Mathellaneous by Norman Do

The mathematics of voting

1 Introduction

In the animal kingdom, leaders are chosen by instinct, tradition, and occasionally the bumping of heads. We humans, on the other hand, have moved away from such primitive and barbaric behaviour. Democracy, from the same people who brought us the Olympic Games, necessitated a method to elect leaders in a way that would accurately reflect the will of the people. Thus, voting was born.

Having been practised for thousands upon thousands of years, voting is now rife in our society, whether it be used to opt between holiday destinations, choose a new member for a committee or elect a head of state. And what could be simpler? All it takes is for someone to count the votes and declare whichever candidate has the most to be the winner. Indeed, the subject of voting did not catch the eye of the mathematical world until the late eighteenth century when some clever people noticed that there were strange and paradoxical anomalies lurking behind this seemingly elementary process. Such simple observations gave rise to the marriage of mathematics with the field now known as social choice theory, which analyzes how collective decisions are made by a set of voters.

2 Two Alternatives

Voting between two alternatives is so simple that it can be seen being practised by young children in their playground politics. If more children want to play hide-and-seek than chasey, then it is generally understood that that is the game that they should play. This system, known for obvious reasons as *majority vote*, seems like the optimal way to choose a winner in order to please the most people and displease the fewest. But let us ask ourselves the following question...

> Are there any other ways, besides majority vote, to elect a winner from two alternatives?

What we are on the hunt for is some method of producing from an input of many preferences an output consisting of only one preference. This is an example of what is known in the literature as a *social choice function*. More precisely, a social choice function

- accepts as input a sequence of preference lists which strictly rank the elements of some set, and
- produces as output a winner or a list of tied winners from the set. The input sequence of preference lists is known as a *profile* and the output list of winners is called the *social choice*.

Of course, whether we are voting between hide-and-seek and chasey or between the lesser of two evil candidates for the leader of a nation is totally irrelevant. Also, in this simple case of two alternatives, ranking all of the alternatives is equivalent to choosing a preferred one. So, for convenience, we can give the two alternatives abstract names, such as $+1$ and -1, and a voter's preference list can simply be described as one of these two numbers. Furthermore, a profile can be given by an m-tuple of numbers which are either $+1$ or -1, where m corresponds to the number of voters. As long as there are no ties involved, our problem now translates into finding a function $f : \{-1, +1\}^m \to \{-1, +1\}$. But the number of such functions — which happens to be 2^{2^m} — is phenomenally large, even for reasonably small values of m. The following are just three simple examples of social choice functions for two alternatives.

- Dictatorship: Let one of the voters be the dictator and then let the social choice simply coincide with their preference.
- Constant: No matter how people vote, let the social choice always be $+1$.
- Parity: Let the social choice be $+1$ if an even number of people vote for $+1$ and let it be -1 otherwise.

It should be clear that all of these, although legitimate examples of the abstract notion of a social choice function, are preposterous attempts to accurately reflect the will of the people. The dictatorship does not treat all voters equally while the constant function does not treat both alternatives equally. The parity function allows a voter who changes his or her mind from -1 to $+1$ to make the social choice change in the reverse direction. So we need to put some conditions on our function to avoid these pathological examples. The following are three conditions which seem reasonable along with their translations into the more abstract language of mathematics.

- Anonymous: The social choice function should treat all voters equally.
  If $(x_1, x_2, \ldots, x_m)$ is a permutation of $(y_1, y_2, \ldots, y_m)$, then $f(x_1, x_2, \ldots, x_m) = f(y_1, y_2, \ldots, y_m)$.
- Neutral: The social choice function should treat both alternatives equally.
  For all $(x_1, x_2, \ldots, x_m)$, $f(-x_1, -x_2, \ldots, -x_m) = -f(x_1, x_2, \ldots, x_m)$.
- Monotone: Voting for someone cannot hurt their chances.
  If $x_k \geq y_k$ for all k, then $f(x_1, x_2, \ldots, x_m) \geq f(y_1, y_2, \ldots, y_m)$.

The natural question to ask now is ...

> Besides majority vote, are there any social choice functions for two alternatives which are anonymous, neutral and monotone?

If we restrict our interest to the case where there is an odd number of voters and no ties, then the answer is given in the following...

