Sen's Proof of Arrow's Impossibility Theorem without Tears

Johan E. Gustafsson *

Draft
September 25, 2024

ABSTRACT. Amartya Sen has found a remarkably short proof of Arrow's Impossibility Theorem. It is a strong candidate for being the simplest known proof of the theorem, making little use of technical machinery. A simple proof makes it easy to see not only that the theorem is true but also why it is true. Sen's presentation of the proof, however, is very compact and leaves several steps implicit. In this paper, I present a streamlined version of the proof in a thorough manner with all steps explained, with the aim to make it as easy as possible to follow.

Amartya Sen has found a remarkably elegant proof of Arrow's Impossibility Theorem.[1] It is a strong candidate for being the simplest known proof of the theorem, making very little use of technical machinery.[2] A simple proof makes it easy to see not only that the theorem is true but also why it is true.[3] Sen's presentation of the proof, however, is very compact and leaves several steps implicit. In the following, I will present a streamlined version the proof in a thorough manner with all steps explained, noting where each condition is applied. The aim is to make the proof as easy as possible to follow.

Arrow's Impossibility Theorem can be stated as follows:[4]

* I would be grateful for any thoughts or comments on this paper, which can be sent to me at *johan.eric.gustafsson@gmail.com*.

[1] It may very well be a proof from *The Book* — the transfinite book, imagined by Erdős, containing the best proof of every mathematical theorem. See Alexanderson 1981, p. 254 and Erdős 1984, p. 108.

[2] Sen 1995, p. 4; 2011, pp. 38–9; 2014, pp. 35–7; 2017, pp. 286–8. Another candidate is a short proof by Geanakoplos (2005, pp. 212–23).

[3] Arrow commented in Kelly 1987, p. 59:

> My impossibility theorem ought to be totally obvious when looked at the right way. Yet every proof involves a trick. Maybe not a big trick; I don't think it's a mathematically hard theorem. But somehow if you had the right way of approaching it, it should be trivial.

[4] Arrow 1963, pp. 27, 30, 96–7 and Sen 1977, pp. 58–9.

Arrow's Impossibility Theorem If there are at least three options and a finite number of individuals, then the following conditions cannot all be true:

- *Unrestricted Domain* For any logically possible specification of each individual's preference ordering of the options, there is a unique social preference ordering of the options.
- *Independence of Irrelevant Alternatives*[5] The social preference ranking of any options x and y only depends on the individuals' preference rankings of x and y.
- *Unanimity*[6] For all options x and y, whenever everyone prefers x to y, then x is socially preferred to y.
- *Non-Dictatorship* There is no individual such that, for all options x and y, whenever this individual prefers x to y, then x is socially preferred to y.

Here, an option x is *preferred* to an option y if and only if x is at least as preferred as y and y is not at least as preferred as x. A *preference ordering* is a preference ranking that is complete and transitive. A preference ranking is *complete* if and only if, for all options x and y, it holds that x is at least as preferred as y or y is at least as preferred as x. And a preference ranking is *transitive* if and only if, for all options x, y, and z, it holds that, if x is at least as preferred as y and y is at least as preferred as z, then x is at least as preferred as z.

Sen's proof of Arrow's Impossibility Theorem makes use of the notions of decisiveness and global decisiveness. Let us say that a set of individuals G is *decisive* for an option x against an option y if and only if x is socially preferred to y whenever everyone in G prefers x to y.[7] And let us say that a set of individuals G is *globally decisive* if and only if, for all distinct options x and y, it holds that G is decisive for x against y.[8]

[5] The name of the condition is misleading. The condition must not be conflated with following condition regarding the irrelevance of other alternatives:

> *Contraction Consistency* If option x is choice-worthy from a set of options O and x is in a subset of O, then x is also choice-worthy from the subset.

As Ray (1973) notes, Arrow's Impossibility Theorem does not hold if Independence of Irrelevant Alternatives is replaced by Contraction Consistency.

[6] Unanimity is also know as the Pareto principle. See Arrow 1963, p. 96.

[7] Arrow 1951, p. 52.

[8] Sen 1995, p. 4; 2014, p. 35.

