Social Choice and the Arrow Conditions

Allan Gibbard

Harvard University

Autumn Semester, 1968–69

For the Joint Seminar of Kenneth Arrow, John Rawls, and Amartya Sen

Scanned 2010 January 22[1]

Allan Gibbard

Department of Philosophy

University of Michigan, Ann Arbor

gibbard@umich.edu

+http://www-personal.umich.edu/~gibbard/

[1] This paper has been widely cited but never published. I now scan a photocopy of the original to post on my website.

A. Gibbard

SOCIAL CHOICE AND THE ARROW CONDITIONS

The Arrow impossibility theorem is somewhat unintuitive. The following dialog will probably not make it much more clear what is going on, but it may show something. I present a proof in a two person case of the lemma which seems to me to be at the heart of the proof of the theorem: that xDy implies xDz.

SOCRATES. Good morning, Meletus. That is a beautiful ice cream cone you have. Are you planning to eat it?

MELETUS. Indeed, Socrates.

SOCRATES. Yet are you not a man of justice, Meletus?

MELETUS. The most just man in ancient Greece, Socrates!

SOCRATES. Then if I can show you that justice demands that you let me eat it, you will certainly let me.

MELETUS. If you could show such a thing, I would give you the cone. But since it is mine and I want to eat it myself, I don't see how you can possibly show that justice demands that I let you eat it.

SOCRATES. I admit the task will not be easy. But tell me, Meletus, do you accept that justice is determined by a social choice function which satisfies the conditions of collective rationality, unanimity, and independence of irrelevant alternatives?

MELETUS. Of course, Socrates.

SOCRATES. Now Meletus, suppose for some reason you could not eat the ice cream cone, and I wanted to eat it. Do you agree that the social choice should be to let me eat it, regardless of your preference in the matter?

MELETUS. Well I don't know, Socrates. It's my cone, after all.

SOCRATES. Certainly it is your cone, and possibly we shall decide that that gives you the right to eat it if you can and want to. But if you can't eat it, our choice function shouldn't allow you to throw it away out of pure spite when I want it.

MELETUS: All right, I accept what you say, Socrates.

SOCRATES. Then if I prefer eating it to throwing it away, the social preference will rank my eating it above throwing it away, regardless of your preferences.

MELETUS. That seems a small thing to concede, especially since in fact I can eat it, and I want to, and there is no reason to suppose that a just social choice function would deny me my first choice in the matter.

SOCRATES. We shall see. But now, Meletus, suppose I preferred to eat the cone, but preferred throwing it away to letting you eat it. And suppose you were not hungry and preferred to throw the cone away, but would rather eat it than let me eat it.

MELETUS. My preferences would never be so malicious, and I am surprised that you admit that yours might be!

SOCRATES. I was not accusing you or myself of malice. But malicious preferences are not impossible, and our social choice function must be able to handle them. Suppose, then, we had the malicious preferences I have described. Then I would prefer eating the cone to throwing it away, and we have agreed that my preferences rule between those two alternatives, whatever your preferences may be. Furthermore, given the preferences I have pictured, we would both prefer throwing the cone away to having you eat it. Thus the social choice function would rank throwing it away above your eating it in the case I have pictured. Thus given these supposed individual preferences, justice would rank my eating the cone above throwing it away and throwing it away above your eating it. By transitivity, it would rank my eating it above your eating it.

MELETUS. Your logic is impeccable, Socrates. But fortunately my preferences are not those in the supposition.

SOCRATES. I have not forgotten that. But see how things would be if our preferences were as I have pictured them. I would prefer eating the cone to letting you eat it, you would prefer eating it to letting me eat it, and justice would demand that I eat it.

MELETUS. Yes, Socrates.

SOCRATES. But the situation as it stands now is exactly as I have just described. I prefer eating it to letting you eat it and you prefer eating it to letting me eat it. In deciding whether you or I shall eat it, how we feel about throwing it away is irrelevant. If justice demanded that I eat the cone given the preferences I pictured earlier, it demands that I eat the cone given our actual preferences, for our actual preferences differ from those hypothetical preferences only in the way we rank throwing the cone away, which is not in question here.

MELETUS. The ice cream has melted, Socrates.

