

Computational Social Choice 2021

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

[<http://www.illc.uva.nl/~ulle/teaching/comsoc/2021/>]

Plan for Today

To illustrate a further application of the *axiomatic method*, today we are going to review three of the classical *impossibility theorems* in the domain of voting and preference aggregation:

- *Arrow's Theorem* (1951)
- *Sen's Theorem* on the Impossibility of a Paretian Liberal (1970)
- the *Muller-Satterthwaite Theorem* (1977)

They all show that it is impossible to simultaneously satisfy certain intuitively appealing axioms when designing a voting rule.

Full details of all proofs are available in my review paper (cited below).

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*, College Publications, 2011.

Axiom: The Pareto Principle

A voting rule $F : \mathcal{L}(A)^n \rightarrow 2^A \setminus \{\emptyset\}$ is called (weakly) *Paretian* if, whenever all voters rank alternative x above y , then y cannot win:

$$N_{x \succ y}^{\mathbf{R}} = N \text{ implies } y \notin F(\mathbf{R})$$

Axiom: The Principle of Liberalism

Think of A as the set of all possible ‘social states’. Certain aspects of such a state will be some individual’s private business. Example:

If x and y are identical states, except that in x I paint my bedroom white, while in y I paint it pink, then I should be able to dictate the relative social ranking of x and y .

Remark: For examples of this kind, it makes more sense to think of F as a *social choice function* rather than a *voting rule*.

F is called *liberal* if, for every individual $i \in N$, there exist two distinct alternatives $x, y \in A$ such that i is *two-way decisive* on x and y :

$$i \in N_{x \succ y}^{\mathbf{R}} \text{ implies } y \notin F(\mathbf{R}) \text{ and } i \in N_{y \succ x}^{\mathbf{R}} \text{ implies } x \notin F(\mathbf{R})$$

The Impossibility of a Paretian Liberal

Bad news:

Theorem 1 (Sen, 1970) *For two or more individuals, there exists **no** social choice function that is both **Paretian** and **liberal**.*

As we shall see, the theorem holds even when liberalism is enforced for only two individuals. The number of alternatives does not matter.

A.K. Sen. The Impossibility of a Paretian Liberal. *Journal of Political Economics*, 1970.

Proof Sketch

Let F be a SCF that is Paretian and liberal. Get a contradiction:

Take two distinguished individuals i_1 and i_2 , with:

- i_1 is two-way decisive on x_1 and y_1
- i_2 is two-way decisive on x_2 and y_2

Assume x_1, y_1, x_2, y_2 are pairwise distinct (other cases: easy).

Consider a profile with these properties:

- (1) Individual i_1 ranks $x_1 \succ y_1$.
- (2) Individual i_2 ranks $x_2 \succ y_2$.
- (3) All individuals rank $y_1 \succ x_2$ and $y_2 \succ x_1$.
- (4) All individuals rank x_1, x_2, y_1, y_2 above all other alternatives.

From liberalism: (1) rules out y_1 and (2) rules out y_2 as winner.

From Pareto: (3) rules out x_1 and x_2 and (4) rules out all others.

Thus, there are no winners. Contradiction. ✓

Resolute Social Choice Functions

For the remainder of today, we focus on *resolute SCF's*:

$$F : \mathcal{L}(A)^n \rightarrow A$$

The axioms we have seen already can be easily adapted to this slightly simpler model. For example, this is the Pareto Principle:

$$N_{x \succ y}^{\mathbf{R}} = N \text{ implies } y \neq F(\mathbf{R})$$

The next result we are going to see, Arrow's Theorem, originally got formulated for so-called *social welfare functions* instead:

$$F : \mathcal{L}(A)^n \rightarrow \mathcal{L}(A)$$

This change in framework does not affect the essence of the result, and it makes it fit better with our overall storyline . . .

Axiom: Independence of Irrelevant Alternatives

If alternative x wins and y does not, then x is *socially preferred* to y .

If both x and y lose, then we cannot say.

Whether x is socially preferred to y should *depend* only on the relative rankings of x and y in the profile (not on other, irrelevant, alternatives).

These considerations motivate our next axiom:

F is called *independent* if, for any two profiles $\mathbf{R}, \mathbf{R}' \in \mathcal{L}(A)^n$ and any two distinct alternatives $x, y \in A$, it is the case that $N_{x \succ y}^{\mathbf{R}} = N_{x \succ y}^{\mathbf{R}'}$ and $F(\mathbf{R}) = x$ imply $F(\mathbf{R}') \neq y$.

Thus, if x prevents y from winning in \mathbf{R} and the relative rankings of x and y remain the same, then x also prevents y from winning in \mathbf{R}' .

Arrow's Impossibility Theorem

A resolute SCF F is a *dictatorship* if there exists an $i \in N$ such that $F(\mathbf{R}) = \text{top}(R_i)$ for every profile \mathbf{R} . Voter i is the dictator.