The basic outline of the proof runs as follows. We begin by proving Spread of Decisiveness — a principle that says that, if a set of individuals is decisive for some option against another option, then that set is globally decisive. Then, with the help of Spread of Decisiveness, we prove Contraction of Decisive Sets — a principle that says that, if a set of more that one individual is globally decisive, then a smaller subset of that set is globally decisive as well. Next, we note that Unanimity entails that the set of all individuals is globally decisive. And then with repeated application of Contraction of Decisive Sets, we find that smaller and smaller sets of individuals are globally decisive. Eventually, we find that a set of just one individual is globally decisive set, which violates Non-Dictatorship.

1. Spread of Decisiveness

Spread of Decisiveness says that, if a set of individuals is decisive for an option against another option, then that set is globally decisive:[9]

> *Spread of Decisiveness* If x and y are distinct options and a set of individuals G is decisive for x against y, then G is globally decisive.

To prove Spread of Decisiveness, we will first prove three intermediate results.

We will start by proving Against-Side Spread of Decisiveness — which says that, if a set of individuals is decisive for option x against another option y, then that set is also decisive for x against any other option. Then we will prove, the analogous, For-Side Spread of Decisiveness — which says that, if a set of individuals is decisive for option x against another option y, then that set is also decisive for any other option against y. And, then with the help of both For-Side and Against-Side Spread of Decisiveness, we will prove Symmetry of Decisiveness — which says that, if a set of individuals is decisive for option x against option y, then that set is also decisive for y against x. Finally, using these intermediate results, we will prove Spread of Decisiveness.

1.1 AGAINST-SIDE SPREAD OF DECISIVENESS

First, we consider the case where we replace the option that a set of individuals is decisive *against* with a third option. We will prove the following

[9] Sen 1995, p. 4; 2014, p. 35.

result:

> *Against-Side Spread of Decisiveness* If x, y, and z are distinct options and a set of individuals G is decisive for x against y, then G is decisive for x against z.

Let x, y, and z be distinct options, and suppose that G is a set of individuals that is decisive for x against y. Unrestricted Domain entails that there is a unique social preference ordering for any logically possible specification of each individual's preference ordering. So we can let everyone in G prefer x to y, y to z, and x to z. And we can let all individuals outside G (if there are any such individuals) prefer y to z, with their other preference rankings left unspecified. Letting '$x \succ y$' denote that x is preferred to y, we have the following specification of preferences:

Specification One

	$\{x, y\}$	$\{y, z\}$	$\{x, z\}$
G	$x \succ y$	$y \succ z$	$x \succ z$
Others		$y \succ z$	

Since everyone in G prefers x to y and G is decisive for x against y, it follows that x is socially preferred to y. And, since everyone prefers y to z, it follows, by Unanimity, that y is socially preferred to z. Since x is socially preferred to y and y is socially preferred to z, it follows, by the transitivity of the social preference ordering, that x is socially preferred to z.[10]

Now, note that Independence of Irrelevant Alternatives entails that the social preference ranking of x and z only depends on the individuals' preference rankings of x and z. But the only preference rankings that

[10] This follows since transitivity (the principle that, for options x, y, and z, if x is at least as preferred as y and y is at least as preferred as z, then x is at least as preferred to z) entails *transitivity of strict preference* — the principle that, for all options x, y, and z, if x is preferred to y and y is preferred to z, then x is preferred to z. To prove this, suppose that option x is preferred to option y and y is preferred to option z. Consequently, x is at least as preferred as y, and y is at least as preferred as z. So, by transitivity, we then find that x is at least as preferred as z. Now, suppose — for proof by contradiction — that z is at least as preferred as x. Then, since y is at least as preferred as z and z is at least as preferred as x, it follows, by transitivity, that y is at least as preferred as x. But this contradicts that x is preferred to y. So we reject our assumption that z is at least as preferred as x. Accordingly, x is at least as preferred as z, and z is not at least as preferred as x. Hence x is preferred to z — which concludes the proof of transitivity of strict preference.

have been specified between x and z in Specification One are those of the individuals in G. As for the individuals outside G, it has only been assumed that they all prefer y to z, and this is compatible with any pattern of preference rankings these individuals outside G might have between x and z.[11] Whatever that pattern might be, the social preference ranking of x and z remains the same: x is socially preferred to z whenever all individuals in G prefer x to z. It follows that G is decisive for x against z — which concludes the proof of Against-Side Spread of Decisiveness.

