Social Choice Theory

An analysis of collective decision making

Michael Osborne

November 29th, 2016

West Virginia University

Introduction

Election

Zola	$a \succ b \succ c \succ d \succ e$
Marcela	$b \succ c \succ d \succ e \succ a$
Piotr	$c \succ d \succ e \succ a \succ b$
Geoffrey	$d \succ e \succ a \succ b \succ c$
Shifu	$e \succ a \succ b \succ c \succ d$

Voting Paradox

Zola	$a \succ b \succ c \succ d \succ e$
Marcela	$b \succ c \succ d \succ e \succ a$
Piotr	$c \succ d \succ e \succ a \succ b$
Geoffrey	$d \succ e \succ a \succ b \succ c$
Shifu	$e \succ a \succ b \succ c \succ d$

As a group, we have that $a \succ b$, $b \succ c$, $c \succ d$, $d \succ e$, and $e \succ a$. So, our individual preference orderings have produced a cyclic ordering when conceived as a group.

What is social choice theory?

Utilitarian Viewpoint

	1	2	3	4	5	Total	Score
Hillary Clinton	21.14% 26	15.45% 19	11.38% 14	8.13% 10	43.90% 54	123	2.62
Gary Johnson	5.69% 7	38.21% 47	27.64% 34	23.58% 29	4.88% 6	123	3.16
Bernie Sanders	25.20% 31	28.46% 35	21.95% 27	22.76% 28	1.63% 2	123	3.53
Jill Stein	0.00% 0	11.38% 14	34.15% 42	39.02% 48	15.45% 19	123	2.41
Donald Trump	47.97% 59	6.50% 8	4.88% 6	6.50% 8	34.15% 42	123	3.28

The social choice problem

Social choice problem

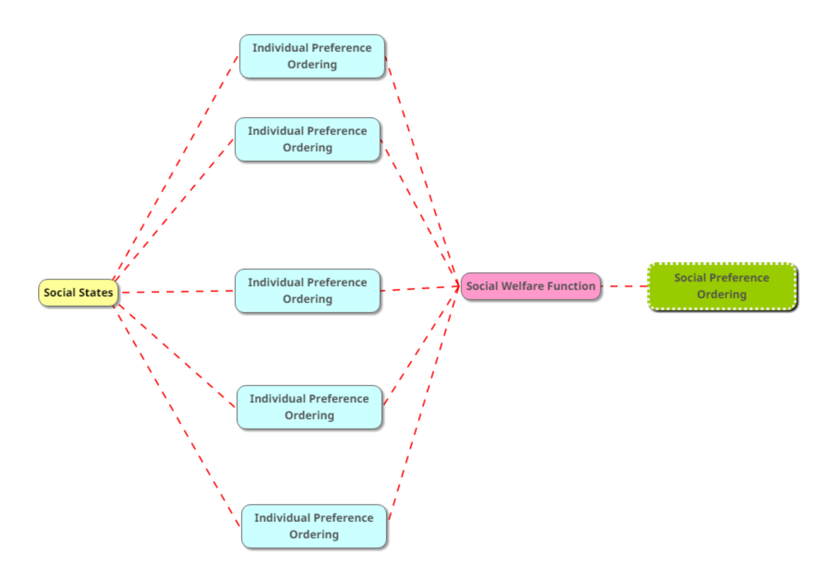

Social Welfare Function(SWF) Example 1

	1	2	3	4	5	Total	Score
Hillary Clinton	21.14% 26	15.45% 19	11.38% 14	8.13% 10	43.90% 54	123	2.62
Gary Johnson	5.69% 7	38.21% 47	27.64% 34	23.58% 29	4.88% 6	123	3.16
Bernie Sanders	25.20% 31	28.46% 35	21.95% 27	22.76% 28	1.63% 2	123	3.53
Jill Stein	0.00% 0	11.38% 14	34.15% 42	39.02% 48	15.45% 19	123	2.41
Donald Trump	47.97% 59	6.50% 8	4.88% 6	6.50% 8	34.15% 42	123	3.28

Social Welfare Function(SWF) Example 2

49% of the electorate:	Bush $\succ$ Gore $\succ$ Nader
20% of the electorate:	Gore $\succ$ Bush $\succ$ Nader
20% of the electorate:	Gore $\succ$ Nader $\succ$ Bush
11% of the electorate:	Nader $\succ$ Gore $\succ$ Bush

Social Welfare Function(SWF) Example 2

49% of the electorate:	Bush $\succ$ Gore $\succ$ Nader
20% of the electorate:	Gore $\succ$ Bush $\succ$ Nader
20% of the electorate:	Gore $\succ$ Nader $\succ$ Bush
11% of the electorate:	Nader $\succ$ Gore $\succ$ Bush

There is no majority winner. Nader receives the fewest first place votes, so we move his name to the last position of every individual preference ordering. Doing so allows Gore to receive the 11% that Nader once had. So, now Gore has 51% of the vote and is the winner of the IRV election!

