confused his independence condition for another, and for a contrary argument see (Bordes and Tideman 1991). For discussion of several independence notions whose differences have not always been appreciated, see (Ray 1973).

Sometimes, the following condition is called Independence of Irrelevant Alternatives:

> (I^*) For all x and y, and all $<R_i>$ and $<R_i^*>$ in the domain of f, if for all i: xR_iy if and only if xR^*_iy, then $x\, f{<R_i>}\, y$ if and only if $x\, f{<R_i^*>}\, y$.

If the intention is to express Arrow's condition this is a mistake because condition I^*, though similar in appearance to I, has a different content. I says that whenever everybody's preferences concerning a pair of options are the same in one profile as they are in another, the derived social preference for this pair must also be the same in the two profiles. This is not what I^* says, because the embedded antecedent 'for all i: xR_iy if and only if xR^*_iy', is satisfied not only when everybody's preferences among x and y are the same in $<R_i>$ as they are in $<R_i^*>$ but in many other instances as well. For example, suppose that in $<R_i>$ everybody is indifferent between some social state T and another state S (in which case for all i, both $T\,R_i\,S$ and $S\,R_i\,T$), while in $<R_i^*>$ everybody strictly prefers T to S (for all i, $T\,R^*_i\,S$ but not $S\,R^*_i\,T$). Then the antecedent 'for all i: $T\,R_i\,S$ if and only if $T\,R^*_i\,S$' is satisfied, although individual preferences among T and S are not the same in the two profiles. I^* constrains f in some instances when I does not and is the stronger of the two conditions.

The additional demands of I^* are sometimes excessive. Say T is the result of reforming some tried and true status quo S. We might reasonably favor reform if it is generally thought that change will be for the better, but not otherwise. In this case, we will seek a social welfare function f that derives a strict social preference for T above S if everybody strictly prefers T to S, but a strict preference for S above T if everyone is indifferent between these states. It is not difficult to see that I^* rules out every social welfare function that conforms to this desideratum.

There is another stronger condition that might be mistaken for Independence. What *I* requires, it might be said, is that the collective ranking of any given pair of alternatives, say social states, shall depend only on individual preferences among this pair. This is correct but leaves room for misunderstanding. *I* says that the only individual preferences that count are those concerning just these two social states. But that doesn't mean that individual preferences are all that counts. Various features of the two states may as far as *I* is concerned also factor into their collective ranking, such as the fact that one of them is the status quo, the equality among people there, respect for rights, and more. The doctrine that preferences are the *only* basis for comparing the goodness of social states is welfarism (briefly mentioned already in Section 4.3). An example illustrates how nasty welfarism can be.

In the status quo *S*, Peter is filthy rich and Paul is abjectly poor. Would it be better to take from Peter and give to Paul? Let *T* be the social state resulting from transferring a little of Peter's wealth to Paul. Paul prefers *T* to *S* ("I need to eat") and Peter prefers *S* to *T* ("not my problem"). This is one case. Compare it to another. Social state *T** arises from a different status quo *S**, also by taking from Peter and giving to Paul. This time, though, their fortunes are reversed. In *S** it is Peter who is poor and Paul is the rich one, so this is a matter of taking from the poor to give to the rich. Even so, we may assume, the pattern of Peter's and Paul's preferences is the same in the second case as it is in the first, because each of them prefers to have more for himself than to have less. Paul prefers *T** to *S** ("I need another Bugatti") and Peter prefers *S** to *T** ("wish it were my problem"). Since everybody's preferences are the same in the two cases, welfarism requires that the relative social goodness is the same as well. In particular, it allows us to count *T* socially better than *S* only if we also count *T** better than *S**. Whatever we think about taking from the rich to give to the poor, though, taking from the poor to give to the rich is quite another thing. As Samuelson said of a similar case, "[o]ne need not be a doctrinaire egalitarian to be speechless at this requirement" (Samuelson 1977, p. 83).

Condition *I* does not express welfarism. Applied to this example, *I* states that there is to be no change in the collective ordering of the status quo *S* with respect to the result *T* of redistribution unless Peter's preferences among these states change, or Paul's do (assume they are the only two involved). In this sense, the social ordering of these states may be said to depend "only" on individual preferences among them. *I* says the same about *S** and *T** or about any other pair of alternatives. But *I* is silent about how the social ordering of *S* with respect to *T* is to compare with that of *S** with respect to *T**. In particular, it leaves a social welfare function free to count *T* socially better than *S* (for *in*creasing equality), but *T** worse than *S** (for *de*creasing equality), even though the pattern of individual preferences is for the two pairs the same. Intuitively speaking, *I* allows the social welfare function to "shift gears" as we go from comparing one pair of social states to the next.

The condition that expresses welfarism is:

> *Strong Neutrality* (*SN*): For all alternatives x, y, z and w, and all profiles $<R_i>$ and $<R_i*>$: IF for all i: xR_iy if and only if $zR*_iw$, and yR_ix if and only if $wR*_iz$, THEN $x\, f<R_i>\, y$ if and only if $z\, f<R_i*>\, w$, and $y\, f<R_i>\, x$ if and only if $w\, f<R_i*>\, z$.

SN is more demanding than *I*.[6] *I* requires consistency for each pair of alternatives separately, as we go from one profile in the domain to the next. *SN* requires this and also consistency as we go from one *pair* to the next, whether that is within a single profile or among several different ones. This is how *SN* keeps non-welfare features out of the collective ordering: by compelling the social welfare function to treat any two pairs of alternatives the same way, if the pattern of individual preference is for each pair the same.

