There are many social choice rules : plurality, the Borda count, plurality with a runoff..... As the example in the first note of this section shows, these different rules matter, in that changing the social choice rule may change the alternative which is chosen.

So which rule should we use?

Rather than answering that question, here a slightly different question is considered : "What constitutes a 'good' social choice rule?".

Here are a few properties that might be required of a social choice rule, if we wanted it to work well.

UD ("universal domain") : for **any** group of people, and **any** pair of alternatives, the rule must choose a winner ; it must provide a ranking of the alternatives (such as "alternative A is ranked better than alternative B", or "alternative B is better than alternative C", or "the two alternatives are equally good") ; this ranking of the alternatives should depend (only) on how the different voters rank the alternatives

T : the ranking must be **transitive** : if A is ranked above B by the rule, and if B is ranked above C, then A has to be ranked above C (That is, there should be no endless cycles : C defeats A, which defeats B, which defeats C, which defeats A, and so on forever.)

P : the rule must obey the **Pareto principle** : if every single person thinks that alternative A is better than alternative B, then the rule must rank A above B ; this is a very minimal way of requiring that the rule somehow pays attention to the actual rankings of the voters

ND : the rule cannot be **dictatorial** ; the rule can give a lot of influence to person 1, for example ; but the rule can't be just "let's use person 1's ranking"

IIA : the ranking of alternatives A compared with alternative B done by the rule must be **independent of irrelevant alternatives** ; that is, whether A is ranked above B, or vice versa, should depend only on how the different people rank A compared to B, and not on how they rank A compared to some third alternative C, for example

For example, suppose that there is a bunch of people, and each of them ranks all the alternatives, and the rule uses those people's preferences to make its overall ranking, and then the rule ranks A above B. Now if some of the people's preferences change a little, but no person changes her opinion about A versus B, then the rule cannot change its ranking of A over B, since the changes in people's preferences are really irrelevant to the contest between A and B.

Arrow's Impossibility Theorem says that there is **no** social choice rule satisfying all the five properties listed above. That is, take any possible social choice rule : then there are some circumstances under which that rule does not satisfy one of the properties.

Examples

1. Plurality voting satisfies all the properties except for the independence of irrelevant alternatives.

It always can rank all the alternatives, no matter what are people's preferences, so it satisfies UD. The overall ranking of alternatives is based solely on the number of votes each alternative gets : so if A gets more votes than B, and B gets more votes than C, it must be true that A has received more votes than C : plurality ruling provides a transitive ordering (it satisfies T). If everyone ranks alternative A above alternative B, then alternative B is a first choice for nobody. It gets zero votes, and is ranked last (or tied for last) in the overall ranking. So property P is satisfied : if everyone prefers alternative A to alternative B, B cannot finish above A in the ranking provided by plurality voting. Plurality is also non–dictatorial (property ND) : if one person ranks alternative A first, and if everyone else ranks B first, then B will win the election.

To see why the outcome is dependent on irrelevant alternatives, suppose that there were 3 voters, and 4 alternatives, and that the preferences of the voters were

$$\begin{array}{ccc} 1 & 2 & 3 \\ A & A & C \\ B & B & D \\ C & C & B \\ D & D & A \end{array}$$

Using plurality voting, alternative A gets the votes of voters 1 and 2, and person 3 votes for C. A wins, and C finishes second, with B and D tied for third. The rule ranks alternative A above alternative C.

Now suppose preferences were to change to

$$\begin{array}{ccc} 1 & 2 & 3 \\ B & B & C \\ A & A & D \\ C & C & B \\ D & D & A \end{array}$$

Now voters 1 and 2 both vote for B, and voter 3 still votes for C. B is ranked first, then C, and then A and D tied for third. So the change in preferences (people 1 and 2 deciding that B was actually better than A) has led to a change in the ranking of A versus C ; before A was ranked higher than C, and now it's ranked lower.

But no–one has changed her relative rankings of A versus C. All that changed is that people 1 and 2 have changed their relative rankings between A and B. They still both think that A is better than C, as they did before. And person 3 still feels that C is better than A, as she did before. Yet the rise in the ranking (of people 1 and 2) of an irrelevant alternative, alternative B, has changed the way that the social choice rule (plurality voting) ranks alternative A relative to alternative C.