Formally defining a social choice problem

Individual preference ordering can be conceived of as a vector (i.e. a list of ordered objects), I = [a, b, c, d, e, f , g, h, . . .] , ordered by preference.

Instead of using this somewhat awkward notation, for the vector J=[a,b,c] we will write $I : a \succ b \succ c$ and so on.

Then, all individual preference orderings in society can be represented by a set of vectors G = I, K, L, ...

Formally defining a social choice problem

$$M : a \succ b \succ c \succ d \succ e \succ f$$
$$N : b \succ a \succ c \succ e \succ d \succ f$$
$$O : a \sim b \succ d \succ c \succ e \succ f$$
$$P : a \succ b \succ c \succ d \sim e \succ f$$

Using majority rule as our SWF, we find the social preference ordering $S : a \succ b \succ c \succ d \succ e \succ f$. This is exactly our individual preference ordering for M.

Formally defining a social choice problem

$$M : a \succ b \succ c \succ d \succ e \succ f$$
$$N : b \succ a \succ c \succ e \succ d \succ f$$
$$O : a \sim b \succ d \succ c \succ e \succ f$$
$$P : a \succ b \succ c \succ d \sim e \succ f$$

If we changed O to $O : a \sim b \succ d \succ c \succ e \succ f$ we would no longer have it that S is the same as M. This means that the SWF is **non-dictatorial**. Should S always be the same as M, it would be **dictatorial**.

Social choice theorists claim that every normatively reasonable SWF should be non-dictatorial.

Decisive

A group of people D (which may be a single-member group), which is part of the group of all individuals G, is decisive with respect to the ordered pair of social states (a, b) if and only if state a is socially preferred to b whenever everyone in D prefers a to b. A group that is decisive with respect to all pairs of social states is simply decisive.

Non-dictatorship

No single individual (i.e. no single-member group D) of the group G is decisive.

Ordering

For every possible combination of individual preference orderings, the social preference ordering must be complete, asymmetric and transitive.

The SWF majority rule is ruled out by this condition, as some social preference orderings generated by the majority rule are cyclic.

Arrow’s impossibility theorem

Kenneth Arrow

In his 1951 doctoral thesis, Arrow proved a great theorem which later earned him the Nobel Prize.

Arrow’s impossibility theorem

There is no social welfare function (SWF) that meets the conditions of non-dictatorship and ordering, as well as two additional conditions.

Pareto condition

Arrow explicitly required that if everyone in the group prefers a to b, then the group should prefer a to b. This is known as the Pareto condition and can also be stated as ”the group of all individuals in society is decisive.”

Independence of irrelevant alternatives

Zola	$a \succ b \succ c$
Marcela	$b \succ a \succ c$

Independence of irrelevant alternatives

Old Zola	$a \succ b \succ c$
Old Marcela	$b \succ a \succ c$

New Zola	$c \succ a \succ b$
New Marcela	$a \succ c \succ b$

Since Zola and Marcela still agree that a is better than b, the New society must also prefer a to b. How c is ranked is irrelevant when it comes to determining the social preference between a and b.

Independence of irrelevant alternatives

If all individuals have the same preference between a and b in two different sets of individual preference orderings G and G', then society's preference between a and b must be the same in G and G'.

Independence of irrelevant alternatives: Controversy

IIA effectively excludes all SWFs that are sensitive to relational properties of the individual preference orderings.