The seminal result in SCT, here adapted from SWF's to SCF's:

Theorem 2 (Arrow, 1951) Any *resolute SCF* for ≥ 3 alternatives that is *Paretian* and *independent* must be a *dictatorship*.

Remarks:

- You should be surprised by this and refuse to believe it (for now).
- Not true for $m = 2$ alternatives. (Exercise: Why?)
- Common misunderstanding: dictatorship \neq 'local dictatorship'
- Impossibility reading: independence + Pareto + nondictatoriality
- Characterisation reading: dictatorship = independence + Pareto

K.J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 2nd edition, 1963. First edition published in 1951.

Proof Plan

For full details, consult my review paper, which includes proofs both for SWF's and SCF's (the latter within the proof for the *M-S Thm*).

Let F be a SCF for ≥ 3 alternatives that is Paretian and independent.

Call a *coalition* $C \subseteq N$ *decisive* for (x, y) if $C \subseteq N_{x \succ y}^{\mathbf{R}} \Rightarrow y \neq F(\mathbf{R})$.

We proceed as follows:

- *Pareto* condition = N is decisive for all pairs of alternatives
- C with $|C| \geq 2$ *decisive* for all pairs \Rightarrow some $C' \subset C$ as well
- By induction: there's a decisive coalition of size 1 (= *dictator*).

Remark: Observe that this only works for finite sets of voters. (*Why?*)

The step in the middle of the list is known as the *Contraction Lemma*.

To prove it, we first require another lemma ...

U. Endriss. Logic and Social Choice Theory. In A. Gupta and J. van Benthem (eds.), *Logic and Philosophy Today*, College Publications, 2011.

Contagion Lemma

Recall: $C \subseteq N$ *decisive* for (x, y) if $C \subseteq N_{x \succ y}^{\mathbf{R}} \Rightarrow y \neq F(\mathbf{R})$

Call $C \subseteq N$ *weakly decisive* for (x, y) if $C = N_{x \succ y}^{\mathbf{R}} \Rightarrow y \neq F(\mathbf{R})$.

Claim: C weakly decisive for $(x, y) \Rightarrow C$ decisive for *all* pairs (x', y') .

Proof: Suppose x, y, x', y' are all distinct (other cases: similar).

Consider a profile where individuals express these preferences:

- Members of C : $x' \succ x \succ y \succ y'$
- Others: $x' \succ x$, $y \succ y'$, and $y \succ x$ (note: x' -vs.- y' not specified)
- All rank x, y, x', y' above all other alternatives.

From C being weakly decisive for (x, y) : y must lose.

From Pareto: x must lose (to x') and y' must lose (to y).

Thus, x' must win (and y' must lose). By independence, y' will still lose when everyone changes their non- x' -vs.- y' rankings.

Thus, for every profile \mathbf{R} with $C \subseteq N_{x' \succ y'}^{\mathbf{R}}$ we get $y' \neq F(\mathbf{R})$. \checkmark

Contraction Lemma

Claim: If $C \subseteq N$ with $|C| \geq 2$ is a coalition that is decisive on all pairs of alternatives, then so is some nonempty coalition $C' \subset C$.

Proof: Take any nonempty C_1, C_2 with $C = C_1 \cup C_2$ and $C_1 \cap C_2 = \emptyset$.

Recall that there are ≥ 3 alternatives. Consider this profile:

- Members of C_1 : $x \succ y \succ z \succ \text{rest}$
- Members of C_2 : $y \succ z \succ x \succ \text{rest}$
- Others: $z \succ x \succ y \succ \text{rest}$

As $C = C_1 \cup C_2$ is decisive, z cannot win (it loses to y). Two cases:

- (1) The winner is x : Exactly C_1 ranks $x \succ z \Rightarrow$ By independence, in any profile where exactly C_1 ranks $x \succ z$, z will lose (to x) $\Rightarrow C_1$ is weakly decisive on (x, z) . So by Contagion Lemma: C_1 is decisive on all pairs.
- (2) The winner is y , i.e., x loses (to y). Exactly C_2 ranks $y \succ x \Rightarrow \dots \Rightarrow C_2$ is decisive on all pairs.

Hence, one of C_1 and C_2 will always be decisive. \checkmark

Axioms: Weak and Strong Monotonicity

Two axioms for a resolute SCF F :

- F is called *weakly monotonic* if $x^* = F(\mathbf{R})$ implies $x^* = F(\mathbf{R}')$ for any alternative x^* and any two profiles \mathbf{R} and \mathbf{R}' with $N_{x^* \succ y}^{\mathbf{R}} \subseteq N_{x^* \succ y}^{\mathbf{R}'}$ and $N_{y \succ z}^{\mathbf{R}} = N_{y \succ z}^{\mathbf{R}'}$ for all $y, z \in A \setminus \{x^*\}$.
- F is called *strongly monotonic* if $x^* = F(\mathbf{R})$ implies $x^* = F(\mathbf{R}')$ for any alternative x^* and any two profiles \mathbf{R} and \mathbf{R}' with $N_{x^* \succ y}^{\mathbf{R}} \subseteq N_{x^* \succ y}^{\mathbf{R}'}$ for all $y \in A \setminus \{x^*\}$.