Why is that dependence on irrelevant alternatives a bad result? Suppose that A, B, C and D are four different candidates for a job. But we don't know whether they would really be willing to accept the job. So we'll vote, using plurality voting, and then offer the job to the candidates in order : so we'll give it to the candidate with the second–most votes if we are turned down with the candidate with the most votes.

What if candidate B isn't really interested in the job. How people rank this irrelevant candidate, who won't even accept the job, has changed the ranking of the two candidates who would accept the job, A and C. And that does not seem a very attractive property for a choice rule.

2. Plurality voting with a runoff also satisfies all the properties except for the independence of irrelevant alternatives.

It always provides a ranking of the alternative, based on the number of votes they get. (The overall winner is ranked first, followed by the loser of the runoff election, followed by the other alternatives in order of the votes which they received in the first ballot.) So UD holds. So does T, since there are no cycles in this ordering based on the number of votes received by candidates. P holds, since if A is preferred by everyone to B, then (as in the case of plurality voting), B gets no votes and finishes last, or tied for last. Similarly, ND holds, since a candidate preferred by only one voter cannot win the election.

So plurality with a runoff must not have the IIA property. Why? Two ways of explaining this. One, Arrow's Impossibility Theorem says that this must be the case : any social choice rule which always satisfies UD, T, P and ND must sometimes violate IIA, otherwise the theorem would not be true. The other way of showing IIA is violated? Use the same example used above in the case of plurality voting. This example also shows that plurality voting with a runoff does not satisfy the independence of irrelevant alternatives axiom.

3. Pairwise majority rule **does** satisfy IIA, since the ranking of A versus B depends only on how many people prefer A to B ; the whole point of "pairwise" is that the relative ranking of alternatives A and B is decided by a vote between those two choices (and only those two choices).

Pairwise majority rule satisfies UD. It also satisfies P ; if every person prefers alternative A to alternative B, then alternative A will win a vote between those two alternatives. It also is non–dictatorial (ND holds) : if only one person prefers A to B, and everyone else prefers B to A, then A wins the pairwise vote.

So the property which does not hold for pairwise majority rule is transitivity (T). The simplest example of the failure of pairwise majority rule to generate a transitive ordering is a situation in which there are three people (named 1,2, and 3), three different alternatives (alternatives A, B and C), and the preference orderings of the people are

1	2	3
A	B	C
B	C	A
C	A	B

In this case, A will defeat B in a pairwise vote, since people and 3 both rank A above B. B defeats C in a pairwise vote, since 1 and 2 both rank B above C. But A does not defeat C : both 2 and 3 rank C above A, so that C defeats A by a 2–to–1 vote. The overall ranking is not transitive. There is a (potentially endless) cycle : A defeats B which defeats C which defeats A which defeats B and on and on.

4. The Borda count provides a transitive ordering of all the alternatives : they are ranked by the number of points each alternative gets. So property T holds. UD holds, since the Borda scoring rule can be applied for any set of preference ordering of alternatives by voters. It also satisfies the Pareto principle P ; if everyone likes A better than B, then A will get more points than B. ND is satisfied as well : if one person ranks alternative A first, and everyone else ranks alternative A last, then A will not be ranked first under the Borda count rule.

So (like plurality voting, and plurality voting with a runoff), it must violate the independence of irrelevant alternatives axiom IIA. The first example above (under "plurality voting") confirms this. Using a 3-2-1-0 scoring system (in which people's top choice gets 3 points,their second choice 2 points, and so on), with the first set of preferences in the example, A gets a score of 6, B and C each get 5, and D gets 2. When the preferences change to the second set in the example, now B gets 7, C gets 5, A gets 4 and D gets 2. So changing the preferences moves A from being ranked above C to being ranked below C, even though no person changed her ranking of A versus C.

5. The unanimity rule certainly is non–dictatorial (ND holds) : A is ranked above B only if everyone ranks A at least as high as B. It obeys the Pareto principle P, since it will rank A above B exactly when every voter ranks A at least as high as B. IIA holds as well : whether A is ranked at least as high as B depends only on whether every voter ranks A at least as high as B, and not at all how any voter feels about any third alternative.

