8 The Problem of Social Choice

One can argue that social choice theory asks the most fundamental question in economics. That, is, consider any economy. Then, provided that we are economists and insists on full rationality assumptions we have no conceptual difficulty in thinking about how an individual agents evaluates all possible *alternatives* (think of allocations, outcomes) in that economy. While the set of possible allocations in some more realistic model of an actual economy would be amazingly large, at least we understand the rationalistic benchmark, and, studying particular problems, we are often quite confident that we can ignore all but what we view as central to the particular problem when analyzing, say the individual preferences for clean air.

Now, the Pareto criterion is a quite useful (but incomplete) concept for thinking about "what is good for society as a whole". Arguably, one shortcoming is that it doesn't include fairness. However, this isn't the only issue. It is also the case that the Pareto criterion is problematic as a way to reach a decision. One problem is that starting with a given (non-Pareto optimal) allocation x, there is generally a whole range of allocations that Pareto dominate x. So, as a society, which one should we pick?

This leads to the following question; is there any reasonable procedure that can be used to "aggregate preferences"? That is, can we find a rule, where if we feed in the preferences of all individuals, the rule delivers a preference ordering for society? Now, clearly, one way of doing it is to pick an agent in society and just use that agents preferences. That doesn't seem fair, so we don't quite like that. Another way, would be to say that society has some given preferences regardless of what people actually think. That doesn't seem good either, becasuse we'd miss out on possible Pareto improvments. Hence, what the theory of social choice considers is a question which vaguely can be stated as "is it possible to find a rule that aggregates individual preferences that satisfies some minimal criteria for what should be considered a good rule?".

9 Social Welfare Functions

Let:

- A be a set of social alternatives
- $\succeq$ be a *preference ordering* or *weak ordering* on A. This means that $\succeq$ describes some "sensible ranking" of the alternatives in A. By "sensible" we mean that $\succeq$ satisfies a couple of axioms which together imply that $\succeq$ agrees with our common sense usage of the word ranking. These axioms are

  Axiom 1 *For all $x, y \in A$ either $x \succeq y$ or $y \succeq x$*

  This is referred to as "completeness" and simply means that all alternatives can be compared. Note that (since we are considering *weak* orderings) *both* $x \succeq y$ and $y \succeq x$ may be true.

  Axiom 2 *For all $x, y, z \in A$. If $x \succeq y$ and $y \succeq z$, then $x \succeq z$.*

  This is called transitivity. As you may recall, these two axioms together generates what we think of as rational preferences. A strict preference and an indifference relation can be derived if desired.

Now, consider a society with n agents and write $\succeq_i$ for the preference ordering of agent i. We are now in the position to formally describe a social welfare function. From a collection $(\succeq_1, .., \succeq_n)$ of individual ordering we want to derive a *social ordering* $\succeq$. Very importantly and sometimes misunderstood: the problem of social choice is to find a process *assigning a social ordering for any collection of individual orderings in society.* Formally:

Definition 1 *A social welfare function is a process that assigns a social ordering $\succeq$ for each individual ordering $(\succeq_1, ..., \succeq_n)$, i.e., it is a function $F(\cdot)$, where each $F(\succeq_1, ..., \succeq_n)$ is a complete and transitive preference ordering.*

Now, exactly what would constitute desirable properties of this process of preference aggregation is a matter of debate. However, it seems that we would dislike a society where a single agent would determine the preferences completely. If the function F has such an agent we call this agent a dictator;

Definition 2 *$i \in N$ is a dictator for F if the social preferences agree with is preferences whenever is preferences are strict. That is, in our notation, if $x \succ y \Rightarrow xF(\succeq_1, ..., \succeq_n)\, y$ regardless of the preferences of the other agents.*

Definition 3 *F satisfies non-dictatorship if F has no dictator.*

Next, we would like to ensure that we always get Pareto efficiency;

Definition 4 *F satisfies the Pareto principle if society prefers x to y whenever all agents agree that x is better than x (that is $xF(\succeq_1, ..., \succeq_n)\, y$ whenever $x \succeq_i y$ for all agents.*