May's Theorem (1952): Suppose that we have a social choice function for two alternatives which

- has an odd number of voters;
- does not allow ties; and
- is anonymous, neutral and monotone.

Then the social choice function is a majority vote.

Problem: Prove May's Theorem by finding all functions $f : \{-1, +1\}^m \to \{-1, +1\}$ for odd m which are anonymous, neutral and monotone.

3 Three or More Alternatives

May's Theorem tells us that voting between two alternatives offers no surprises, so let us turn our attention to the more interesting case of three or more alternatives. Of course, one might be tempted to think that we are just about ready to conquer the world of voting, as long as we can generalize May's theorem to a larger number of alternatives. The most obvious way to do this is to simply take as the social choice the alternative whom most people think is the best. This method is known as *plurality voting* and is widely adopted, most notably in the United States presidential elections. Despite this fact, it is generally accepted in the social choice theory community that plurality voting is flawed. The following profile of voter preferences from a hypothetical election will give just one reason of many as to why this is so.

Suppose that there are fifteen voters and three alternatives. It may be the case that six of the voters prefer A to B to C, while five prefer C to B to A and the rest prefer B to C to A. This information can be conveniently captured in the following table.

6 voters	5 voters	4 voters
A	C	B
B	B	C
C	A	A

If this election were to use plurality voting as the social choice function, then it is clear that alternative A would be the winner. However, notice that A is the last choice for nine of the fifteen voters…does choosing A accurately reflect the will of the people? Most people would be inclined to think not and would perhaps even believe that A is the worst possible social choice. One might be tempted to think that such an anomaly is particular to this case and a handful of other concocted examples. But this is far from the truth and the flaws of plurality voting have been witnessed in many situations, most notably in the controversial US presidential race of 2000.

To avoid this flaw in plurality voting, it seems sensible to take advantage of the full preference lists of the voters, and voting systems which use this information are known as *preferential voting systems*. However, with this extra information at our disposal, there are far more options for us in determining a winner. Social choice functions abound and have been constructed by such voting conscious mathematicians as Charles Lutwidge Dodgson[1] and Edward John Nanson[2]. The following are three simple examples of social choice functions.

Borda Count: In order to overcome the deficiencies of plurality voting, Jean-Charles de Borda introduced in the late eighteenth century a voting system which would take advantage of each individual's intensity of preference for each alternative. He proposed assigning a number of points to each alternative, equal to the distance from the bottom of each voter's preference list. Thus, an alternative would receive 0 points for each last place vote, 1 point for each next-to-last place vote, all the way up to $n-1$ points for each first place vote, where n is the number of alternatives. The winner is, of course, the alternative that has been awarded the most points. You may be wondering what is so special about the numbers $0, 1, 2, 3, 4, \ldots$ that they should be the number of points assigned to the alternatives. Why

[1] Charles Lutwidge Dodgson (1832–1898) is better known by his pseudonym Lewis Carroll, under which he wrote the popular book "Alice's Adventures in Wonderland". Less well known is the fact that he was a mathematician and particularly interested in the mathematics of voting. In 1876, his exploration into voting led him to design a social choice function, but since then mathematicians have shown that the problem of calculating a winner by his method is NP-hard.

[2] Edward John Nanson (1850–1936) was a professor at The University Of Melbourne for 48 years. He was a strong advocate of preferential voting in the years leading up to the federation of Australia.

not use a sequence like $1, 2, 3, 5, 8, \ldots$ or $1, 2, 5, 14, 42, \ldots$? Such social choice functions which generalize the Borda count are known as *positional* voting systems.

The Hare System: Also known as single transferable vote, this voting system was introduced by Thomas Hare in 1861. The utilitarian philosopher John Stuart Mill described it as "among the greatest improvements yet made in the theory and practice of government" and it is currently in use to elect officials in Australia, Malta and Ireland. The idea behind the Hare system is that the winner should be voted in by a majority of the voters. But this is a rare occurrence when there are many alternatives to choose between, so what can we do? Simple — just delete some of the alternatives! More precisely, if any alternative is at the top of a majority of preference lists, then they are declared the winner. If not, then the alternative which appears at the top of the fewest preference lists is eliminated and the process is repeated. Of course, it may be the case that there is more than one alternative tied with the fewest number of votes and in that case, we can delete all of them. The process terminates when all remaining alternatives are at the top of the same number of preference lists. These remaining alternatives are the winners.