1.2 FOR-SIDE SPREAD OF DECISIVENESS

Now, we consider the case where we replace the option that a set of individuals is decisive *for* with a third option. We will prove the following result:

> *For-Side Spread of Decisiveness* If x, y, and z are distinct options and a set of individuals G is decisive for x against y, then G is decisive for z against y.

The proof proceeds in much the same way as the proof of Against-Side Spread of Decisiveness.

Like before, let x, y, and z be distinct options, and suppose that G is a set of individuals that is decisive for x against y. Note, again, that Unrestricted Domain entails that there is a unique social preference ordering for any logically possible specification of each individual's preference ordering. So we can let everyone in G prefer z to x, x to y, and z to y. And we can let all individuals outside G (if there are any such individuals) prefer z to x, with their other preference rankings left unspecified:

Specification Two

	$\{x,z\}$	$\{x,y\}$	$\{y,z\}$
G	$z \succ x$	$x \succ y$	$z \succ y$
Others	$z \succ x$		

[11] The preferences that have been specified for the individuals outside G can be extended to a preference ordering where x is preferred to z given any preference between x and y. And they can be extended to a preference ordering where z is preferred to x given that y is preferred to x. Finally, they can be extended to a preference ordering where x is *equally preferred* as z (that is, x is at least as preferred as z and z is at least as preferred as x) given that y is preferred to x.

Since everyone in G prefers x to y and G is decisive for x against y, it follows that x is socially preferred to y. And, since everyone prefers z to x, it follows, by Unanimity, that z is socially preferred to x. Since z is socially preferred to x and x is socially preferred to y, it follows, by the transitivity of the social preference ordering, that z is socially preferred to y.[12]

Independence of Irrelevant Alternatives entails that the social preference ranking of y and z only depends on the individuals' preference rankings of y and z. But the only preference rankings that have been specified between y and z in Specification Two are those of the individuals in G. As for the individuals outside G, it has only been assumed that they all prefer z to x, and this is compatible with any pattern of preference rankings these individuals outside G might have between y and z.[13] Whatever that pattern might be, the social preference ranking between y and z remains the same: z is socially preferred to y whenever all individuals in G prefer z to y. It follows that G is decisive for z against y — which concludes the proof of For-Side Spread of Decisiveness.

1.3 SYMMETRY OF DECISIVENESS

Next, we will prove that decisiveness is symmetric:

> *Symmetry of Decisiveness* If a set of individuals G is decisive for x against y, then G is decisive for y against x.

Let x and y be distinct options, and suppose that G is a set of individuals that is decisive for x against y. First, by Against-Side Spread of Decisiveness, we find that G is decisive for x against z, where z is an option that is distinct from each of x and y. Then, by For-Side Spread of Decisiveness, we find that G is decisive for y against z. And then, by Against-Side Spread of Decisiveness, we find that G is decisive for y against x — which concludes the proof of Symmetry of Decisiveness.

[12] This follows by transitivity of strict preference. See note 10.

[13] The preferences that have been specified for the individuals outside G can be extended to a preference ordering where y is preferred to z given that y is preferred to x And they can be extended to a preference ordering where z is preferred to y given any preference between x and y. Finally, they can be extended to a preference ordering where y is equally preferred as z (see note 11) given that y is preferred to x.

Finally, we will prove Spread of Decisiveness.

Again, let x and y be distinct options, and suppose that G is decisive for x against y. Then, by Symmetry of Decisiveness, we find that G is decisive for y against x. Let z be any option that is distinct from each of x and y. Then, by Against-Side Spread of Decisiveness, we find that G is decisive for x against z. And, by For-Side Spread of Decisiveness, we find that G is decisive for z against y. Next, by Symmetry of Decisiveness, we find that G is decisive for z against x and decisive for y against z.

If there are only three options, then there are no further combinations to consider. And, if there are more than three options, then we also need to show that G is decisive for option u against option v, where u, v, x, and y are all distinct. We have supposed that G is decisive for x against y. Then, by Against-Side Spread of Decisiveness, we find that G is decisive for x against v. And then, by For-Side Spread of Decisiveness, we find that G is decisive for u against v.