Non-dictatorship and the Pareto principle are straightforward requirements. However, the final requirement we will discuss is more subtle. It says, in essence, that the pairwise rank between any two alternatives should not be affected by any other alternative:

Definition 5 *Consider two preference profiles $(\succeq_1, ..., \succeq_n)$ and $(\succeq_1', ..., \succeq_n')$ which are such that $x \succeq_i y \Leftrightarrow x \succeq_i' y$. Then, F satisfies independence of irrelevant alternatives if $xF(\succeq_1, ..., \succeq_n)\, y \Leftrightarrow xF(\succeq_1', ..., \succeq_n')\, y$*

While Social choice is about aggregation of preferences and not (directly) about preference revelation, the rationale that I like the most for this requirement is that it is necessary (relatively easy to understand) and sufficient (too hard to explain the reason for this in an undergraduate course) for being able to induce truthful revelation of preferences in the case where these are unknown.

9.1 Some Examples

9.1.1 Majority Rule with 2 Alternatives

Consider the case when $A = \{x, y\}$. In this case there are only 3 possible weak orderings ($x \succ y$,indifference and $y \succ Px$). Majority rule then simply says that society prefers x to y if and only if a majority prefers x to y (a good exercise is to introduce notation so that F can be described formally). To check whether this rule satisfies the criteria above, let's observe that;

1. Suppose that i prefers x to y and all others prefer y to x. Then, y is preferred by a majority. Thus, for all i there exists some preferences for the others such that the societal preferences go against i. This means that majority rule satisfies non-dictatorship.

2. Suppose that all agents agree that x (or y) is the preferred alternative. Then, the societal preferences also prefers x. Hence, the Pareto principle is satisfied.

3. Finally, since the only alternatives are x and y, if two preference profiles agree on the choice between x and y, the preference profiles are identical. Hence, the independence axiom holds by "default".

9.1.2 The Condorcet Paradox

In the example where $A = \{x, y, z\}$ and the preference profile is given by

1s ranking	2s ranking	3s ranking
x	y	z
y	z	x
z	x	y

,

we have already noted that, in pairwise majority rule, x beats y that beats z that beats x...We labelled this a voting cycle in the discussion above. Now, the question is, which of the three criteria is violated when pairwise majority rule is applied?

The answer to the question is none of them! However, pairwise majority rule nevertheless fails, because it simply does not generate a social welfare function. That is, when we define a social welfare function we define it as a function that assigns societal preferences for every preference profile in society. This builds in the restriction that *preferences for society are transitive and complete,* which is what fails when you get a voting cycle.

9.1.3 Borda Rules (Rank Order Voting)

Suppose that we for each agent, associate the number 1 with the most preferred alternative, 2 with the second and so on, and then rank the social alternatives in accordance to how many points each alternative gets when we sum over all agents (Obviously we can as well assign the highest number to the most preferred alternative as in College Football Polls). For the Condurcet example above we then get

1s ranking	2s ranking	3s ranking
$x\,(1)$	$y\,(1)$	$z\,(1)$
$y\,(2)$	$z\,(2)$	$x\,(2)$
$z\,(3)$	$x\,(3)$	$y\,(3)$

So, we find that each alternative scores 6 points in total so we get that "society is indifferent" between the alternatives. Is this a "solution" to the voting paradox? No, at least not if we think that independence of irrelevant alternatives is a reasonable requirement. To see this replace 3s ranking with the strict ranking $z \succ y \succ x$. Then we get then

1s ranking	2s ranking	3s ranking		Societal Preferences
$x\,(1)$	$y\,(1)$	$z\,(1)$		$y\,(5)$
$y\,(2)$	$z\,(2)$	$y\,(2)$	$\Longrightarrow$	$z\,(6)$
$z\,(3)$	$x\,(3)$	$x\,(3)$		$x\,(7)$

so society all of a sudden strictly prefer y to z, which is a violation of independence of irrelevant alternatives.

10 Arrows Theorem

Theorem 6 *If $|A| \geq 3$ there exists no social welfare function satisfying Non-dictatorship, the Pareto principle and Independence of irrelevant alternatives.*

No step in the proof below is particularily difficult, but it may still feel a bit overwhelming due to teh number of steps.