In short, as soon as Meletus concedes that Socrates has a substantive right—that between throwing the cone away and having Socrates eat it, Socrates' wishes should prevail—he has conceded Socrates dictatorial power. I think we might be convinced that something was wrong here even if Professor Arrow had not shown that his conditions on social welfare functions logically preclude their existence. I want in this paper to discuss what we should do to the Arrow conditions to avoid results like these, and incidentally to avoid the impossibility theorem itself.

One way to avoid such results is, of course, to drop the condition of indepencence of irrelevant alternatives. Ways of doing this have been widely discussed. If the only information we have is individual preferences among available alternatives, though, then the social choice must be independent of what individual preferences among other alternatives would be. We ought still to be able to find a satisfactory social choice function. The question I shall discuss is this: If we accept the argument for the independence of irrelevant alternatives, how can we most reasonably modify the other conditions to get an acceptable method for social choice?

In the spirit of Professor Sen's work on "social decision functions", let me define a social choice function in terms of choice sets. This gives more generality than a definition in terms of a binary relation, without sacrificing the purpose of a social choice function. To avoid problems of making choices from sets of alternatives which are not topologically closed, I shall confine my discussion to cases where the set of alternatives is finite. Given a finite set of alternatives, an n-place <u>social choice function</u> over V is a function f which for each n-tuple $R_1,\ldots,R_n$ of weak orderings of V and each subset S of V, gives a set $C \subseteq S$ of optimal alternatives.

$$C = f(S, R_1, \ldots, R_n) \ \& \ \forall S, R_1, \ldots, R_n : C \subseteq S$$

I am going to discuss whether various combinations of conditions are reasonable for such a function. Obviously whether a condition on the function is reasonable depends on what the function and its arguments are supposed to represent. I shall assume that the arguments $R_1,\ldots,R_n$ are non-ethical preference orderings of members of a society. Given these and the set S of available alternatives, C is the set of alternatives which are, let us say, just for the society. In an ideal society, f would represent the unanimous belief of members of the society as to what makes an alternative just.

Now given this view of what a social choice function is, some of the Arrow conditions seem undoubtedly right to me. I accept that the domain of f should contain all n-tuples of orderings of V, that social choices should be Pareto-optimal, and that no one should be dictator. Even if some of these conditions are not ethically valid for all social choices, we can surely find at least three alternatives to which they jointly apply, and this is all we need to prove the impossibility theorem.

It is the two remaining conditions, the independence of irrelevant alternatives and collective rationality, that I want to discuss. The independence condition says that the set B of optimal alternatives is a function of the preferences over available alternatives alone.

$$\exists g\ \forall S, R_1, \ldots, R_n : \ C = g(S, R_1 \uparrow S, \ldots, R_n \uparrow S).$$

The social choice is thus independent of what individual preferences would be between any pair of alternatives such that at least one member is unavailable.

Let me break collective rationality into three conditions. the condition I shall call <u>fixity of social preference</u> states that the social choice is generated by a binary relation R whose ranking of any pair of alternatives does not depend on the availability or unavailability of other alternatives.

$$\exists h\ \forall R_1, \ldots, R_n\ \exists R : [R = h(R_1, \ldots, R_n)\ \&\ R \text{ is connected } \&\ \forall S : C = \{x \mid x \in S\ \&\ \forall y [y \in S \rightarrow xRy]\}].$$

The function h which this condition asserts to exist is essentially the Arrow social welfare function or Sen social decision function. For each n-tuple of individual orderings $R_1,\ldots,R_n$, it gives a binary relation R. Thus R is independent of S. For each S, though, R generates the choice set $f(S,R_1,\ldots,R_n)$.

It would be easy to confuse the independence and fixity conditions. In the first place, they both make claims of independence. The "independence" condition says that $f(S,R_1,\ldots,R_n)$ is independent of $R_1-(R_1\uparrow S),\ldots,R_n-(R_n\uparrow S)$, the parts of the individual orderings which involve non-available alternatives. The fixity condition says that $f(S,R_1,\ldots,R_n)$ is generated on S by a relation R which is independent of S. Thus while they both make claims of independence, the claims are different. In the second place, Professor Arrow's Condition 3, the independence of irrelevant alternatives, really involves both independence and fixity. His social welfare function takes an ordering R as its value, and the choice set C is defined as the set generated by R from S, the set of available alternatives. Once C has been defined this way, his Condition 3 says the same thing about it as my independence condition. I hope I can show that separating the two conditions clarifies the question of social choice.