Independence of irrelevant alternatives: Controversy Example

Consider the following group of individual preference orderings:

$$U : c \succ b \succ a$$
$$V : b \succ a \succ c$$
$$W : a \succ c \succ b$$

Independence of irrelevant alternatives: Controversy Example

Consider the following group of individual preference orderings:

$$U : c \succ b \succ a$$
$$V : b \succ a \succ c$$
$$W : a \succ c \succ b$$

$$S : a \sim b \sim c$$

Independence of irrelevant alternatives: Controversy Example

Now, consider the group of individual preference orderings with c's in a different position:

$$U' : c \succ b \succ a$$
$$V' : c \succ b \succ a$$
$$W' : c \succ a \succ b$$

Independence of irrelevant alternatives: Controversy Example

Now, consider the group of individual preference orderings with c's in a different position:

$$U' : c \succ b \succ a$$
$$V' : c \succ b \succ a$$
$$W' : c \succ a \succ b$$

$$S' : c \succ b \succ c$$

Proof of Arrow's Impossibility Theorem

Field-Expansion Lemma If a group D is decisive with respect to any pair of states, then it is decisive.

Group-Contraction Lemma If a group D (which is not a single-person group) is decisive, then so is some smaller group contained in it.

Proof of Arrow's Impossibility Theorem Pareto tells us that the group of all individuals is decisive. Since the number of individuals in society was assumed to be finite we can apply it over and over again (to a smaller decisive group). At some point, we will eventually end up with a decisive single-member group, that is, a dictator.

Sen on liberalism and Pareto principle

Amartya Sen

In 1998, Amartya Sen was awarded the Nobel Prize in Economics for his lifetime work in social choice and welfare economics.

Liberalism and the Pareto Principle

Sen argued that the Pareto principle is incompatible with the basic ideals of liberalism. If correct, this indicates that there is either something wrong with liberalism, or the Pareto principle, or both.

Example of liberalism

Consider an individual who prefers to have pink walls rather than white walls at home. In a liberal society, we should permit this somewhat unusual preference, even if the majority would prefer to see white walls.

More precisely put, Sen proposes a minimal condition of liberalism, according to which there is at least one pair of alternatives (a, b) for each individual such that if the individual prefers a to b, then society should prefer a to b, no matter what others prefer.

Minimal liberalism

Minimal liberalism: There are at least two individuals in society such that for each of them there is at least one pair of alternatives with respect to which she is decisive, that is, there is a pair a and b, such that if she prefers a to b, then society prefers a to b (and society prefers b to a if she prefers b to a).

The paradox of the Paretian liberal

Theorem: No SWF satisfies minimal liberalism, Pareto and the ordering condition.

Robert Nozick believes that liberalism shouldn't be a property of a SWF. What do you think?

The paradox of the Paretian liberal

Theorem: No SWF satisfies minimal liberalism, Pareto and the ordering condition.

Robert Nozick believes that liberalism shouldn't be a property of a SWF. What do you think?

If the decision is going to be made based on one individuals preferences, why even pose it as a decision to be made by society?

Proof: The paradox of the Paretian liberal

Let the two individuals referred to in the condition of minimal liberalism be X and Y, and let the two decisive pairs of alternatives be (a, b) and (c, d), respectively. Obviously, (a, b) and (c, d) cannot be the same pair of alternatives, because if X's preference is a $a \succ b$ and Y's is $b \succ a$, then the social preference ordering would be $a \succ b \succ a$, which contradicts the ordering condition.

Proof: The paradox of the Paretian liberal

This leaves us with two possible cases: (a, b) and (c, d) either have one element in common, or none at all.

Proof: The paradox of the Paretian liberal

Let us first consider the case in which they have one element in common, say $a = c$. Suppose that X's preference is $a \succ b$ and that Y's is $d \succ c (= a)$. Also suppose that everyone in society, including X and Y, agrees that $b \succ d$. Because of the ordering condition, X's preference ordering is $a \succ b \succ d$, while Y's is $b \succ d \succ a$. The ordering condition guarantees that this set of individual preference orderings is included in the domain of every SWF. However, minimal liberalism entails that society prefers $a \succ b$ and $d \succ c$, and since we assumed that $c = a$, it follows that society prefers $d \succ a$. Finally, Pareto implies that society prefers $b \succ d$. Hence, the social preference ordering is $a \succ b \succ d \succ a$, which contradicts the ordering condition.

Proof: The paradox of the Paretian liberal

To complete the proof, we also have to consider the case in which the pairs (a, b) and (c, d) have no common elements. Let X's preference ordering include $a \succ b$, and let Y's include $c \succ d$, and let everyone in society (including X and Y) prefer $d \succ a$ and $b \succ c$. Hence, X's preference ordering must be as follows: $d \succ a \succ b \succ c$, while Y's is $b \succ c \succ d \succ a$. However, minimal liberalism entails that society prefers $a \succ b$ and $c \succ d$, whereas Pareto entails that $d \succ a$ and $b \succ c$. It follows that the social preference ordering is $d \succ a \succ b \succ c \succ d$, which contradicts the ordering condition.