First of all, it is sufficient to consider the case with three alternatives. This is not totally obvious, but I will take it for granted.

The first step in in proving the result is to define a "decisive group" relative to the decision between x and y :

Definition 7 *A group E is decisive for the (ordered) pair (x, y) if $x \succ y$ whenever $x \succ_i y$ for all i in the group E and $y \succ_i x$ for all other agents.*

The first thing to ask is if such a decisive group exists, and the answer is that the set of all agents must be decisive, so we are not talking about something non-existent here.

We then claim;

Claim 8 *If E is decisive for (x, y), then E is decisive for all other pairs.*

There are more elegant ways to prove this, but I will approach it by constructing preference profiles that verifies that E must be decisive in all the five other possibilities $(x \succ_i z, y \succ_i z, y \succ_i x, z \succ_i x, z \succ_i y)$.

Proof. If E is the set of all agents, this follows trivially from the Pareto principle. Suppose not and let all agents in E have preferences $\succ_E$ and assume that all other agents have preferences $\succ_N$, where

$\succ_E$	$\succ_N$
x	y
y	z
z	x

Then;

1) $x \succ y$ since E is decisive over (x, y)

2) $y \succ z$ by the Pareto principle.

3) By transitivity, it follows that $x \succ z$.

4) Since all agents in E prefers x to z and all others have the opposite preferences, E is decisive also over (x, z)

5) By independence of irrelevant alternatives it follows that $x \succ z$ also for the profile

$$\begin{array}{cc} \succ_E & \succ_N \\ y & z \\ x & y \\ z & x \end{array}.$$

6) By the Pareto principle $y \succ x$

7) By transitivity $y \succ z$.

8) Since all agents in E prefers y to z and all others have the opposite preferences, E is decisive also over (y, z)

Now, you may worry that E may not be decisive if you would reverse the rankings, so that, for example, all agents in E think that y is better than x and the opposite holds outside E. That is, if you consider

$$\begin{array}{cc} \succ_E & \succ_N \\ y & z \\ z & x \\ x & y \end{array}.$$

it is not obvious that $y \succ x$. But,

1) We already know that E is decisive over (y, z) , so $y \succ z$

2) $z \succ x$ by the Pareto principle

3) Hence, $y \succ x$ by transitivity. Thus, E is decisive over (x, y) in the sense that if Es preferences are opposite of the preferences outside E, the social decision follows E, regardless of whether the agents in E think that x or y is the better alternative.

In the same vein, consider

$\succ_E$	$\succ_N$
z	x
y	z
x	y

Here;

1) $y \succ x$ since E is decisive over (y, x)

2) $z \succ y$ by the Pareto principle

3) Transitivity implies that $z \succ x$, which means that E is decisive over (z, x) (i.e., over the pair regardless of the direction of preferences)

The final possibility is when all agents in E prefer z to y. Then consider

$\succ_E$	$\succ_N$
z	y
x	z
y	x

Again;

1) $x \succ y$ since E is decisive over (x, y)

2) $z \succ x$ by the Pareto principle

3) Transitivity implies that $z \succ y$, which means that E is decisive over (z, y). ■

In the notion of a decisive group, we so far only know that the decision follows the decisive group if the other agents have the opposite preferences. We now show that the opposite preferences of the others is not needed.

Claim 9 *Suppose E is a decisive group and $x \succ_i y$ for all i in E. Then $x \succ y$*

Proof. Consider a profile **R** where the strict preference ordering implied for all agents in E are as follows

$$\begin{array}{c} \succ_E \\ x \\ z \\ y \end{array}$$

where t is an arbitrary alternative different from x and y. For agents outside E, let z be the most preferred alternative for everybody, but the rest of the preference ordering is arbitrary for agents outside E. Then,

1) E is decisive and the preferences over the pair (x, z) is opposite in E and outside E. Hence $x \succ z$

2) The Pareto principle guarantees that $z \succ y$

3) Hence, transitivity implies that $x \succ y$

The conclusion follows since the relative rank between x and y outside E is arbitrary. ■

Claim 10 *Suppose that neither A or B is a decisive group. Then $A \cup B$ is not a decisive group either.*

Proof. Suppose all agents in A share the preferences $\succ_A$, the agents in B share the preferences $\succ$, where these preferences are given by,

$$\begin{array}{cc} \succ_A & \succ_B \\ x & z \\ y & x \\ z & y \end{array}$$

If there are agents that are in neither A or B, the preferences for these agents are completely arbitrary.