But what of the other two properties UD and T it violates depends on what we do when preference is not unanimous. We could say that we simply won't rank A versus B, unless one of them is preferred unanimously to the other. In that case, property UD is violated, since some pairs of alternatives just can't be ranked.

Alternatively, we could say that it will be called a tie between A and B, unless one is preferred unanimously to the other. In this case we can rank any pair of alternatives, so that UD holds. But the ranking is not transitive. If the preferences were

1	2	3
B	C	C
C	A	A
A	B	B

then A would be ranked at least as good as B ; they would be tied, since two people prefer A to B, and one person prefers B to A. Similarly B is ranked at least as high as C. They're tied again, since one person prefers B to C and two people prefer C to B. Property T requires that if A is ranked at least as high as B, and B is ranked at least as high as C, then A must be ranked at least as high as C. But that's not the case in this example : A is ranked below C, since all three people prefer C to A.

6. Here's a nice simple rule not discussed earlier : rank the alternatives in alphabetical order. UD works, since we can always use the rule. So does T : alphabetical ordering is transitive. IIA holds, since the relative ranking of any two alternatives depends only on where those two alternatives stand in the alphabet, not on any third alternative. ND holds, since a person who ranks the alternatives in some other order (other than alphabetical order) will not have her way.

So even though the rule seems a very silly one, more or less random, it obeys 4 of the 5 properties. The only property violated is P. If every person agreed that B was better than A, then the rule would not respond to this unanimous preference, instead ranking A higher because it comes higher in the alphabet.

Strategic Behaviour

Arrow's impossibility theorem says that we have to give up one of the 5 properties listed above in deciding on a social choice rule. It may seem that, if we have to give up one "nice" property, then the independence of irrelevant alternatives (IIA) should be the one to go. It seems to be the most complicated, and perhaps the least vital. Dictatorship does not seem an attractive system, so that ND should not be sacrificed. We want a rule that is at least somewhat responsive to the actual feelings of the voters, so we should keep P. And it is important to be able to get an answer to the question of the overall ranking, no matter what are people's preferences — so UD and T are pretty essential.

But it turns out that IIA is also not a property which we want to lose. So far, there has been an implicit assumption maintained when evaluating how rules work, that people vote **sincerely**. Whoever is administering the social choice rule does not actually know the people's rankings. She is just asking the people to announce their rankings, and trusting that the answers given are the actual rankings. As in public good provision, it is very important for efficiency that we induce people to reveal their true rankings.

Is it in people's own interest to vote always for their preferred alternative, if the social choice rule is pairwise majority rule? Is it in their own interests to vote always for their most preferred of all the alternatives under plurality rule?

If people find that they may be better off voting for an alternative which is not their top choice, or listing alternatives (under a Borda count rule) differently than their actual preferences, then they are said to be voting **strategically**. (Sometimes this behaviour is described as "sophisticated voting".)

There is nothing necessarily wrong, or immoral, in strategic voting. People do behave in many elections. For example, suppose that the choice rule is plurality voting, the rule which we actually do use in electing representatives in Canada. Suppose that person 1 regards A as the best candidate, followed by B, and C as the worst candidate. Suppose as well that she has some idea of how other people feel, and that she observes that most of the other voters are split, almost evenly, between two types of voter, types 2 and 3, with the preferences as follow (where I have also listed person 1's own true ranking of the candidates) :

1	2	3
A	B	C
B	A	B
C	C	A

(For example, suppose that A is the NDP candidate, B the Liberal, and C the Conservative.)

What should voter #1 do? She expects that there are very few voters who feel as she does, and who rank A as the best candidate.She expects the others are about evenly split between those who are likely to vote for B, and those who are likely to vote for C. If she votes "sincerely", by voting for A, her vote will be wasted in a sense. She does not expect many other people to vote for A. But if the others are split pretty closely between the B supporters and the C supporters, then there is a chance that she can influence the election. She may feel that she is better off voting for B, rather than for A. By voting for B, she may possibly help C get defeated. And C is her least–liked alternative.