Harsanyi’s utilitarian theorems

John Harsanyi

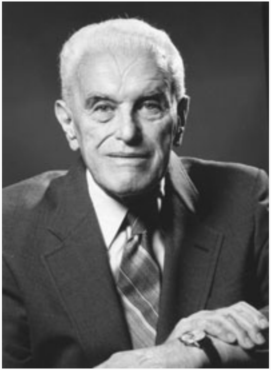

In 1994, John Harsanyi was awarded the Nobel Prize in Economics along with John Nash for his lifetime work in utilitarian ethics and equilibrium selection.

John Harsanyi's Utilitarian Solution

Harsanyi defended a utilitarian solution to the problem of social choice, according to which the social preference ordering should be entirely determined by the sum total of individual utility levels in society.

Examples of the Utilitarian Solution

Suppose a single individual strongly prefers a high tax rate over a low tax rate, and all others disagree. Then, by the utilitarian solution, society should nevertheless prefer a high tax rate given that the preference of the single individual is sufficiently strong.

Suppose a doctor can save five dying patients by killing a healthy person and transplanting her organs to the five dying ones – without thereby causing any negative side-effects (such as decreased confidence in the healthcare system) – then the doctor should kill the healthy person.

Individual Rationality

Harsanyi rejected Arrow's view that individual preference orderings carry nothing but ordinal information. He said that it is reasonable to assume that individual preference orderings satisfy the von Neumann and Morgenstern axioms for preferences over lotteries (or some equivalent set of axioms). This directly implies that rational individuals can represent their utility of a social state on an interval scale.

The Chairperson

Harsanyi asks us to imagine an individual (who may or may not be a fellow citizen) who evaluates all social states from a moral point of view. Let us refer to this individual as the Chairperson. If the Chairperson is a fellow citizen, then he has two separate preference orderings:

- a personal preference ordering over all states that reflects his personal preference ordering
- a separate preference ordering over the same set of social states that reflects the social preference ordering.

The Chairperson

What can be concluded about the Chairperson's social preference ordering, given that it fulfills certain structural conditions?

All that matters is that we somehow know that the Chairperson's preferences, whatever they are, conform to the structural conditions proposed by von Neumann and Morgenstern.

Rationality of social preferences and Pareto

Rationality of social preferences: The Chairperson's social preference ordering satisfies the von Neumann and Morgenstern axioms for preferences over lotteries.

Pareto: Suppose that a is preferred to b in at least one individual preference ordering, and that there is no individual preference ordering in which b is preferred to a. Then, a is preferred to b in the Chairperson's social preference ordering. Furthermore, if all individuals are indifferent, then so is the Chairperson in his social preference ordering.

Utility functions

From individual rationality it follows that individual preference orderings can be represented by utility functions that measure utility on an interval scale, and from rationality of social preferences it follows that the same holds true of the social preference ordering.

Harsanyi's first theorem

Let $u_i(a)$ denote individual i's utility of state a, and let $u_s(a)$ denote the utility of a as reflected in the Chairperson's social preference ordering. Furthermore, let α be a real number between 0 and 1.

Individual rationality, rationality of social preferences and Pareto together entail that:

$$u_s(a) = \sum_{i=1}^{n} \alpha_i u_i(a) \text{ with } \alpha_i > 0 \text{ for } i = 1, \cdots, n.$$

This theorem tells us that society's utility of state a is a weighted sum of all individuals' utility of that state.

Harsanyi's first theorem

The theorem guarantees that each individual preference ordering is assigned some weight. However, utilitarians typically argue that all individual preference orderings should be assigned the same weight. Harsanyi thinks he can solve this problem by introducing a further assumption.

Equal treatment of all individuals: If all individuals' utility functions $u_1, \cdots, u_n$ are expressed in equal utility units (as judged by the Chairperson, based on interpersonal utility comparisons), then the Chairperson's social utility function u_c must assign the same weight to all individual utility functions.

Harsanyi's second theorem

Harsanyi's second theorem: Given equal treatment of all individuals, the coefficients in Harsanyi's first theorem will be equal:

$$\alpha_1 = \cdots = \alpha_n$$

Questions?