1) Since A is not a decisive group, $z \succeq x$

ii) Since B is not a decisive group, $y \succeq z$

iii) By transitivity, $y \succeq x$. Notice that all agents in A and B strictly prefer x to y, so we conclude that the union of A and B is not decisive over the pair (x, y), which by the previous claim means that the group is not decisive over any pair of alternatives. ■

To finish the proof we just need to observe that non-dictatorship implies that any single agent cannot be a decisive group. But that means, by repeated use the claim above that the set of all agents is not a decisive group either. This violates the Pareto principle since it says that if all agents think x is better than y, then society cannot rank x above y ■

10.1 What does Arrows theorem tell us?

The force of the result is that it says that any conceivable procedure to aggregate preferences will violate at least one of the assumptions, so the search for the universal aggregation procedure is a dead end, unless we are willing to commit to the sin of making interpersonal comparisons of utilities.

In principle, one could argue that the postulates of what is a reasonable social welfare function are too strict, but waving the unanimity requirement opens the possibility for a welfare function that picks an arbitrary alternative, irrespective of individual preferences. Similarly, waving non-dictatorship doesn't seem very interesting. The only remaining postulate to relax is thus the independence axiom.

If independence is waved there are lots of social welfare functions that fulfill the remaining axioms. For example, Borda rules would work just fine.

It was also noted (by Nobel Laurete Vickrey) that there is a close relation between the Independence requirement and possibilities for truthful revelation of preferences. Vickreys first insight is that if the independence axiom and a "non-perversity" postulate would be satisfied, then there would be no gains from misreporting preferences. In words, the non-perversity means that a change in preferences for an individual shall never change the social preferences in the opposite order (which is implied by the other assumptions).

Now, suppose that non-perversity is satisfied and we are considering the relative merits of alternatives x and y. Exactly like in majority voting between two alternatives there is then absolutely no point for a person who likes x to claim that her actual preferences are the opposite (we are now supposing that the alternative ranked the highest by society is actually implemented so that people care about the aggregation procedure). It is a dominant strategy to tell the truth since no matter what the others are doing: non-perversity implies that an individual who likes x to y and as a result of reporting this society nevertheless implements y, then the choice can't be x if the individual reports y. By the independence requirement, changes in reported preferences concerning other alternatives can have no effect on the relative ranking between x and y, so no matter what the other agents are doing there are no incentives to misreport preferences.

Vickrey conjectured that the converse is also true. If the independence criterion is not satisfied, then the social welfare function is not immune against strategic misrepresentations of preferences (which by Arrows theorem implies that no social welfare function over more than 3 alternatives is immune against strategic misrepresentations of preferences). This conjecture was proved some ten year later.

10.2 What Should a Public Economist Do?

One may argue that the problem of social choice as formulated by Arrow is too ambitious. Maybe society doesn't need a social preordering at all. For example, the Pareto criterion has proved very useful both in Public and many other fields.

Another much used "solution" is to put restrictions on individual preferences. Arrows' theorem relies crucially on the requirement that the welfare function must work for all possible collections of individual orderings that are conceivable. In some cases this may be an unnecessarily stringent assumption. For example, sometimes it may be natural to the problem that alternatives will be ranked in some systematic fashion.

The by far most popular restriction on individual preferences in the literature is that of "single-peaked" preferences that we have already discussed. Under single-peakedness assumptions, it is possible to find reasonable ways to aggregate preferences. Majority voting is the far most popular rule.

Another approach, common in particular in the macro literature is to use cardinal properties of individual preferences to generate a social welfare function. However, this is tricky business if the planner would have to rely on reported preferences since it is hard to avoid preferences to exaggerate preferences.

11 Bargaining over Public Goods and Externalities

Consider the following simplistic model of a production externality;

1. Bart can build a plant producing, say, timber. If building the plant some toxic waste will be emitted into the sea, killing a significant number of fish. Assume that Bart earns a profit of v_B if the timber plant is set up.