A good way to remember the difference:

- *weak monotonicity* = raising the winner preserves the winner
- *strong monotonicity* = lowering a loser preserves the winner

Strong monotonicity is also known as *Maskin monotonicity*.

Example

Even *weak monotonicity* is not satisfied by some common voting rules.

Under *plurality with runoff* the two alternatives with the highest plurality score enter a second round and the majority winner of that round is the winner (used to elect the French president). Example:

15 voters: $a \succ b \succ c$

24 voters: $c \succ a \succ b$

14 voters: $b \succ c \succ a$

So b is eliminated in the first round and c beats a 38:15 in the runoff. But if 2 of the voters in the first group *raise c to the top*, then b wins.

But many other rules (e.g., *plurality*) do satisfy weak monotonicity.

How about *strong monotonicity*?

The Muller-Satterthwaite Theorem

More bad news:

Theorem 3 (Muller and Satterthwaite, 1977) Any *resolute* SCF for ≥ 3 alt. that is *surjective* and *strongly monotonic* is a *dictatorship*.

Here, a resolute SCF F is called *surjective* (or *nonimposed*) if for every alternative $x \in A$ there exists a profile \mathbf{R} such that $F(\mathbf{R}) = x$.

Exercise: Show that surjectivity is required for this theorem to hold.

Proof: Next, we are going to show:

- strong monotonicity implies independence
- surjectivity and strong monotonicity imply the Pareto Principle

The claim then follows from Arrow's Theorem. ✓

E. Muller and M.A. Satterthwaite. The Equivalence of Strong Positive Association and Strategy-Proofness. *Journal of Economic Theory*, 1977.

Deriving Independence

Recall: F is *independent* if, for $x \neq y$, we have that $N_{x \succ y}^{\mathbf{R}} = N_{x \succ y}^{\mathbf{R}'}$ and $F(\mathbf{R}) = x$ together imply $F(\mathbf{R}') \neq y$.

Claim: If F is strongly monotonic, then F is also independent.

Proof: Suppose F is SM, $x \neq y$, $N_{x \succ y}^{\mathbf{R}} = N_{x \succ y}^{\mathbf{R}'}$, and $F(\mathbf{R}) = x$.

Construct a third profile \mathbf{R}'' :

- All individuals rank x and y in the top-two positions.
- The relative rankings of x vs. y are as in \mathbf{R} , i.e., $N_{x \succ y}^{\mathbf{R}''} = N_{x \succ y}^{\mathbf{R}}$.
- Rest: whatever

By strong monotonicity, $F(\mathbf{R}) = x$ implies $F(\mathbf{R}'') = x$.

By strong monotonicity, $F(\mathbf{R}') = y$ would imply $F(\mathbf{R}'') = y$.

Thus, we must have $F(\mathbf{R}') \neq y$. ✓

Deriving the Pareto Principle

Recall: F is *Paretian* if $N_{x \succ y}^{\mathbf{R}} = N$ implies $F(\mathbf{R}) \neq y$.

Claim: If F is surjective and SM, then F is also Paretian.

Proof: Suppose F is surjective and SM (and thus also independent).

Take any two alternatives x and y .

From surjectivity: x will win for *some* profile \mathbf{R} .

Starting in \mathbf{R} , have everyone move x above y (if not above already).

From strong monotonicity: x still wins.

From independence: y does not win for *any* profile where all individuals continue to rank $x \succ y$. ✓

The Bigger Picture

As a deeper analysis reveals, Arrowian impossibilities arise as the result of the interaction between two forces:

- *Axioms*, particularly independence, that directly constrain the behaviour of the aggregation rule.
- *Collective rationality*, i.e., the requirement for the output to satisfy certain structural requirements (here: having a single winner).

This perspective is useful for COMSOC and AI, as it helps understand the dynamics of aggregating other types of structures, such as social networks, argument graphs, or nonstandard (incomplete) preferences.

U. Endriss and U. Grandi. Graph Aggregation. *Artificial Intelligence*, 2017.

Summary

Making heavy use of the *axiomatic method*, we have presented and proved three of the classic *impossibility theorems* of SCT. They all establish the incompatibility of certain desirable axioms:

- *Sen*: Pareto and liberalism
- *Arrow*: Pareto and independence
- *Muller-Satterthwaite*: surjectivity and strong monotonicity

In one case, the combination in question is completely impossible, in the other two it leads to a dictatorship for resolute voting rules.

What next? More axiomatic method, to analyse strategic behaviour.