Dictatorship: This social choice function is one of the simplest to implement! One of the voters is assumed to be a dictator and the social choice is simply whoever is at the top of the dictator's preference list.

The following problem highlights the distinction between these different social choice functions. In particular, notice that in many elections, the person who wins can be heavily dependent on the type of voting system in use.

Problem: Consider the following simple profile with seven voters and four alternatives. Check that A wins using plurality, B wins using the Borda count, C wins using the Hare system and D wins using a dictatorship by voter 7. Who would you choose as the winner?

Voter 1	Voter 2	Voter 3	Voter 4	Voter 5	Voter 6	Voter 7
A	A	A	C	C	B	D
B	B	B	B	B	C	C
C	D	D	D	A	A	B
D	C	C	A	D	D	A

4 Condorcet's Voting Paradox

One of the first surprise results in voting was noticed by the Marquis de Condorcet, a contemporary of Borda, who considered the possibility of determining a social choice by using pairwise elections. He proposed that the alternative which would beat all others in a one-on-one majority vote should be the social choice. This alternative is known in the literature as the *Condorcet winner*. This voting system offers a natural generalization to majority vote which overcomes the flaws of plurality voting by using each individual's full list of preferences.

However, such a voting system is not without its flaws, as can be demonstrated by considering the following voting profile.

23 voters	17 voters	2 voters	10 voters	8 voters
A	B	B	C	C
B	C	A	A	B
C	A	C	B	A

Suppose now that candidate C dropped out of the running, leaving A and B in a head-to-head contest. Then using majority vote, as May's Theorem suggests we should, shows us that A would beat B by 33 votes to 27. In other words, a majority of voters prefer A to B, and we can encapsulate this statement in the convenient notation $A \succ B$. Now by comparing B and C, we find that $B \succ C$, winning by 42 votes to 18. It seems clear from this analysis that society prefers A to B and B to C, thereby making A the logical social choice.

But wait a minute... why did we neglect to compare alternative C with alternative A? Indeed, it turns out that in a direct comparison, C would have beaten A by 35 votes to 25. So we have the relations

$$A \succ B, \qquad B \succ C, \qquad C \succ A,$$

and it turns out that for this particular profile that there is no Condorcet winner. Common sense dictates that if a person prefers an apple to a banana and a banana to a cherry, then they should prefer an apple to a cherry. However, Condorcet's voting paradox tells us that society as a whole does not obey such common sense. If society prefers A to B and B to C, it may also occur that society prefers C to A. Of course, this argument relies on the assumption that society prefers one option to another if a majority of voters do. But what other ways are there in which to decide, given that we only have everyone's preference listing to deal with?

Condorcet himself realized the possibility for this type of intransitive behaviour to occur in the following much simpler example, which is now known as the Condorcet profile.

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

Even on this small scale, we have the relations $A \succ B$, $B \succ C$ and $C \succ A$, which are reminiscent of the famous children's game known as rock-paper-scissors.

Condorcet's Voting Paradox: There are particular profiles in which a social choice function must choose a particular alternative X as the winner, even though a majority of people prefer some other alternative Y.

This result from the late eighteenth century gave mathematicians a preview of the paradoxes that were in store for future social choice theorists. The upshot of it is that even though voting between a pair of alternatives is easy, pairwise voting falls apart when three or more alternatives are involved.

Problem: Show that for three voters and three alternatives, the probability that a Condorcet winner exists is $\frac{17}{18}$.

5 Arrow's Theorem

Thus far, our exploration of voting has only dealt with social choice functions, voting systems which turn a profile into a set of tied winners. In this section, we will be interested in more general voting systems which turn a profile into a ranking of the alternatives. More precisely, a *social welfare function*

- accepts as input a sequence of individual preference lists of some set, and
- produces as output a listing (perhaps with ties) of the set. This list is called the *social preference list*.