As an example of a social choice function which satisfies independence but not fixity, consider a system of rank order voting among available candidates. Each voter gives two votes to his first choice among the available candidates, and one vote to his second choice. Now how members feel about unavailable candidates has no bearing on the choice among available candidates. Thus the method satisfies the independence condition. Using an example similar to Professor Arrow's, though, we can show that the method violates fixity. Let two voters order x, y, and z alphabetically, while three voters order them zxy. If all three are available, x gets 7 votes, y gets 2, and z gets 6.

x is the winner. If y is unavailable, however, then x still gets 7 votes, but z gets 8, and wins. Thus the social choice is not generated by a fixed relation. For if it were, since z beats x in a two man contest, we would have zPx. Then x could not be the winner in a three man contest which includes z.

The two other conditions I want to discuss are strengthenings of the fixity condition. That is, they are conditions on the relations R asserted to exist by the fixity condition. We can require quasi-transitivity of these relations: that xPy and yPz implies xPz. A social choice function which satisfies quasi-transitivity is generated by a Sen social decision function. Finally, we can require these relations R to be fully transitive. A social choice function which satisfies the condition of full transitivity is generated by an Arrow social welfare function. While full transitivity implies quasi-transitivity and quasi-transitivity implies fixity, the converses, Professor Sen has shown, do not hold.

We know that unrestricted domain, Pareto optimality, non-dictatorship, independence, and full transitivity are incompatible. Accepting the first three conditions, I want to assume independence and discuss quasi-transitivity and fixity. I shall show that quasi-transitivity has unacceptable consequences, and that the consequences of the weaker condition fixity are still bad. Then I shall argue that if we accept the argument which makes the independence condition plausible, we have no good reason to expect fixity to hold. Thus, as we would expect, with a fuller understanding of what is involved, the Arrow paradox ceases to be paradoxical.

I. Quasi-transitivity: The Liberum Veto Oligarchy

I want to show in this section that even though there are social choice functions which satisfy quasi-transitivity and the other Arrow conditions, they are not ethically satisfactory.

In the first place, even with transitivity reduced to quasi-transitivity, giving someone as little as one substantive right still makes him dictator. Professor Arrow's proof that xDy implies uDv requires only quasi-transitivity. Socrates, for example, uses quasi-transitivity but not full transitivity to talk Meletus out of his ice cream cone. The problem is not confined to small scale choice situations, though. Suppose we say that a person has a right to medical care. This presumably means at a minimum that if situation x differed from situation y only in that in x individual k gets medical care and in y he doesn't, then if k prefers x to y, society should prefer x to y even if everyone other than k prefers y to x. In other words, k is decisive for x against y. Given the Arrow conditions with transitivity reduced to quasi-transitivity, though, this makes k dictator.

Worse results yet follow from quasi-transitivity. The reader might be willing to accept that a satisfactory social choice function cannot embody rights. To show beyond question that relaxing transitivity to quasi-transitivity does not permit a satisfactory social choice function, I show in the mathematical appendix that it creates what we might call a liberum veto oligarchy. There is a unique set A of individuals, the theorem says, such that for all x and y,

(i) $[\forall i(i \in A \longrightarrow xP_i y)] \longrightarrow xPy$

(ii) $[\exists i(i \in A \longrightarrow xP_i y)] \longrightarrow xRy$

If the members of the oligarchy A are unanimous between two alternatives, what they say goes, regardless of what anybody outside the oligarchy prefers. This consequence alone is acceptable, for the oligarchy might include a majority, or even everyone. In addition, though, each member of the oligarchy can veto any social preference. For any pair of alternatives x and y, if a member of the oligarchy prefers x to y, society cannot prefer y to x. ~~Thus if the oligarchy is large,~~ The choice function says nothing unless everyone in the oligarchy can agree. Thus if the oligarchy is large, social preference requires extreme non-ethical consensus. Since I have not been able to find a simple proof of this theorem, I place the proof I have found in an appendix.