Since both *I** and *SN* are logically stronger than *I*, obviously a version of Arrow's theorem can be had using either one of them instead of *I*. Such a version will be less interesting, though—not only because it is weaker theorem but also because, as we have seen, the stronger conditions are often unreasonable.

---

[6] *SN* reduces *I* when we equate x with z, and equate y with w, so that we are in effect dealing with just a single pair of alternatives.

Given that the import and ramifications of Independence are not readily overseen, it is surprising just how little has been said to justify imposing this condition. In a typical statement of the impossibility theorem it is remarked that *I* seems reasonable but often that is all. We have already considered Arrow's attempt to motivate this condition using the example of the dead candidate. In the second edition of *Social Choice and Individual Values* he offered another rationale. There he argued that Independence, like the reliance on ordinal information about individual preferences, embodies the principle that welfare judgments are to be based only on observable behavior. Having expressed approval for Bergson's use of indifference maps, Arrow continued:

> The Condition of Independence of Irrelevant Alternatives extends the requirement of observability one step farther. Given the set of alternatives available for society to choose among, it could be expected that, ideally, one could observe all preferences among the available alternatives, but there would be no way to observe preferences among alternatives not feasible for society. (Arrow 1963, p. 110)

Arrow seems to be saying that social decisions have to be made on the basis of preferences for feasible alternatives because these are the only ones that are observable. Arguably, this is insufficient support. His independence condition, as we have seen, concerns *all* environments *S*. Arrow's observability argument, though, apparently just concerns some "given" feasible alternatives. See (Hansson 1973, p. 38) on this point. It is significant in this connection that Arrow's own formulation of Independence (Arrow 1951, Condition 3, p. 27) uses the ambiguous expression 'a given environment *S*' at the critical point. It is as if Arrow justifies Independence under one reading, but then relies on the much stronger universal reading to prove the theorem.

Gerry Mackie (2003) argues that there has been equivocation on the notion of irrelevance. It is true that we often take nonfeasible alternatives to be irrelevant. That presumably is why in elections we do not ordinarily elicit voter preferences for dead people, by putting their names on ballots along with those of live candidates. But *I* also excludes from consideration information on preferences for alternatives that, in an ordinary sense, *are* relevant. An example illustrates this point. Al Gore, George W. Bush, and Ralph Nader ran in the United States presidential election of 2000. Say we want to know whether there was a collective preference for Gore above Bush. Independence requires that this question be answerable independently of whether the people preferred either of them to, say, Abraham Lincoln, or preferred George Washington to Lincoln. This seems reasonable enough. Neither Lincoln nor Washington ran for President that year. They were, intuitively, irrelevant alternatives. But Independence also requires that the ranking of Gore with respect to Bush should be independent of voters' preferences for Nader, and this is different because he *was* on the ballot and, in the ordinary sense, he was a relevant alternative to them. Certainly Arrow's observability criterion does not rule out using information on preferences for Nader. They were as observable as any in that election.

A different rationale has been given for imposing Independence specifically in the case of voting. Many voting procedures are known to present opportunities for voters to manipulate outcomes by misrepresenting their preferences. Section 5.2 discusses the example of Borda counting, which allows voters to promote their favorite candidates by strategically putting others' favorites at the bottom of their lists. Borda counting, it will be seen, violates *I*. Proofs of the Gibbard-Sattherthwaite theorem (Gibbard 1973, Sattherthwaite 1975) associate vulnerability to strategic voting systematically with violation of *I*, and Iain McClean argues on this ground that voting methods ought to satisfy this condition: "Take out [Independence] and you have gross manipulability" (McClean 2003, p. 16).

Arrow himself did not in the end think that Independence must constrain all social welfare functions. He allowed that information relating even to *intuitively* irrelevant alternatives might sometimes be used for making social decisions:

> The austerity imposed by this condition is perhaps stricter than necessary; in many situations, we do have information on preferences for nonfeasible alternatives. It can be argued that, when available, this information should be used in social choice (Arrow 1963, p. 110).

Many have agreed that *I* constrains social welfare functions too tightly. Section 5.2 discusses two kinds of preference information that can be used in social choice but which *I* rules out: information about the *positions* of alternatives in individual orderings, and information about the *fairness* of social states.

5. Possibilities

Arrow's theorem, it has been said, is about the impossibility of trying to do too much with too little information about people's preferences. This remark points out two main avenues leading from Arrovian pessimism toward a sunnier view of the possibilities for collective decision making: not trying to do so much, and using more information to do it with. One way of asking less of social welfare functions is to soften the demand of the domain condition *U* that there be a collective ordering for each "logically possible" list of individual orderings. Section 5.1 develops this "escape route" in more detail. Alternatively, by loosening the independence constraint *I* we can release to social welfare functions more of the information carried by individual orderings. See Section 5.2. Finally, we can make available information about the strength of individual preferences that is not represented in Arrow's framework. This escape route is the topic of Section 5.3.

5.1 Domain Restrictions

Sometimes, in the nature of the alternatives and how individual preferences are determined, not all individual preferences can arise. When studying such a case within Arrow's framework there is no need for the social welfare function to derive a collective ordering from each and every *n*-tuple of individual orderings. Some but not all profiles are admissible, and the domain is said to be *restricted*. In fortunate cases it is possible to find a social welfare function that meets all assumptions and conditions of Arrow's theorem—apart, of course, from *U*. Such domains are said to be *Arrow consistent*. This Section considers some important examples.