Social welfare functions are easy to find, and it turns out that we already have a few examples of them. For example, the Borda count can easily be converted from picking a winner to determining a social preference list by ordering the alternatives from highest to lowest depending on the number of points that they have received. The Hare system can also be converted into a social welfare function. To determine the set of people tied for first place is easy — just take the set of winners. To determine the set of people tied for second place, just delete the alternatives tied for first and repeat the Hare system. Iterating the procedure yields an ordered social preference list, perhaps with ties, as desired. This construction is not specific to the Hare system, but actually shows that every social choice function gives rise to a social welfare function. Of course, the reverse is also true, by obvious reasons. Note also that, in terms of social welfare functions, a dictatorship assumes that there is a voter who is a dictator and the social preference list coincides exactly with the dictator's preference list.

Now that we have the definition and some examples of a social welfare function, it seems natural for us to ask the question...

Can we find a *reasonable* social welfare function?

Of course, the answer to this question will depend on what exactly we mean by the imprecise and subjective term "reasonable". It is not overly difficult to find properties that almost everyone will agree are desirable in a social welfare function. One such property is illustrated in the following dialogue, which takes place in a dimly-lit fancy restaurant.

Waiter:	Good evening, madam. Where would you and your family like to sit — in the smoking section or the non-smoking section?
Madam:	I'd prefer non-smoking, but let me ask my family first... [turning to family] Where would you all like to sit?
Father:	I'd rather non-smoking.
Son:	Yeah, I hate the smoking section.
Daughter:	Me too!
Waiter:	Well, in that case, let me put you all in the smoking section!

Anyone in their right mind would find the waiter's decision to be incomprehensibly absurd. And it is surely reasonable for us to exclude from our consideration all social welfare functions which are incomprehensibly absurd. We do this by requiring the following property to be satisfied.

Unanimity: If every voter has the same preference list, then the social preference list should agree with it.

Another desirable property that we might impose on our social welfare function is illustrated in the following dialogue, which takes place in the very same dimly-lit fancy restaurant only minutes later.

Waiter:	Good evening, madam. Can I get you a drink to start off with?
Madam:	Yes, please. What kind of juices do you have?
Waiter:	We have apple juice and orange juice.
Madam:	OK, well I'll have the orange juice, thanks.
Waiter:	Excellent choice, madam! Oh, I just remembered... we also have cranberry juice available.
Madam:	You also have cranberry? Well in that case, I'll have the apple juice!

This time, it is the waiter's turn to be surprised by the lady's absolute irrationality. Why on earth would the presence of a third alternative alter which one of the first two is preferred? Again, it is surely reasonable for us to exclude from our consideration all social welfare functions which are absolutely irrational. We do this by requiring the following property to be satisfied.

> Independence of Irrelevant Alternatives: The relative positions of X and Y in the social choice should depend only on the relative positions of X and Y in the preference list of each voter. In other words, if every voter changes their preference list but decides to keep the relative positions of X and Y the same, then the social choice should keep the relative positions of X and Y the same.

Now we seem to have two desirable properties that any reasonable social welfare function should possess. Of course, we could go searching for more, but then we might be here all day without knowing when to stop. Anyway, it might pay to not be so greedy for the moment. So let us now ask the question…

> Can we find a social welfare function which satisfies unanimity and independence of irrelevant alternatives?

The quick-witted reader might already have noticed that there is an obvious, although undesirable, candidate for such a social welfare function — namely, a dictatorship. So for those quick-witted readers, let us rephrase our question just slightly.

> Can we find a social welfare function which satisfies unanimity and independence of irrelevant alternatives without being a dictatorship?

This question was answered by Kenneth Arrow in 1950 with a surprising and resounding, "No!"

Arrow's Theorem (1950): Suppose that we have a social welfare function which

- has at least three alternatives;
- satisfies unanimity; and
- satisfies independence of irrelevant alternatives.

Then the social welfare function is a dictatorship.

From our earlier discussion, we were led to the fact that any reasonable social welfare function should satisfy unanimity and independence of irrelevant alternatives, and probably a host of other desirable properties as well. But Arrow's Theorem tells us that even with these two obvious conditions, a dictatorship is unavoidable. In 1972, Kenneth Arrow was awarded the Nobel Prize in Economics for "pioneering contributions to general economic equilibrium theory and welfare theory".