Together, then, quasi-transitivity and independence are too restrictive. This does not tell us which of the two principles is the culprit. Let me propose an argument which lays the blame squarely on quasi-transitivity. The argument does not convince me in the end, but if it fails, it at least shows something about the kinds of rights we can allow an individual if we accept quasi-transitivity. I purport to show, using a generally recognized right, that quasi-transitivity and unanimity alone lead to unacceptable results, even if we drop the independence condition.

Edwin proposes to Angelina, and says that if she refuses to marry him, he will marry Beatrice, who wants whatever will make him happy. Angelina argues that he is obligated to remain single. "You have admitted that I have a right not to marry you if I don't want to. So if the choice is between your marrying me and remaining single, and I prefer that you remain single, the social welfare function ranks your bachelorhood above our marriage. Now I do prefer your remaining single to marrying me, but I would rather marry you than have you marry Beatrice. Therefore since society ranks your bachelorhood above your marrying me and marrying me above marrying Beatrice, by quasi-transitivity it must rank bachelorhood above marrying Beatrice. I refuse your proposal, and you have no right to marry Beatrice."

In this case, of course, if Edwin is unrestrained by considerations of justice, Angelina will marry him. The threat that he will marry Beatrice if Angelina refuses him will be enough to make Angelina accept. If he accepts that she has a right not to marry him in the sense she claims, and accepts unanimity and quasi-transitivity, he must remain single. The conclusion that he must remain single seems unacceptable.

I do not regard this as a telling argument against quasi-transitivity, for I may not have correctly described the right not to marry. When we say that Angelina has a right not to marry Edwin, we may mean roughly that if she ranks marrying him last, then society must prefer his remaining single to his marrying her. Now to assert a right in this sense is to violate either the independence condition or quasi-transitivity. For as I shall show in the next section, if a social choice function satisfies independence and fixity, then whether xRy depends solely on individual preferences between x and y. Thus if we make the social choice between Edwin marrying Angelina or remaining single depend on whether she ranks marrying him ahead of or behind his marrying Beatrice, we violate either independence or fixity. If we retain quasi-transitivity we retain fixity. Thus if we want to retain quasi-transitivity, and for that reason define rights as I have suggested in this paragraph, we must drop the independence condition. If we are willing to drop the independence condition, though, then nothing in Angelina's argument forces us to drop quasi-transitivity.

II. The Fixity Condition

I have shown that if we retain the independence condition, quasi-transitivity is unacceptable. What happens when we weaken the condition to fixity?

If a social choice function satisfies both the independence and fixity conditions, then whether xPy, xIy, or yPx depends solely on the individual preferences between x and y. For by independence, the social choice set C for the set $\{x,y\}$ of available alternatives depends only on the individual preferences between x and y, and by fixity, xPy, xIy, or yPx according as $f(\{x,y\},R_1,\ldots,R_n)$ is $\{x\}$, $\{x,y\}$, or $\{y\}$ respectively.

Now I must admit that given independence and fixity, I cannot derive as much of a disaster as I could with independence and quasi-transitivity. Still, the consequences of combining independence and fixity are bad. If we accord someone as little as one substantive right, he becomes much more dictatorial than he ought to be. If k is decisive for x against y, then no alternative can be ranked higher than x and none lower than y unless k consents.

If we say, for example, that other things being equal, k has a right to medical care even if the others would prefer him to remain sick, we say that he has a right to medical care at any cost. For let x, y and z be the following situations:

x: k gets medical care, and everybody else is healthy.
y: The same as x, except that k does not get medical care.
z: k gets medical care, but because of that, half the people in the country die a hideous death. The people who die do not include k or any of his friends.

Since we are assuming both independence and fixity, to say that k has a right to medical care can only be to say that society must prefer x to y if k does, regardless of the preferences of the others. In other words, xDy. Now suppose, very hypothetically, that k's preferences were zxy and everybody else's were yzx. Of course we hope that nobody would have such preferences, for to prefer z to x is either misanthropic or suicidal. The social choice function, however, is supposed to tell us what to do for all possible combinations of preferences, and what it says for a strange preference configuration is related to what it says for the individual preferences we can expect. Now given the preferences I have hypothesized, we know xPy, because xP_ky, and other things being equal, k is entitled to medical care. Also zPx, because by supposition everybody prefers z to x. Therefore we cannot have yPz. If yPz, we would have a cycle xPy, yPz, zPx, and there would be no choice set—there would he no alternative such that nothing else in the set is preferred to it. Since we cannot have yPz, we have zRy. In the hypothetical case, then, k prefers z to y, everybody else prefers y to z, and society finds z no worse than y.