Take to begin with a profile in which most voters completely agree:

1. *ABC*
2. *ABC*
3. *CBA*

Reckoning the collective preference by pairwise majority decision, we derive the ordering of the majority: *ABC*. Consider now a domain that is made up entirely of such profiles in which two of the three voters have the same ordering. On such a domain, pairwise majority decision is a social welfare function. And it is nondictatorial one as well if the domain, though restricted, still retains a certain variety among profiles. In the above profile, *3* strictly prefers *B* to *A*. The collective ordering reckoned by majority decision ranks *A* above *B*, though, so *3* is not a dictator provided this profile is in the domain. If each of the three voters disagrees in this way with both of the others, in some or other profile, pairwise majority decision satisfies *D*. (The case of Zelig in Section 4.4 is not like this because, being a conformist, his preferences never conflict with both of the others'.) Pairwise majority decision always satisfies *WP* and *I*. On such a domain, then, it satisfies all of Arrow's non-domain conditions. The domain is Arrow consistent.

Full *identity* of preferences is not needed for Arrow consistency. Sometimes it is sufficient that everybody's preferences are in a certain sense similar, even if no two voters ever entirely agree. An example illustrates the important case of *single peaked* domains.

Suppose three bears get together to decide how hot their common pot of porridge will be. Papa bear likes hot porridge, the hotter the better. Mama bear likes cold porridge, the colder the better. Baby bear most likes warm porridge; hot porridge is next best as far as he is concerned ("it will always cool off"), and he doesn't like cold porridge at all. The bears' preferences among hot, warm and cold porridge can be represented as a list of orderings:

Papa: *Hot Warm Cold*

Mama: *Cold Warm Hot*

Baby: *Warm Hot Cold*

Or else they can be pictured like this:

[Picture the bears' preferences]

The bear's preference profile is *single peaked*. Each bear has a "bliss point" somewhere along the ordering of the alternatives by their temperature; and each bear likes alternatives less as we move along this common ordering away from the bliss point, on either side. Single peaked *preferences* arise naturally with respect to various dimensions such as the position of candidates in an election on the left-right political spectrum, the cost of alternative public works and so on. Single peaked *profiles*, in which everybody's preferences are single peaked with respect to *the same* ordering of the alternatives, arise naturally when everybody agrees about what it is in the options that matters—their temperature, their left-right orientation, their cost—even if there is no further consensus about which options are better than which.

Duncan Black (1948) showed that if the number of voters is odd, and their preference profile is single peaked, pairwise majority decision always delivers up a weak ordering. Given an odd number of individuals, then, and a domain made up entirely of single peaked profiles, pairwise majority decision is a social welfare function.[7] Furthermore, as Black showed, the maximum of the ordering obtained by pairwise majority decision is the bliss point of the *median* voter—the voter whose bliss point has, on the common ordering, as many voters' bliss points to one side as it has to the other. In the example this is Baby bear, and warm porridge is the collective maximum. Singlepeakedness with respect to a meaningful underlying ordering brings the possibility of compromise.

Provided a single peaked domain is sufficiently inclusive (so that for each *i* there is within the domain some profile in which *i* is not the median voter) pairwise majority decision satisfies *D*. It always satisfies *WP* and *I*, so such a domain is Arrow consistent.

Pairwise majority decision does not always result in orderings when the number of voters is even. It is not in general a social welfare function, even if the domain is single peaked. For example, suppose there are just two voters and that their individual orderings are:

1. *CAB*
2. *BCA*

Pairwise majority decision derives from this profile a weak social preference for *A* to *B*, for there is one who weakly prefers *A* to *B,* and one who weakly prefers *B* to *A*. Similarly, it derives a weak social preference for *B* to *C*. Transitivity of the social ordering requires a weak collective preference for *A* to *C*, but there is none. On the contrary, by this reckoning there is a strict social preference for *C* above *A*, since that is the unanimous preference of the two voters. But this profile is single peaked with respect to the ordering: B-C-A:

[Picture the single peaked profile]

[7] Since pairwise majority decision does not derive a weak ordering from the profile that gives rise to Condorcet's paradox, in Section 1, it follows that this profile is not single peaked. It is a simple but useful exercise to verify, by checking the different ones, that there is no linear ordering of the alternatives with respect to which all three rankings of that profile are single peaked.

Majority decision with "phantom" voters yields a suitable social welfare function when the number of people is even. Let there be $2n$ people, and let each profile in the domain be single peaked with respect to the same ordering of the alternatives. Let R_{2n+1} be an ordering that also is single peaked with respect to this common ordering. R_{2n+1} represents the preferences of a "phantom" voter. Now expand any profile $<R_1,\ldots R_{2n}>$ in the domain into $<R_1\ldots R_{2n}, R_{2n+1}>$, by adding R_{2n+1}. The set of all the expanded profiles is a single peaked domain, and because the people together with the phantom are odd in number Black's result applies to it. Let g be pairwise majority rule for the expanded domain. We obtain a social welfare function f for the original domain by assigning to $<R_1,\ldots R_{2n}>$ the collective ordering that pairwise majority decision assigns to its expansion $<R_1\ldots R_{2n}, R_{2n+1}>$. That is, we set $f<R_1,\ldots R_{2n}>$ equal to $g<R_1\ldots R_{2n}, R_{2n+1}>$. This f satisfies *WP* because there are more "real" voters than there are phantoms ($2n$ to 1; we could have used any odd number of phantoms smaller than $2n$); f satisfies *I* because the phantom ordering is the same in all profiles of the expanded domain. If the domain included sufficient variety among profiles then f satisfies *D* and the domain is Arrow consistent. Moulin (1980) introduced phantom voters into single peaked domains to characterize a class of voting schemes that are non-manipulable, in that they do not provide opportunities for strategic voting.