Paul Samuelson, himself a Nobel laureate in Economics, put it this way…

> "The search of the great minds of recorded history for the perfect democracy, it turns out, is the search for a chimera, for logical self-contradiction. New scholars all over the world — in mathematics, politics, philosophy, and economics — are trying to salvage what can be salvaged from Arrow's devastating discovery that is to mathematical politics what Kurt Gödel's 1931 impossibility-of-proving-consistency theorem is to mathematical logic."

Problem: Suppose that we have a social welfare function which has at least three alternatives, satisfies unanimity and satisfies independence of irrelevant alternatives. Show that the social welfare function will never produce ties in the output (without using Arrow's Theorem, of course).

In light of the this fact, Arrow's Theorem boils down to the following problem of purely mathematical content. It is interesting to note that this problem and its solution may never have been uncovered by mathematicians were it not for democracy. The truly adventurous reader may like to try their hand at proving it before reading the next section.

Problem: Let $f : S_n^m \to S_n$ be a function which takes m-tuples of permutations to permutations and which satisfies the following two properties[a].

- For all $p \in S_n$, the function satisfies $f(p, p, \ldots, p) = p$.
- If $(p_1, p_2, \ldots, p_m)$ and $(q_1, q_2, \ldots, q_m)$ are m-tuples of permutations and the integers a and b satisfy

$$\operatorname{sign}[p_i(a) - p_i(b)] = \operatorname{sign}[q_i(a) - q_i(b)]$$

for all i, then

$$\operatorname{sign}[P(a) - P(b)] = \operatorname{sign}[Q(a) - Q(b)]$$

where $P = f(p_1, p_2, \ldots, p_m)$ and $Q = f(q_1, q_2, \ldots, q_m)$.

Then there exists an integer k such that for every m-tuple of permutations $(x_1, x_2, \ldots, x_m)$

$$f(x_1, x_2, \ldots, x_m) = x_k.$$

[a] As usual, S_n denotes the symmetric group on n elements and we will consider an element of S_n to be a bijection $f : \{1, 2, \ldots, n\} \to \{1, 2, \ldots, n\}$.

6 A proof of Arrow's Theorem in five bite-sized pieces

We have now come to the meaty part of the exposition, into which much of the mathematical argument has been condensed. For easier digestion, the proof of Arrow's Theorem has been divided into five bite-sized pieces of steadily increasing size.

Bite-sized piece 1: If every voter ranks X over Y, then society ranks X over Y.

Consider the profile where every voter has the same preference list with X at the top and Y second from the top. By unanimity, society must also have X at the top and Y second from the top. But independence of irrelevant alternatives tells us that whether society ranks X over Y or not depends only on each voter's relative ranking of X and Y. In particular, if everyone were to rearrange their list while keeping X over Y, then the social choice should not change. So we conclude that for any profile where every voter ranks X over Y, society must also rank X over Y.

Bite-sized piece 2: If every voter ranks A at the top or bottom of their preference list, then society ranks A at the top or bottom.

Let us assume on the contrary that there is a profile where every voter ranks A at the top or bottom but society does not. Then there must be three distinct alternatives A, B and C such that society ranks B at least as high as A and A at least as high as C.

But consider now what happens if every voter moves their ranking of C to be just over B. Since A occupies an extremal position in every preference list, this does not disturb any voter's relative ranking between A and B. Nor does it disturb any voter's relative ranking between A and C. So society must continue to rank B at least as high as A and A at least as high as C. Hence, society must rank B at least as high as C. But as we have shown above, if everyone ranks C over B, society must also rank C over B. Since society cannot simultaneously rank B at least as high as C and C over B, we have the desired contradiction.

Bite-sized piece 3: There exists a profile at which A is at the bottom of the social ranking and a voter who can move A to the top of the social ranking by changing his or her preference list.

Consider a profile where every voter has the same preference list with A placed at the bottom. By unanimity, society must also place A at the bottom. Now consider what happens to the social ranking when each voter in turn moves A from the bottom of their ranking to the top. After this process has finished, A is at the top of every voter's ranking and hence, must be at the top of the social ranking as well. So there must have been some voter $V(A)$ whose change caused A to move from the bottom of the social ranking. Let profile I denote the profile just prior to $V(A)$ moving A from the bottom to the top and let profile II denote the profile just after.