This conclusion, to be sure, supposes a strange configuration of preferences. Suppose, though, that only alternatives y and z are available, and k, as we would expect, orders z above y in his non-ethical preference scheme while everybody else orders y above z. By fixity, the unavailability of one alternative does not alter the social preference among the others, as long as all the individual orderings remain unchanged. By the independence condition, what the individual preferences would have been with respect to x is irrelevant if x is not available. Thus since the individual preferences between y and z are unchanged, the social preference is unchanged. If k prefers z to y and everybody else prefers y to z, then zRy. Society cannot prefer to leave k sick for an extra week to avert the hideous death of half its population; at best it can be indifferent.

I have used this example to argue that if we accept independence, we should reject fixity. Perhaps, though, it shows rather that if we accept independence, we should reject the unrestricted domain condition. To derive my unacceptable conclusions, I asked what would happen in the example if individuals had preferences which it is hard to imagine their having. If the social choice function does not give an answer for such cases, the argument does not go through.

I think I can still pin the blame on the combination of independence and fixity. Even if we reject the unrestricted domain condition, and require simply that the social choice function include in its domain any configuration of preferences which we can imagine individuals having, independence and fixity still lead to strange results. If there is a right to medical care, other things being equal, then when medical care is not available, there is no right to freedom of religion. Let k be a preacher who falls ill, and consider these alternatives.

x: No one goes to k's church, and k gets medical care.

y: No one goes to k's church, and k does not get medical care.

z: Everyone goes to k's church, and k does not get medical care.

We can imagine k with a preference ordering zxy. He wants medical care, but it is more important to have everybody in church. We can imagine preference ordering yzx for everyone else. They do not like the preacher, and would prefer to see him sick. They like church, but because they dislike the preacher, they would slightly prefer not to go. This preference, though, is weaker than their preference for seeing the preacher sick.

Now as in the previous example, zPx by unanimous preference, and xPy by k's right to medical care. Thus to avoid a cyclic social preference, we must have zRy. Thus again, if only y and z are available, society cannot prefer y to z. If no medical care is available and everybody is required to attend church, those who would prefer not to have no grounds for complaint no matter how strong their preferences may be. The situation in which they are in his church against their wills is no worse than the one in which they are by choice not in his church.

We might exclude all malicious preferences from the domain of the social choice function, and thus get rid even of the argument in the last example. If we do exclude all malicious preferences, though, then the social choice function will not be able to handle situations which actually arise. For even moral men can have malicious non-ethical preferences. Since they are moral, of course, their malicious non-ethical preferences are overridden by correct ethical beliefs, but that still does not prevent their non-ethical preferences from being malicious. The social choice, then, should allow some malicious preferences in its domain.

When I say this, I do not mean that the social choice function should give malicious individual preferences weight. Probably if k has a malicious preference for x over y, the social choice should be what it would be if he were indifferent between x and y. To ignore malicious preferences in arriving at a social choice, though, is different from giving no answer at all when faced with malicious preferences. To exclude malicious preferences from the domain of the social choice function is to do the latter, and I do not think we should.

In short, then, if we accept the independence condition, we should reject the fixity condition, for together they give ethically unacceptable results. Even if we drop the condition of unrestricted domain, and admit only individual preferences we can imagine people having, the consequences of indepencence and fixity are still unacceptable.

III. The Information Argument

Fixity and independence are ethically incompatible. If we consider fixity desirable, this means we should reject independence. I shall not discuss whether indeed we should, for the issue has already been widely discussed in the literature. Rather, I shall argue that if Professor Arrow's argument for independence was plausible in the first place, then nothing I have said gives us any reason to reject the independence condition. For if we accept that argument, we have no reason to expect a social choice function to satisfy the fixity condition.

Why should we accept the independence condition? Professor Arrow answers that the condition may be needed to make the criterion for social choice empirical. It does no good to say, "Choose the policies which satisfy condition F" if there is no way, even in principle, of telling which policies satisfy condition F. Now there may be no way of telling, even in principle, what people's preferences are among alternatives not open to society. If that is so, then any meaningful recommendation must give the policies to be chosen as a function of individual preferences among available alternatives alone. I shall call this the "information argument".