Domain restrictions have been the focus of much research in recent decades. Gaertner (2001) provides a general overview of domain restrictions and the possibilities for social choice that they open up. Le Breton and Weymark (2006) survey work on domain restrictions that arise naturally when analyzing economic problems in Arrow's framework. Miller (1992) argued that deliberation can facilitate rational social choice by transforming initial preferences into single peaked preferences, and List et al. (2013) present empirical evidence that sometimes it does indeed bring about such an homogenization of individual preferences.

As Samuelson described it, the single profile approach might seem to amount to the most severe of domain restrictions: "one and only one of the… possible patterns of individuals' orderings is needed … From *it* (not from each of them all) comes a social ordering" (Samuelson 1967, pp. 48-49). According to Sen (1977), though, the Bergson-Samuelson social welfare function has more than a single profile in its domain. It has according to him a completely unrestricted domain, for while according to Samuelson only one profile is needed "it could be *any* one" (Samuelson 1967, p. 49). What distinguishes this approach, on Sen's way of understanding it, is that there is no coordinating the behavior of the social welfare function at several different profiles, by imposing on it interprofile conditions such as *I* and *SN* (see Section 4.5). Either way, though, and just as Samuelson insisted, Arrow's theorem does not apply. Either *U* is inappropriate (if there is a single profile in the domain) or else *I* is inappropriate (if there are no interprofile constraints).

Certain impossibility theorems that are closely related to Arrow's have been thought relevant to single-profile choice even so. These theorems do not use Arrow's interprofile condition *I* but use instead an *intra*profile neutrality condition. This condition says that whenever within any *single* profile the pattern of individual preferences for one pair x,y of options is the same as for another pair, z,w, the collective ordering derived from this profile must also be the same for x,y as it is for z,w:

> *Single-Profile Neutrality* (*SPN*): For any $<R_i>$, and any alternatives x, y, z and w: IF for all i: xR_iy if and only if zR_iw, and yR_ix if and only if wR_iz, THEN $x\, f<R_i>\, y$ if and only if $z\, f<R_i>\, w$, and $y\, f<R_i>\, x$ if and only if $w\, f<R_i>\, z$.

SPN follows from the strong neutrality (*SN*) condition of Section 4.5, on identifying $<R_i>$ with $<R_i^*>$. Parks (1976), and independently Kemp and Ng (1976), showed that there are "single profile" versions of Arrow's theorem using *SPN* instead of *I*. Such theorems were supposed to undermine the Bergson-Samuelson approach by showing that it faces essentially the same aggregation problem that Arrow discovered. Of course, these theorems are no more threatening than their conditions are reasonable, and *SPN* is objectionable for the same reason that *SN* is: both exclude the use of non-welfare information that is relevant to the comparison of social states from an ethical standpoint. Samuelson (1977) ridiculed the notion that *SPN* is reasonable using an example about redistributing chocolate, similar to that of Peter and Paul in Section 4.5.

5.2 More Ordinal Information

Arrow's independence condition *I* limits what information may be used for what. Speaking figuratively, it requires that when the social welfare function goes about the work of assembling collective orderings from individual orderings, it must take each pair of alternatives separately, paying no attention to available information about preferences for alternatives other than them. Some aggregation procedures work this way. Pairwise majority decision does: it counts *x* as weakly preferred to *y* if as many people weakly prefer *x* to *y* as the other way around, and plainly there is no need to look beyond individual preferences for these two alternatives to find this out. All such procedures ignore information that could be used in social decision making. This section considers two kinds of information that they ignore: *positional* information and information about the *fairness* of social states.

Positional voting methods take into account where the candidates come in the different individual orderings—whether they come first, second, … or last. *Borda counting* is an important example. Named after Jean-Charles de Borda, a contemporary of Condorcet, it had already been proposed in the 13th Century by Ramon Lull. Nicholas of Cusa in the 15th Century recommended it for electing Holy Roman Emperors. Borda counting is used in some political elections and on many other occasions for voting, in clubs and other organizations. Consider the profile:

1. *ABCD*
2. *BACD*
3. *BACD*

Let each candidate receive four points for coming first in some voter's ordering, three for coming second, two for a third place and a single point for coming last; the alternatives then are ordered by the total number of points they receive, from all the voters. The Borda count of *A* is then *10* (or *4+3+3*) and that of *B* is *11* (*3+4+4*), so *B* ranks strictly above *A* in the collective ordering. (This method applies with the obvious adaptation to any election with a finite number of candidates.)

But suppose voter *1* moves *B* from second place to last on his own list, and we have the profile:

1. *ACDB*
2. *BACD*
3. *BACD*

Now candidate *B* receives just *9* points (*1+4+4*). *A* receives the same *10* as before, though, and has come to outrank *B* in the collective ordering. This example illustrates two important points. First, Borda counting does not satisfy Arrow's condition *I*, since while each voter's ranking of *A* with respect to *B* is the same in the two profiles, the collective ordering of this pair is not the same. Second, this method provides opportunities for voters to manipulate outcomes by strategic voting. If everybody's preferences happen to be as in the first profile, voter *1* can do better, by his own lights, to misrepresent his preferences by putting *B* at the bottom of his list, for in this way he can promote his own favorite, *A*, to the top of the collective ordering (provided the other voters do not understand what he is up to and adjust their own rankings accordingly, putting *A* at the bottom). This susceptibility of the Borda method to strategic voting has long been known. When this matter was raised as an objection Borda's indignant retort is said to have been that his scheme was only intended for honest men. Lully and Nicholas of Cusa recommended oaths to tell the truth and stripping oneself of sins before voting.