1	2	$\cdots$	$V-1$	V	$V+1$	$\cdots$	N
A	A	$\cdots$	A	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	A	A	$\cdots$	A

Profile I

1	2	$\cdots$	$V-1$	V	$V+1$	$\cdots$	N
A	A	$\cdots$	A	A	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	A	$\cdots$	A

Profile II

Bite-sized piece 4: The voter $V(A)$ is a dictator over any pair B and C not including A. In other words, if $V(A)$ ranks B over C, then society ranks B over C, and if $V(A)$ ranks C over B, then society ranks C over B.

Let us consider profile III which is constructed from profile II in the following way.

- Let $V(A)$ move B and C to be just above and below A, so that his or her first three preferences are B, A and C.
- Let all other voters keep A in the same extremal position as in profile II but let them change the relative order of B and C however they please.

1	2	$\cdots$	$V-1$	V	$V+1$	$\cdots$	N
A	A	$\cdots$	A	B	C	$\cdots$	B
C	B	$\cdots$	C	A	B	$\cdots$	C
B	C	$\cdots$	B	C	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	$\star$	$\cdots$	$\star$
$\star$	$\star$	$\cdots$	$\star$	$\star$	A	$\cdots$	A

Profile III

Note that in profiles I and III every voter has the same relative ranking for A and B. So it must be the case that they yield the same relative ranking in their output, due to independence of irrelevant alternatives. And since the social ranking for profile I places B over A, the social ranking for profile III must also place B over A.

Also note that in profiles II and III, every voter has the same relative ranking for A and C. So it must be the case that they yield the same relative ranking in their output, due to independence of irrelevant alternatives. And since the social ranking for profile II places A over C, the social ranking for profile III must also place A over C.

So the social ranking for profile III places B over A and A over C, and hence, B over C. But remember that we let all voters other than $V(A)$ arrange their relative ranking of B and C arbitrarily. But by independence of irrelevant alternatives, it follows that whenever $V(A)$ places B over C, society must as well. And by a totally analogous argument, if $V(A)$ places C over B, society must as well. So $V(A)$ is a dictator over any pair B and C not including A.

Bite-sized piece 5: The voter $V(A)$ is a dictator over any pair A and B. In other words, if $V(A)$ ranks A over B, then society ranks A over B, and if $V(A)$ ranks B over A, then society ranks B over A.

Now one thing you might be wondering is what is so special about alternative A. Indeed, it is contemptible to play favourites with letters, so it may serve us well to consider what might have happened had we begun our construction with alternative B instead. Supposing we had done this, we would have found out that there exists some voter $V(B)$ who is a dictator over every pair which does not involve B. In particular, $V(B)$ is a dictator over the pair A and C which means that society's relative ranking of A and C must always agree with $V(B)$'s relative ranking of A and C.

But in profiles I and II we noted that $V(A)$ had the power to alter the relative rankings of A and C, while everyone else kept their preference lists the same. So it must be the case that $V(A)$ and $V(B)$ are one and the same person! Of course, there is nothing from stopping us running through the whole process with alternative C and consequently, we would have found out not only that $V(C)$ is a dictator over every pair which does not involve C, but also that $V(A)$, $V(B)$ and $V(C)$ are all one and the same person. It turns out that this particular one person is a dictator over every pair of alternatives and hence, the social welfare function in question must be a dictatorship.

7 Strategy-Proofness and the Gibbard-Satterthwaite Theorem

Consider an election where there are four voters wishing to choose a winner from a set of four alternatives. Suppose that their preferences for the alternatives are as shown in the profile below and that the social choice function to be used is the Borda count.

Voter 1	Voter 2	Voter 3	Voter 4
A	C	C	B
B	D	B	C
C	B	A	D
D	A	D	A

A quick tally of the votes reveals that C wins the election with a count of nine, closely followed by B with eight, then A with four and finally D with a measly three points. If you were voter 1, you probably couldn't help but feel chagrined by such an election outcome. So suppose now that you were crafty and devilish and instead of submitting your true preferences $ABCD$, you instead submitted the insincere preference list $BADC$. Then the resulting profile would have looked a little more like this.

Voter 1	Voter 2	Voter 3	Voter 4
B	C	C	B
A	D	B	C
D	B	A	D
C	A	D	A

And running through the Borda count for this profile shows that B would have won with nine points, followed closely by C with eight points, then D with four and finally A with a measly three points. This quick calculation highlights an important flaw of the Borda count: that a crafty and devilish voter can sometimes obtain a better outcome by voting insincerely!