So under plurality voting, people may not want to vote sincerely. If they feel that their most–preferred candidate has little chance of winning, then they may want to vote for their second (or third, or fourth) choice, if that candidate has a chance of winning, and if there is another candidate who has a chance of winning, whom she really does not like. [Another example of this incentive to vote strategically was the plea made by supporters of Gore in the 2000 American presidential election, that supporters of Ralph Nader should not "waste" their votes by voting for their first choice, but instead should vote for their second choice, Gore, to prevent Bush from winning.]

Is the incentive to vote strategically a problem which is unique to the plurality rule? It is not. Suppose that we use plurality with a runoff. Now suppose that a voter feels that her first choice, A, is virtually certain of finishing first or second in the first round, and making into the runoff. But she feels that it is a close race between B and C for second place (and a place in the runoff election) in this first round. She finds B almost as attractive an alternative as A, but thinks that C is a really bad alternative. So she would really really like to make sure that candidate C does not make it into the runoff election against A, because perhaps C could win that election. What she wants is for A and B to finish in the top two positions in the first round, so that A and B are the choices in the runoff. What is she likely to do in the first round voting? She is confident that her preferred candidate, A, will get into the runoff without her help. So she may choose to vote for B in the first round, to prevent C from getting into the runoff election. So she may vote strategically in the first round. In the runoff, she will have no incentive to vote strategically : with just two candidates, why would she vote for anyone but the one she prefers?

So this example shows that strategic voting may be an attractive option for people in the first round of plurality voting with a runoff. The example also shows some of the problems in voting strategically. What if a lot of A supporters feel the same way as this voter? : they like A best, they hate C, they are confident that A will make it into the runoff, but are worried that C might make it as well. Then many of them may vote strategically for B, rather than A, in the first round, confident that the other A supporters will vote sincerely. But if a lot of A supporters behave this way, perhaps not enough of them will be left to vote sincerely. Even though A has lots of supporters, if they all vote strategically (for B) in the first round, then A might not make it into the runoff : precisely because her supporters are so sure that she will make it into the runoff.

Strategic voting is also a problem with the Borda count. Suppose this time that nearly all voters find C to be the worst of three alternatives. But they are closely split over whether A or B is the best. That is, there are roughly equal numbers of type–1 and type–2 voters, with preferences

1	2
A	B
B	A
C	C

A type–1 voter has a strong incentive now to vote strategically : to mark A first, but then C second, rather than B, and B third. This voter knows that the real race is between A and B for first choice : C is clearly going to finish last by a wide margin. But if she moves C up, and B down, in her rankings, then she can help A defeat B in the race for first. Her altered ranking gives B the fewest points that she can give him,which will help A defeat B. This strategy does give C more points than she would like, but she is confident that C is so unpopular that her extra points for C will not be nearly enough to move C up from third place.

So again there is an incentive to vote strategically, with the Borda count. And again, voting strategically can pose some problems. What if most voters figure out the same clever strategy : to rank C second to help their preferred alternative (A or B) win the race for first? if enough people behave this way (and if, perhaps, there are a few genuine C supporters), then this sophisticated behaviour might actually get C to win, not an outcome that many of the voters wanted.

The incentive to vote strategically is probably something which we would like to avoid. It forces people to make decisions based on what they think other people are doing, and makes results very sensitive to polls, rumours and speculation. So a nice property which we might want a social choice rule to have would be

S : each voter's best strategy, given the choice rule, should be to vote **sincerely**

Some people might even find S a more useful property than IIA. Unfortunately, many of the rules we use, such as plurality voting, do not satisfy S. [However, some rules do, most notably pairwise majority rule.]

More unfortunately, there is another impossibility theorem. It follows from the fact that S is (more or less) equivalent to IIA. replacing IIA by S does not make it possible to find a "good" social choice rule.

Gibbard's and Satterthwaite's impossibility theorem : there is no social choice rule which satisfies all 5 of : UD, T, P, ND and S. If a rule (such as plurality) satisfies UD, T, P and ND, then there will be some circumstances under which voters are better off voting strategically.