Hence, for any distinct options x and y, it holds that G is decisive for x against y. That is, G is globally decisive — which concludes the proof of Spread of Decisiveness.

2. Contraction of Decisive Sets

Contraction of Decisive Sets says that, if a set of at least two individuals is globally decisive, then a smaller subset is so too:[14]

> *Contraction of Decisive Sets* If a set of individuals G is globally decisive and contains at least two individuals, then a smaller subset of G is globally decisive.

Having proven Spread of Decisiveness, we will now prove Contraction of Decisive Sets.

Let G be a globally decisive set of at least two individuals. Divide G into two subsets G_1 and G_2. Note, once more, that Unrestricted Domain entails that there is a unique social preference ordering for any logically possible specification of each individual's preference ordering. So, letting x, y, and z be distinct options, we can let everyone in G_1 prefer x to y and x to z, with their preference rankings of y and z left unspecified, and

[14] Sen 1995, p. 4; 2014, p. 36.

let everyone in G_2 prefer x to z and y to z, with their preference rankings of x and y left unspecified. And we can leave the preference rankings of individuals outside G (if there are any such individuals) completely unspecified:

Specification Three

	$\{x, y\}$	$\{x, z\}$	$\{y, z\}$
G_1	$x \succ y$	$x \succ z$	
G_2		$x \succ z$	$y \succ z$
Others			

Now, there are two possible cases, which we will consider in turn. It can either be that (i) the preference rankings in Specification Three are sufficient to derive that x is socially preferred to y or be that (ii) they are not sufficient to derive that x is socially preferred to y.

First, consider case (i). Independence of Irrelevant Alternatives entails that the social preference ranking of x and y only depends on the individuals' preference rankings of x and y. But the only preference rankings that have been specified between x and y in Specification Three are those of the individuals in G_1. As for the individuals outside G (that is, the individuals outside both G_1 and G_2), their preference rankings have not been specified at all, and, as for the individuals in G_2, it has only been assumed that they all prefer x to z and y to z — but this is compatible with any pattern of preference rankings which the individuals outside G_1 might have between x and y. Whatever that pattern might be, the social preference ranking between x and y remains the same: x is socially preferred to y whenever all individuals in G_1 prefer x to y. It follows that G_1 is decisive for x against y. Then, by Spread of Decisiveness, it follows that G_1 is globally decisive.

Next, consider case (ii) — that is, the case that the preference rankings in Specification Three are not sufficient to derive that x is socially preferred to y. Then there must be some extension of Specification Three such that, given that extended specification of individual preference rankings, x is not socially preferred to y. Remember that Independence of Irrelevant Alternatives entails that the social preference ranking of x and y only depends on the individuals' preference rankings of x and y. Hence the only part of the extension that is needed to let us derive that x is not socially preferred to y is the individuals' preference rankings between x and y. Accordingly, there is an extension of Specification Three where the only addition is that all individuals' preference rankings between x and y

are specified so that the extended specification lets us derive that x is not socially preferred to y. Letting '$x * y$' denote a specified preference ranking between x and y in this extension (leaving it open what the specified preference ranking is), we have the following:

Specification Three Extended

	$\{x, y\}$	$\{x, z\}$	$\{y, z\}$
G_1	$x \succ y$	$x \succ z$	
G_2	$x * y$	$x \succ z$	$y \succ z$
Others	$x * y$		

Since we extended Specification Three for this purpose, the individual preference rankings in Specification Three Extended let us derive that x is not socially preferred to y. Then, by the completeness of the social preference ordering, it follows that y is socially at least as preferred as x. Since G is globally decisive and everyone in G (that is, everyone in G_1 and G_2) prefers x to z, it follows that x is socially preferred to z. Then, since y is socially at least as preferred as x and x is socially preferred to z, it follows, by the transitivity of the social preference ordering, that y is socially preferred to z.[15]

Independence of Irrelevant Alternatives entails that the social preference ranking of y and z only depends on the individuals' preference rankings of y and z. The only preference rankings between y and z which are specified in Specification Three Extended are those of the individuals in G_2. As for the individuals outside G, only their preference rankings between x and y have been specified, and, as for the individuals in G_1, it has only been assumed that they all prefer x to y and x to z — but this is compatible with any pattern of preference rankings these individuals outside G_2 might have between y and z. Whatever that pattern might be, the social preference ranking of y and z remains the same: y is socially preferred to z whenever all individuals in G_2 prefer y to z. It follows that G_2 is decisive for y against z. Then, by Spread of Decisiveness, it follows that G_2 is globally decisive.