For further discussion of positionalist voting methods see the entries on Voting Methods and Social Choice Theory, and for an analytical overview see Pattanaik's (2002) handbook article. Barberà (2010) reviews what is known about strategic voting.

Mark Fleurbaey (2007) has shown that social welfare functions need more ordinal information than *I* allows them if they are to respond appropriately to the fairness of social states. He gives the example of

Ann, who has ten apples and two oranges, and Bob, with three apples and eleven oranges. This allocation is envy free if, intuitively, she would be at least as happy with his basket of fruit as she is with her own, and he would be as happy with hers. Let the distribution of fruit in one social state S be as described, and consider the state S^* in which the allocations are reversed. That is, in S^* it is Ann that has three apples and eleven oranges, while Bob has ten apples and two oranges. More technically, S is *envy free* if Ann weakly prefers S to S^*, and Bob does too. Plainly, the envy freeness of social states is a matter of individual preferences. In general, it will vary from one profile in the domain of a social welfare function to the next.

We might expect that, other things being equal, the envy freeness of a social state will promote it in the social ordering above an alternative that is not envy free. But, as Fleurbaey has shown, I does not allow this. Starting from the status quo S, consider whether it would be better, socially, to take an apple and an orange from Bob and give both of them to Ann. Let T be the state arising from this transfer. We may assume for the sake of the example that Ann always strictly prefers having more for herself to having less, and that Bob's preferences are similarly self-serving, so that in all admissible profiles Ann strictly prefers T to S, while Bob strictly prefers S to T. Their preferences among these states are opposite and, absent relevant differences between the states, it would appear that there is no basis for a social preference one way or the other. This is where fairness might be expected to come in. Relative to one profile of individual preferences—in which both Ann and Bob weakly prefer S to S^*—the status quo S is envy free but T is not. A social welfare function that promotes envy-freeness will come out against transfer by ranking S strictly above T. Relative to another profile though, in which T is the envy free state —Ann and Bob weakly prefer T to the result T^* of a swap—instead T will end up strictly above S in the social ranking. In direct conflict with I, the social preference among S and T will switch as we go from one profile to the other, although all individual preferences among S and T stay the same. The social ranking of S and T turns on preferences for the "irrelevant" S^* and T^* because the fairness of these states does.

Fleurbaey recommends a weaker independence condition, attributing it to Hansson (1973) and to Pazner (1979):

> *Weak independence*: Social preferences on a pair of options should only depend on the population's preferences on these two options and on what options are indifferent to each of these options for each individual (Fleurbaey 2007, p. 23).

Fleurbaey (2007) discusses social welfare functions satisfying weak independence together with Arrow's conditions apart, of course, from I.

5.3 Cardinal Information

Another way to have collective orderings in spite of Arrow's theorem is to derive them from more information about individual preferences than is available in Arrow's profiles. Sen in particular has argued that social decisions should be based on richer information than just *orderings* of the alternatives according to individual preferences. Restricting the domains of social welfare functions (Section 5.1) and allowing them to use more ordinal information (Section 5.2) are ways of getting around Arrow's theorem while working within his framework. Developing this idea requires modifications to it.

Sen (1970) extended Arrow's framework by representing the preferences of individuals i not as orderings R_i but as cardinal utility functions U_i that map the alternatives onto real numbers: $U_i(x)$ is the utility that i obtains from x. A utility function U_i contains at least as much information as an individual preference ordering because we can reduce it to an ordering by putting xR_iy if $U_i(x) \geq U_i(y)$. There is in general more information, though, because we cannot always go in reverse: different utility functions reduce to the same ordering. A preference profile in Sen's framework is a list $<U_1, \dots U_n>$ of utility functions, and a domain is a set of these. An aggregation function, now a social welfare *functional*, maps each profile in some domain onto a weak ordering of the alternatives.

Sen showed how to study various assumptions concerning the measurability and interpersonal comparability of utilities by coordinating the collective orderings derived from profiles that, depending on these assumptions, carry the same information. For instance, ordinal measurement with interpersonal noncomparability—built into his framework by Arrow as fixtures—amount, in Sen's more flexible set up, to a requirement that the same collective ordering is to be derived from any utility profiles that reduce to the same list of orderings. At the other extreme, utilities are measured on a ratio scale with full interpersonal comparability if those profiles yield the same collective ordering that can be obtained from each other by rescaling, or multiplying all utility functions by the same positive real number. Sen explored many combinations of such assumptions.

One important finding was that having cardinal utilities is not by itself enough to avoid an impossibility result. In addition, utilities have to be interpersonally comparable. Intuitively speaking, to put information about preference strengths to good use it has to be possible to compare the strengths of different individuals' preferences. See (Sen 1970, Theorem 8*2). Interpersonal comparability opens up many possibilities for aggregating utilities and preferences. Two important ones can be read off from classical utilitarianism and Rawls's difference principle. See (Sen 1970, 1979) and also the entry on Social Choice Theory for details.

6. Extensions and Reinterpretations

The Arrow-Sen framework lends itself naturally to the study of a range of aggregation problems other than that of collective preferences. This Section briefly describes some of them.