A social choice function is said to be *strategy-proof* if there is no profile in which one of the voters can vote insincerely in order to obtain a better outcome. Unfortunately, this definition is somewhat incomplete, since it is difficult to define what exactly a better outcome for a voter is. Suppose, for example, that your true preference list is $ABCD$ and that voting sincerely produces A and D as tied winners, while voting insincerely might produce B and C as tied winners. Which of these two choices is a better outcome for you? If, for example, ties were to be broken by the flip of a coin, it might depend on whether you are a pessimist or an optimist. But instead of worrying about whether the glass is half-empty or half-full, let us just sweep the problem under the rug and focus exclusively on social choice functions in which there are no ties. In such cases, it is obvious which one of two outcomes is better for a particular voter.

Now there are several reasons why it is desirable to have a social choice function which is strategy-proof, such as the following.

(1) Insincere voting introduces an element of randomness into collective decisions.
(2) Unequal strategic skills amongst voters means that they will not be treated equally.
(3) Voters will waste resources in making strategic calculations.
(4) Voters are encouraged to conceal their preferences, reducing a flow of information that might aid in collective decision making.

So the obvious question to ask now is...

> Can we find a social choice function which has no ties and is strategy-proof?

An example of a strategy-proof social choice function is the constant one in which one particular alternative is always declared the winner, no matter how people vote. Of course, this is terribly unfair to the other candidates, so to avoid this pathological example, we might look for social choice functions which allow all alternatives to win. So the question we are trying to answer is now the following...

> Can we find a social choice function which has no ties, is strategy-proof and allows all alternatives to win?

A little bit of thought might suggest that a dictatorship satisfies all of these conditions, so let us refine our question just one more time...

> Can we find a social choice function which has no ties, is strategy-proof, allows all alternatives to win, but is not a dictatorship?

This question was answered independently by Allan Gibbard and Mark Satterthwaite with a surprising and resounding, “No!” Thus did mathematics deal another devastating blow to democracy.

Gibbard-Satterthwaite Theorem (1973): Suppose that we have a social choice function which

- has at least three alternatives;
- does not allow ties;
- allows all alternatives to win; and
- is strategy-proof.

Then the social choice function is a dictatorship.

8 The Future of Voting and Social Choice Theory

The mathematics of voting has been inspired by a trilogy of theorems that we have presented here — Condorcet's Voting Paradox, Arrow's Theorem and the Gibbard-Satterthwaite Theorem. It is surprising, though somewhat disheartening, that social choice theory seems to be based on a bunch of negative results. Does the mathematics suggest that we should give up on democracy? Should we appoint a dictator? Should we give up on society altogether, become hermits and eke out the rest of our lives as turnip farmers? Unless you happen to have a particular penchant for dictators or turnips, then such drastic measures are probably not required. Instead, these results should be taken as a warning that voting is not as simple as it seems and as a mathematical challenge to design and analyze more democratic voting procedures. They also highlight the power of mathematics to uncover beautiful structures and surprising paradoxes in areas where intuition tells us there should be none.

References

[1] K. Arrow, *Social Choice and Individual Values* (Wiley New York 1963).
[2] Geanakoplos, *Three brief proofs of Arrow's impossibility theorem*, Cowles Discussion Paper 1123R, (Cowles Foundation Yale University 1996).
[3] J. Kelly, *Arrow impossibility theorems* (Academic Press New York 1978).
[4] J. Kelly, *Social choice theory: an introduction* (Springer-Verlag New York 1987).
[5] P.J. Reny, *Arrow's Theorem and the Gibbard-Satterthwaite Theorem: A Unified Approach*, Economics Letters (2001), 99–105.
[6] D. Saari, and F. Valognes, *Geometry, Voting and Paradoxes*, Mathematics Magazine **71** (1998), 243–259.
[7] A. Taylor, *Mathematics and Politics: Strategy, Voting, Power and Proof* (Springer-Verlag New York 1995).

Department of Mathematics and Statistics, The University of Melbourne, VIC 3010
E-mail: N.Do@ms.unimelb.edu.au