[15] To see this, suppose that y is at least as preferred as x and x is preferred to z. Then x is at least as preferred as z. And, since y is at least as preferred as x and x is at least as preferred as z, it follows, by transitivity, that y is at least as preferred as z. Now, assume — for proof by contradiction — that z is at least as preferred as y. Then, since z is at least as preferred as y and y is at least as preferred as x, it follows, by transitivity, that z is at least as preferred as x. But this contradicts that x is preferred to z. So we reject our assumption that z is at least as preferred as y. Accordingly, y is at least as preferred as z, and z is not at least as preferred as y. Hence y is preferred to z.

Hence, in each of cases (i) and (ii), we find that one of G_1 and G_2 is globally decisive. Since these cases are exhaustive, it follows that one of these smaller subsets of *G* is globally decisive — which concludes the proof of Contraction of Decisive Sets.

3. Proof of Arrow's Impossibility Theorem

Having proven Contraction of Decisive Sets, the proof of Arrow's Impossibility Theorem is straightforward. Let *G* be the set of all individuals. Then, by Unanimity, it follows that, for all options *x* and *y*, whenever everyone in *G* prefers *x* to *y*, then *x* is socially preferred to *y*. Accordingly, *G* is globally decisive. Then, by repeated application of Contraction of Decisive Sets, it follows that smaller and smaller subsets of *G* are also globally decisive — until we, after a finite number of steps (since the number of individuals is finite), reach a subset with just one individual which is globally decisive. But then there is an individual such that, for all options *x* and *y*, whenever this individual prefers *x* to *y*, then *x* is socially preferred to *y*. But this conclusion contradicts Non-Dictatorship. It follows that Unrestricted Domain, Independence of Irrelevant Alternatives, Unanimity, and Non-Dictatorship cannot all be true — which concludes the proof of Arrow's Impossibility Theorem.

I wish to thank Gustav Alexandrie, Harvey Lederman, Jacob Nebel, Martin Peterson, and Wlodek Rabinowicz for valuable comments.

References

Alexanderson, G. L. (1981) 'An Interview with Paul Erdős', *The Two-Year College Mathematics Journal* 12 (4): 249–259.

Arrow, Kenneth J. (1951) *Social Choice and Individual Values*, New York: Wiley.

— (1963) *Social Choice and Individual Values*, New York: Wiley, second edn.

Erdős, Paul (1984) 'Convocation Address', in *Indian Statistical Institute: Fifty-Second Annual Report April 1983–March 1984*, pp. 101–108, Calcutta: Indian Statistical Institute.

Geanakoplos, John (2005) 'Three Brief Proofs of Arrow's Impossibility Theorem', *Economic Theory* 26 (1): 211–215.

Kelly, J. S. (1987) 'An Interview with Kenneth J. Arrow', *Social Choice and Welfare* 4 (1): 43–62.

Ray, Paramesh (1973) 'Independence of Irrelevant Alternatives', *Econometrica* 41 (5): 987–991.

Sen, Amartya (1977) 'Social Choice Theory: A Re-Examination', *Econometrica* 45 (1): 53–89.

— (1995) 'Rationality and Social Choice', *The American Economic Review* 85 (1): 1–24.

— (2011) 'The Informational Basis of Social Choice', in Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, eds., *Handbook of Social Choice and Welfare Volume Two*, pp. 29–46, Amsterdam: Elsevier.

— (2014) 'Arrow and the Impossibility Theorem', in Eric Maskin and Amartya Sen, eds., *The Arrow Impossibility Theorem*, pp. 29–42, New York: Columbia University Press.

— (2017) *Collective Choice and Social Welfare: An Expanded Edition*, Cambridge, MA: Harvard University Press.