6.1 Judgment Aggregation

On an epistemic conception, the value of democratic institutions lies, in part, in their tendency to arrive at the truth in matters relevant to public decisions (see Estlund 2008, but compare Peter 2011). This idea receives some support from Condorcet's jury theorem which tells us, simply put, that if individual people are more likely than not to judge correctly in some matter of fact, independently of one another, then the collective judgment of a sufficiently large group, arrived at by majority voting, is almost certain to be correct (Condorcet 1785). The phenomenon of the "wisdom of crowds," facilitated by cognitive diversity among individuals, provides further and arguably better support for the epistemic conception (Landemore 2012). But there are theoretical limits to the possibilities for collective judgment on matters of fact. Starting with Kornhauser and Sager's (1986) discussion of group deliberation in legal settings, work on the theory of judgment aggregation has explored paradoxes and impossibility theorems closely related to those that Condorcet and Arrow discovered in connection with preference aggregation. See (List 2012) and the entry Social Choice Theory for overviews of this rapidly developing field of research.

6.2 Multi-Criterial Decision

In many decision problems there are several criteria by which to compare alternatives and, putting these criteria in place of people, it is natural to study such problems within the Arrow-Sen framework. Arrow's theorem, if analogues of its various assumptions and conditions are appropriate, then tells us that there is no procedure for arriving at an "overall" ordering that assimilates different criterial comparisons.

Kenneth May (1954) used Arrow's framework to study the determination of individual preferences. It had been found experimentally that people's preferences, elicited separately for different pairs of options, often are cyclical. May explained this by analogy with the paradox of voting as the result of preferring one alternative to another when it is better by more criteria than not. More generally, he reinterpreted Arrow's theorem as an argument that intransitivity of individual preferences is to be expected when different criteria "pull in different directions." Susan Hurley (1985, 1989) considered a similar problem in practical deliberation when the criteria are moral values. She argued that Arrow's theorem does not apply in this case. One strand of her argument is that, unlike a person, a moral criterion can rank any given alternatives just one way. It cannot "change its mind" about them (Hurley 1985, p. 511), and this makes it inappropriate to impose the analogue of the domain condition *U* on procedures for weighing moral reasons.

Arrow's framework has also been used to study multicriterial evaluation in industrial decision making (Arrow and Raynaud, 1986) and in engineering design (Scott and Antonsson 2000; compare Franssen, 2005).

There are multicriterial problems in theoretical deliberation as well. Okasha (2011) uses the Arrow-Sen framework to study the problem of choosing among rival scientific theories by criteria including fit to data, simplicity, and scope. He argues that the impossibility theorem threatens the rationality of rational theory choice. See (Morreau 2014) for a reason to think that it does not apply to this problem, and (Morreau 2014) for a demonstration that impossibility theorems relevant to single profile choice (see Sections 2.2 and 5.1) might sometimes apply even so. In related work, Jacob Stegenga (2013) argues that Arrow's theorem limits the possibilities for combining different kinds of evidence.

6.3 Overall Similarity

Things are more similar to each other in one respect, less similar in another. Much philosophy relies on notions of aggregate or "overall" similarity and Arrow's framework has been used to study these. According to Karl Popper (1963), scientific theories, though false, can get closer and closer to the truth. Work on his notion of *verisimilitude* has distinguished "likeness" and "content" dimensions, and the question arises whether these can be combined into a single ordering of theories by their overall verisimilitude. Zwart and Franssen (2007) argue that Arrow's theorem does not apply to this problem but, using a theorem inspired by it, they argue that there is no good way to combine the different dimensions even so. See Schurz and Weingartner (2010) and Oddie (2013) for constructive criticism of their views.

Overall similarity lies at the foundation of David Lewis's metaphysics (Lewis 1968, 1973a, 1973b). He wrote little about how similarities and differences in various respects might go together to yield overall similarities, though (Lewis 1979) gives some idea of what he had in mind. The Arrow-Sen framework lends itself to studying this aggregation problem as well; and an impossibility theorem, if it applies, limits the possibilities for arriving at overall similarities of the sort that Lewis presupposes. Morreau (2010) presents that case that a variant of Arrow's theorem does apply. Kroedel and Huber (2013) take a more optimistic view of the notion of overall similarity.

Bibliography

- Arrow, K.J., 1950, "A Difficulty in the Concept of Social Welfare," *Journal of Political Economy*, 58 (4): 328-346.
- Arrow, K.J., 1951, *Social Choice and Individual Values*, New York: Wiley. 2nd ed., 1963.
- Arrow, K.J. and H. Raynaud, 1986, *Social Choice and Multicriterion Decision-Making*, Cambridge MA: MIT Press.
- Arrow, K.J., A.K. Sen, and K. Suzumura (eds.), 1996, *Social Choice Re-Examined,* vol. 2, London and Basingstoke: Macmillan.
- Arrow, K.J., A.K. Sen, and K. Suzumura (eds.), 2002, *Handbook of Social Choice and Welfare,* vol. 1, Amsterdam: North-Holland.
- Arrow, K.J., A.K. Sen, and K. Suzumura (eds.), 2010, *Handbook of Social Choice and Welfare,* vol. 2, Amsterdam: North-Holland.
- Barberà, S., 2010, "Strategy-Proof Social Choice," in Arrow, Sen and