2. Lisa is a fisherwoman who fishes in the waters that will be polluted by Barts' plant (if set up). Assume that the value of her catch is v_L if the plant is not built and that it is 0 if the plant is built (i.e., she cannot sell salmon with two heads that contain too much PCB).

3. In order to facilitate a clear analysis of the efficiency problem *we assume that Bart and Lisa are the only economic agents in our model economy, that Bart only cares about the profit he can make and that Lisa only cares about the value of her catch.* For example, we assume that there is no agent that assigns any value to a clean ocean etc. This is only to simplify the discussion. If there are agents who assigns a value to a clean ocean, we can easily add such considerations without changing the qualitative insights from the analysis.

4. Bart and Lisa are also endowed with some good x that can be used as a transfer. Barts' utility function is

$$U_B = \begin{cases} x + v_B & \text{if plant built} \\ x & \text{if plant not built} \end{cases}$$

and Lisa has utility function

$$U_L = \begin{cases} x & \text{if plant built} \\ x + v_L & \text{if plant not built} \end{cases}$$

Let the endowments of x be e_B and e_L respectively and assume that the $e_B > v_L$ and $e_L > v_B$

11.1 No Intervention Benchmark

Since the setup is so simple, this is trivial. That is, Bart will build the plant if and only if $v_B \geq 0$. This is economically efficient if and only if $v_B \geq v_L$, whereas if $0 < v_B < v_L$, this leads to an economic inefficiency since if Lisa would transfer, say $\frac{v_S+v_B}{2}$ units of good x to Bart for the promise not to build, then Lisas utility is

$$\underbrace{e_L - \frac{v_L + v_B}{2} + v_S}_{\text{Lisas' utility if bribing not to build}} = e_L + \frac{v_L - v_B}{2} > \underbrace{e_L}_{\text{Lisas' utility if build}}$$

and Barts utility is

$$\underbrace{e_B + \frac{v_L + v_B}{2}}_{\text{Barts' utility if bribed not to build}} > \underbrace{e_B + v_B}_{\text{Barts' utility if build}}$$

11.2 Coases' Solution: Well-Defined Property Rights

One way (that is not quite complete) of explaining the problem that gives rise to the inefficiency is that there is nobody who have property rights to the sea. According to Coase (and most any other economist in the Chicago tradition) the cure to such a problem is to assign well-defined property rights.

Say first that Lisa has the property rights (or, equivalently, that citizens have "rights to clean water"). Then, Lisa will demand at least v_L to allow Bart to pollute. If $v_B < v_L$ Bart would therefore not be willing to Bribe Lisa (which is good from an efficiency point of view since the plant should not be built in this case). If on the other hand $v_B > v_L$, then any bribe in between these values would work, and the plant should be built in this case. With a little leap of faith we may assert that the bargaining outcome in this case should be efficient, in which case we would conclude that the right decision will always be made if Lisa has the right to clean water.

Symmetrically, suppose that we have "rights to pollute". Then, Bruce will not accept any bribe below v_B and Lisa is willing to pay at most v_L, so if $v_B > v_L$, Lisa will be unwilling to Bribe Bart not to build, whereas if $v_B < v_L$, any bribe in between the two values will lead Bruce not to build. In each case (with the same leap of faith as above), economic efficiency is guaranteed.

This example illustrates a very important cornerstone in modern liberal economic thought/ideology (it doesn't quite qualify as theory). The conclusion is that an efficient solution will be achieved *independently of who as assigned the property rights, as long as there is someone who holds the property rights.* Somewhat misleadingly, this is referred to as the Coase Theorem (since there is a lack of a both exact assumptions and proof the word "theorem" is somewhat puzzling). A more complete statement of the assertion is as follows: If property rights are well-defined and there are no transactions cost [my remark: I don't quite know how one should think of "transactions costs"], then bargaining will lead to an efficient solution. An implication of this would thus be that no government intervention is required to deal with externalities (a conclusion that would apply to public goods too, since this is a special case of an externality). All that is needed, is to specify property rights to everything that could possibly be affected by external effects.