I want to show that if we accept the information argument, there is no reason to expect a satisfactory social choice function to satisfy the fixity condition. The social choice function is presumably a formalization of one of our ordinary ethical concepts. The conditions a social choice function should satisfy are the conditions we think that ordinary ethical concept satisfies. Now what ordinary ethical concept does the social choice function formalize? The most obvious candidate is simply the concept 'best', in a sense in which in case of ties we can apply the word to all winners. If we accept the information argument, though, we cannot claim that the choice set C contains the best of the available alternatives in S. For which members of S are best would, given our ordinary ethical notions, depend not simply on who preferred what to what in S, but on the strengths of the preferences. The information argument, though, tells us that strength of preference is not an empirical matter. Therefore there can be no empirical criterion for x actually being a best alternative in S. We might, in fact, on the basis of this argument, want to dismiss the phrase 'actually best' as meaningless.

What the social choice function should give us, if the information argument is correct, is an empirical criterion for the ordinary ethical notion 'best as far as one can tell from the available information', or 'apparently best'. What can we conclude from this? If, on the contrary, the social choice function were supposed to tell us which choices were actually best, the fixity condition, and indeed transitivity, would make sense. Given any set of alternatives, the choice function would simply pick out the best. If, however, we accept the information argument, and regard the choice function as picking out only the apparently best alternatives, then there is no reason to accept even the fixity condition. Which of two alternatives is apparently better may depend on what information is available. According to the fixity argument, the information available depends on what alternatives are possible. Therefore whether xRy holds—whether x is apparently as good as y—may vary with the set of alternatives available, even if all individual preferences remain fixed.

If the way individuals order pairs of alternatives other than $\langle x,y\rangle$ had no bearing on whether x is apparently better than y, then we would still have to accept fixity. However, the other rankings clearly do tell us something about whether x is better than y. First, of course, they give us some hint of preference intensity. In addition, the other rankings may tell us that an individual's preference for x over y is a result of malice, and therefore not to be taken into account. If Meletus, for example, preferred throwing the cone away to eating it himself, then the fact that he prefers eating it himself to letting Socrates eat it would deserve less weight than it does if eating it himself is his first choice. Similarly, if Angelina prefers marrying Edwin to letting him marry Beatrice, then society should not necessarily rank his remaining single above his marrying Angelina, as it would of right if marrying Edwin were her last choice.

In short, then, if the social choice function tells for any set of alternatives which ones are actually best, then for fixed individual preferences we can expect the choices to be generated by a fixed, transitive relation 'actually better than'. If we accept the information argument, though, we can expect the social choice function to tell us only which of the available alternatives are apparently best. This does not tell us, though, which alternatives would be apparently best if more or fewer alternatives were open. For then we would have different information, and x might not be apparently better than y, even though given the information we now have, x is apparently better than y.

If there is anything paradoxical left in the Arrow paradox, it is this: Either the information argument is wrong, or there is no meaningful empirical criterion for whether any particular alternative open to society is actually best. Once we accept that a social choice function which satisfies the independence condition can tell us only what is apparently best, the rest of the paradox disappears. For then we have no reason to accept transitivity, quasi-transitivity, or fixity.

APPENDIX: SOCIAL CHOICE AND QUASI-TRANSITIVITY

DEFINITION: An n-place social choice function over a finite set V of alternatives is a function f which for each n-tuple $R_1, R_2, \ldots, R_n$ of orderings of V and each subset S of V, gives a set $C \subseteq S$ of optimal alternatives.

$$C = f(S, R_1, \ldots, R_n), \quad C \subseteq S.$$

The condition of independence of irrelevant alternatives says that the set C of optimal alternatives is a function of the preferences over available alternatives alone.

Independence: $\exists g \forall S, R_1, \ldots R_n : C = g(S, R_1 \uparrow S, R_2 \uparrow S, \ldots, R_n \uparrow S)$.

The condition of fixity of social preference says that the social choice is generated by a binary relation which ranks each pair of available alternatives independently of the availability or non-availability of any other alternative.