Suzumura 2010: 731-831.
- Bergson, A., 1938, “A Reformulation of Certain Aspects of Welfare Economics,” *The Quarterly Journal of Economics*, 52 (2): 310-334.
- Black, D., 1948, “On the Rationale of Group Decision-making,” *Journal of Political Economy*, 56 (1): 23-34.
- Blackorby, C., D. Donaldson and W. Bossert, 2006, “Anonymous Single-Profile Welfarism,” *Social Choice and Welfare*, 27: 279-287.
- Borda, J.C. de, 1781, “Mémoire sur les Élections au Scrutin,” *Histoire de l’Académie Royale de Sciences,* Paris. In translation: Alfred de Grazia 1953, “Mathematical Derivation of an Election System,” *Isis,* 44 (1-2): 42-51.
- Bordes, G. and N. Tideman, 1991, “Independence of Irrelevant Alternatives in the Theory of Voting,” *Theory and Decision*, 30 (2): 163-186.
- Buchanan, J., 1954, “Social Choice, Democracy, and Free Markets,” *Journal of Political Economy* 62: 114-123.
- Campbell, D.E. and J.S. Kelly, 2002, “Impossibility Theorems in the Arrovian Framework,” in Arrow, Sen and Suzumura (eds.) 2002: 35–94.
- Condorcet, Marquis de, 1785, *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*, Paris.
- Dodgson, C.L. (Lewis Carroll), 1884, *The Principles of Parliamentary Representation*, London: Harrison and Sons.
- Elster J. and Aanund Hylland, 1986, *Foundations of Social Choice Theory,* Cambridge: University Press.
- Estlund D., 1994, “Opinion Leaders, Independence, and Condorcet’s Jury Theorem,” *Theory and Decision* 36: 131-162.
- Estlund D., 2008, *Democratic Authority: a Philosophical Framework*, Princeton: University Press.
- Fleurbaey, M. and P. Mongin, 2005, “The News of the Death of Welfare Economics is Greatly Exaggerated,” *Social Choice and Welfare* 25: 381-418.
- Fleurbaey, M., 2007, “Social Choice and Just Institutions: new perspectives,” *Economics and Philosophy* 23 (1): 15-43.
- Franssen, M., 2005, “Arrow’s Theorem, Multi-criteria Decision Problems and Multi-Attribute Preferences in Engineering Design,” *Research in Engineering Design*, 16, 42-56.
- Gaertner, W., 2001, *Domain Conditions in Social Choice Theory*, Cambridge: University Press.
- Gaertner, W., 2009, *A Primer in Social Choice Theory*, Oxford. Rev. ed.
- Geanakoplos, J., 2005, “Three Brief Proofs of Arrow’s Impossibility Theorem,” *Economic Theory*, 26, 211-215.
- Gibbard, A., 1969, “Social Choice and the Arrow Conditions,” Unpublished Manuscript. The main result is available in (Gibbard 2014).
- Gibbard, A., 1973, “Manipulation of Voting Schemes: A General Result,” *Econometrica,* 41, 587-601.
- Gibbard, A., 2014, “Intransitive Social Indifference and the Arrow Dilemma,” *Review of Economic Design*, 18: 3-10.
- Hansson, B., 1973, “The Independence Condition in the Theory of Social Choice,” *Theory and Decision* 4, 25-49.

- Hurley, S., 1985, "Supervenience and the Possibility of Coherence," *Mind*, 94, 501-525.
- Hurley, S., 1989, *Natural Reasons*," Oxford: University Press.
- Hylland, Aa., 1986, "The Purpose and Significance of Social Choice Theory: some general remarks and an application to the 'Lady Chatterley problem'," in Elster and Hylland (eds.) 1986:45-73.
- Kelly, J.S., 1978, *Arrow Impossibility Theorems*, New York: Academic Press.
- Kemp, M.C. and Y.K. Ng, 1976, "On the Existence of Social Welfare Functions, Social Orderings and Social Decision Functions," *Economica* 43: 59-66.
- Kolm, Serge-Christophe 1996, "Rational Just Social Choice," in Arrow, Sen and Suzumura 1996: 167-196.
- Kornhauser, L. and L Sager 1986, "Unpacking the Court," *The Yale Law Journal,* 96: 82-117.
- Kroedel, T. and F. Huber, 2013, "Counterfactual Dependence and Arrow," *Noûs* 47: 453-466.
- Landemore, H. 2012, *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*, Princeton: University Press.
- Le Breton, M. and J.A. Weymark, 2006, "Arrovian Social Choice Theory on Economic Domains," in Arrow, Sen and Suzumura (eds.) 2011: 191-299.
- Lewis, D. 1968, "Counterpart Theory and Quantified Modal Logic," *Journal of Philosophy*, 65: 113–126.
- Lewis, D. 1973a, *Counterfactuals*, Oxford: Blackwell Publishers and Cambridge: Harvard University Press, 1973, Reprinted with revisions, 1986.
- Lewis, D. 1973b, "Causation," *Journal of Philosophy*, 70: 556–567
- Lewis, D. 1979, "Counterfactual Dependence and Time's Arrow," *Noûs*, 13: 455–476.
- List, C., 2012, "The Theory of Judgment Aggregation: an Introductory Review," *Synthese*, 187: 179-207.
- List, C., R. C. Luskin, J. S. Fishkin, and I. McLean, 2013, "Deliberation, Single-Peakedness, and the Possibility of Meaningful Democracy: Evidence from Deliberative Polls." *Journal of Politics,* 75: 80–95.
- Little, I.M.D., 1952, "Social Choice and Individual Values," *Journal of Political Economy*, 60 (5): 422-432.
- Mackie, G., 2003, *Democracy Defended*, Cambridge: University Press.
- May, K., 1954, "Intransitivity, Utility, and the Aggregation of Preference Patterns," *Econometrica*, 22: 1-13.
- McLean, I., 2003, "The Reasonableness of Independence: A Conversation from Condorcet and Borda to Arrow and Saari," *Nuffield College Working Paper 2003-W6*: http://www.nuffield.ox.ac.uk/politics/papers/2003/McLean%20The%20Reasonableness.pdf.
- Miller, D., 1992, "Deliberative Democracy and Social Choice," *Political Studies*, 40, 54-67.
- Morreau, M. 2010, "It Simply Does Not Add Up: Trouble with Overall Similarity," *The Journal of Philosophy*, 107: 469-490.
- Morreau, M., 2013, "Mr. Fit, Mr. Simplicity and Mr. Scope: From Social Choice

to Theory Choice," *Erkenntnis*: online version, forthcoming in print.