Fixity: $\exists h \forall R_1, \ldots, R_n \; \exists R : [R \text{ is connected } \& \; R = h(R_1, \ldots, R_n) \; \&$
$\forall S : C = \{x \mid x \in S \; \& \; \forall y [y \in S \rightarrow xRy]\}]$.

Quasi-transitivity: Fixity holds, and the R's asserted to exist in the condition are quasi-transitive:

$$\forall x, y, z : xPy \; \& \; yPz \longrightarrow xPz, \text{ where } xPy \text{ means } xRy \sim yRx.$$

Transitivity: Fixity holds, and the R's asserted to exist in the condition are transitive.

THEOREM: Let f be a social choice function which satisfies the conditions of unrestricted domain, Pareto optimality, independence, and quasi-transitivity. Then there is a unique set A of individuals such that for all x and y,

(i) $[\forall i (i \in A \rightarrow xP_i y)] \longrightarrow xPy$

(ii) $[\exists i (i \in A \; \& \; xP_i y)] \longrightarrow xRy$

Definitions: $xD_A y$ iff whenever everyone in A prefers x to y and everyone else prefers y to x, xPy.

$x\bar{D}_A y$ iff whenever everyone in A prefers x to y, xPy, regardless of the preferences of others.

Lemma I. If for some x and y, $xD_A y$, then for all x and y, $x\bar{D}_A y$.

The proof is simply the Arrow proof of the corresponding lemma for D and $\bar{D}$. The proof uses quasi-transitivity but not full transitivity.

Lemma II: If $\exists w,x,y,z\,[w D_A x \,\&\, y D_B z]$, then $\forall x,y: x \overline{D_{A\cap B}}\, y$.

Proof: If $w D_A x$, then by Lemma I, $x \bar{D}_A y$. Consider the following preference pattern.

For individuals in set	Ordering
$A \cap B$	xyz
$A - B$	zxy
$B - A$	yzx
$V - (A \cup B)$	zyx

Then everyone in A prefers x to y, and since $x \bar{D}_A y$, xPy. Everyone in B prefers y to z, and everyone else prefers z to y. Since $y D_B z$, yPz. Therefore by quasi-transitivity, xPz. Now everyone in $A\cap B$ prefers x to z and everyone else prefers z to x. Thus $x D_{A\cap B} z$, and by Lemma I, $\forall x,y: x \overline{D_{A\cap B}}\, y$. This proves Lemma II.

With these two lemmas we prove the theorem. Take two alternatives x and y in V, and let A be the minimal set for which $x D_A y$. We show that A and only A satisfies (i) and (ii).

Part I: A satisfies (i) by Lemma I, for since $x D_A y$, for all x and y, $x \bar{D}_A y$.

Part II: A satisfies (ii). The proof is Arrow-like. Suppose on the contrary that for some $R_1, R_2, \ldots, R_n$ and $k \in A$, $x P_k y$ but yPx. Partition V into three sets as follows.

$$B_1 = \{ i \mid x P_i y \} \quad (\therefore k \in B_1)$$
$$B_2 = \{ i \mid x I_i y \}$$
$$B_3 = \{ i \mid y P_i x \}$$

Consider the following orderings.

B_1: xzy
B_2: $x I y P z$
B_3: yxz

The orderings of x and y are unchanged, and thus by the independence condition, yPx. Everyone prefers x to z, and thus by the Pareto condition, xPz. Therefore by quasi-transitivity, yPz. Thus $y D_{B_2 \cup B_3} z$.

Now since $x D_A y$, by Lemma II, $x D_{A \cap (B_2 \cup B_3)} y$. But $A \cap (B_2 \cup B_3)$ is a proper subset of **A**, since it does not contain k. Since **A** was defined as the minimal set such that $x D_A y$, this is a contradiction, and our supposition that (ii) is violated by **A** is proved false.

Part III: **A** is unique. For let **B** satisfy (i) and (ii). Then from (i), $x D_B y$, and by Lemma II, $x \overline{D_{A \cap B}} y$. Since A is minimal, $A \cap B$ must equal A, and $A \subseteq B$. If $A \neq B$, then let $k \in B$, $k \notin A$, and suppose $x P_k y$ and $\forall i [i \neq k \longrightarrow y P_i x]$. Then from (ii), xRy. From (i), however, yPx. Therefore $A = B$. This proves the theorem.