- Morreau, M., forthcoming, "Theory Choice and Social Choice: Kuhn Vindicated," forthcoming in *Mind*.
- Moulin, H., 1980, "On Strategy-Proofness and Single Peakedness," *Public Choice*, 35: 437-455.
- Murakami Y., 1968, *Logic and Social Choice*, London: Routledge and Kegan Paul.
- Oddie, G., 2013, "The Content, Consequence and Likeness Approaches to Verisimilitude: Compatibility, Trivialization, and Underdetermination," *Synthese* 190: 1647-1687.
- Okasha, S., 2011, "Theory Choice and Social Choice: Kuhn versus Arrow," *Mind* 120: 83-115.
- Parks, R.P., 1976, "An Impossibility Theorem for Fixed Preferences: a Dictatorial Bergson-Samuelson Welfare Function," *Review of Economic Studies* 43: 447-450.
- Pattanaik. P.K., 2002, "Positional Rules of Collective Decision Making," in Arrow, Sen and Suzumura 2002: 361-394.
- Pazner, E., 1979, "Equity, Nonfeasible Alternatives and Social Choice: a reconsideration of the concept of social welfare." In J.J. Laffont (ed.), *Aggregation and Revelation of Preferences,* Amsterdam: North Holland.
- Peter, F. 2011, *Democratic Legitimacy,* New York, London: Routledge.
- Plott, C. R. 1973, "Path Independence, Rationality, and Social Choice," *Econometrica* 41: 1075-1091.
- Popper, K. R., 1963, *Conjectures and Refutations*, London: Routledge.
- Ray, P., 1973, "Independence of Irrelevant Alternatives," *Econometrica*, 44: 987-991.
- Riker, W.H., 1982, *Liberalism against Populism*, San Francisco: W.H. Freeman.
- Roberts, K.W.S., 1980, "Social Choice Theory: The Single- and Multi-Profile Approaches, *The Review of Economic Studies*, 47: 441-450.
- Samuelson, P., 1947, *Foundations of Economic Analysis*, Cambridge, Massachusetts: Harvard University Press.
- Samuelson, P., 1967, "Arrow's Mathematical Politics," in S. Hook (ed.) *Human Values and Economic Policy,* New York: University Press: 41-52.
- Samuelson, P., 1977, "Reaffirming the Existence of 'Reasonable' Bergson-Samuelson Social Welfare Functions," *Economica* 44 (173): 81-88.
- Sattherthwaite, M., 1975, "Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions," *Journal of Economic Theory*, 10: 187-217.
- Schurz, G. and P. Weingartner, 2010, "Zwart and Franssen's Impossibility Theorem holds for Possible-World-Accounts but not for Consequence-Accounts to Versimilitude," *Synthese* 172: 415-436.
- Scott M.J. and E.K. Antonsson, 2000, "Arrow's Theorem and Engineering Design Decision Making," *Research in Engineering Design*, 11: 218-228.
- Sen, A., 1969, "Quasi-transitivity, Rational Choice and Collective Decisions," *Review of Economic Studies* 36: 381-393.

- Sen, A., 1970, *Collective Choice and Social Welfare*, San Francisco: Holden-Day.
- Sen, A., 1976, "Liberty, Unanimity and Rights," *Economica*, 43: 217-245.
- Sen, A., 1977, "On Weights and Measures: Informational Constraints in Social Welfare Analysis," *Econometrica*, 45 (7): 1539-1572.
- Sen, A., 1979, "Utiliarianism and Welfarism," *The Journal of Philosophy*, 76: 463-489.
- Sen, A., 1986, "Foundations of Social Choice Theory: An Epilogue," in Elster and Hylland 1986: 213-248.
- Sen, A., 1999, "The Possibility of Social Choice," *American Economic Review*, 89: 349-378.
- Sen, A., 2010, "The Informational Basis of Social Choice," in Arrow, Sen and Suzumura (eds.) 2010: 29-46.
- Stegenga, J. 2013, "An Impossibility Theorem for Amalgamating Evidence," *Synthese* 190: 2391-2411.
- Weymark, J. A., 1984, "Arrow's Theorem with Social Quasi-Orderings," *Public Choice* 42: 235-246.
- Zwart, S. D. and M. Franssen 2007, "An Impossibility Theorem for Verisimilitude," An Impossibility Theorem for Verisimilitude," *Synthese* 158: 75-92.

Academic Tools

[Auto-inserted by SEP staff]

Other Internet Resources

- [Create link to URL 1 here]
- [Create link to URL 2 here]
- [Create link to URL 3 here]

Related Entries

Philosophy of Economics | Social Choice Theory | Voting Methods | Economics and Economic Justice | Preferences

<Michael Morreau>
<michael.morreau@uit.no>