Parabola Volume 61 Issue 2 (2025), Article 10

A tour of social choice theory[1]

Aaron Manning[2]

1 Introduction

Whether we are electing the next leader of our country, or choosing which movie to watch amongst friends, we see instances of collective decision making throughout our society.

From the origins of democracy until the 18th century, the methods of counting votes in elections were chosen primarily on the basis of intuition. This changed initially due to the work of Nicolas de Condorcet and Jean-Charles de Borda, most notably in [3] and [2] respectively, who are often credited as the founders of a field of study known as *social choice theory*, the formal mathematical study of elections. The research in this area continues today, with some of the greatest contributions arising in the 1950s, 60s and 70s.

One of Cordorcet's contributions which motivated further investigation into this area was the discovery of the so-called *Cordorcet's Paradox* of [3]. Consider the following set of votes, where each voter, represented by a column in our table, ranks the candidates A, B and C in order of their preference.

Ranking	**Ballots**								
1.	A	C	B	A	B	C	B	B	A
2.	B	A	C	B	C	A	C	C	B
3.	C	B	A	C	A	B	A	A	C

Suppose now we just wish to know which candidate is preferred out of A and B, we could do this by ignoring votes for C and producing a set of votes which just includes A and B.

Ranking	**Ballots**								
1.	A	A	B	A	B	A	B	B	A
2.	B	B	A	B	A	B	A	A	B

After counting the votes, we can see that candidate A wins by 5 votes to 4. Likewise, if we choose to consider candidate B against candidate C, ignoring candidate A, we notice that B wins by 7 votes to 2.

[1]This article is based on a recent talk given at the UNSW Mathematics HDR seminars, and is simply a topic of interest to the author.

[2]Aaron Manning is a recent graduate from the pure mathematics honours program at UNSW Sydney.

Ranking	Ballots								
1.	B	C	B	B	B	C	B	B	B
2.	C	B	C	C	C	B	C	C	C

This means that the voting population prefers A to B, and B to C. Hence, of course they prefer A to C, right? Considering a one on one election between A and C leads us to an unfortunate conclusion.

Ranking	Ballots								
1.	A	C	C	A	C	C	C	C	A
2.	C	A	A	C	A	A	A	A	C

C is the preferred candidate to A by the majority of voters! That is, the majority choice behaves in a somewhat irrational and unintuitive way.

The core concern of social choice theory is how we may resolve these so-called paradoxes, and develop a way of counting votes that behaves rationally, and for which we can mathematically prove satisfies many desirable characteristics.

2 Two Alternatives

2.1 Background

From as early as when we attend kindergarten, most of us have an intuition that when choosing between two options, we ought to select the option that the majority of people prefer. This way of counting votes is called the *majority vote*.

Is this really the "best" way to count votes though? Are there any other reasonable alternatives? What trade-offs must we accept if we determine a winner using an alternative method? What should we do when there is an even number of voters?

As this is a mathematics article, here we demonstrate some of the ways in which mathematical methods have been used to shed light on these questions. Of course, in order to reason about something mathematically, we require a mathematical definition which brings our real world understanding of the problem at hand into the world of mathematics. We do this with the notion of a *social choice function*.

Suppose that we conduct an election between two candidates A and B, and receive the following ballot papers.

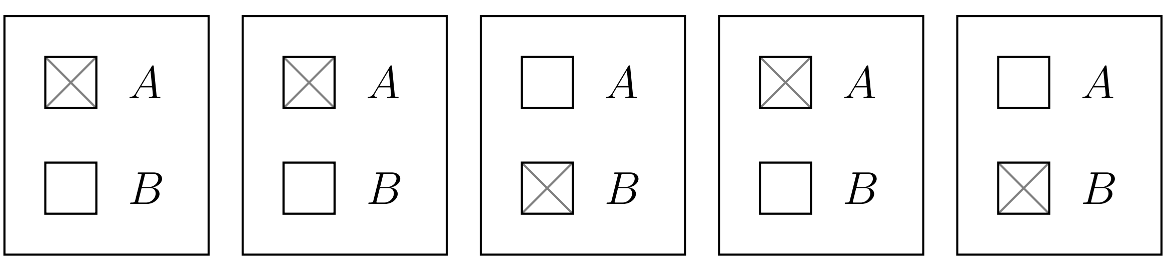

The collection of votes here is called a *voter profile*. In this case, our profile is the tuple (A, A, B, A, B). For the sake of some convenient mathematical notation in the theorems to follow, we choose to identify our two candidates by the integers 1 and -1. Under this labelling, our profile is $(1, 1, -1, 1, -1)$.

An election, or social choice function, is just way of counting votes. This means that it takes as input the voter profile, and returns as output the winner of the election, either 1 or -1. We will also allow the possibility of neither candidate winning, in which case the output will be 0. This will be useful to represent the concept of a tie.

With these ideas in place, we can now formally define what a way of counting votes really is from the mathematical perspective.

Definition 1 (Social Choice Function)**.** *A social choice function between two alternatives is a function $f : \{-1, 1\}^n \to \{-1, 0, 1\}$ for some $n \in \mathbb{Z}^+$.*

Even for relatively modest n, there are a lot of possible choices for functions between these two sets; 3^{2^n} of them, in fact. Most of these do not resemble anything like a sensible way of counting votes. As such, we give a name to a few special social choice functions that we may recognise.

In order to define the first of these, we will use the following additional definition.

Definition 2. *Define the function $\mu : \{-1, 1\} \to \mathbb{Z}^+$ such that, for any vector $\vec{x} \in \{-1, 1\}^n$, $\mu(\vec{x})$ is the number of components of the vector which are 1.*

This function returns the number of voters for the given candidate identified by 1 (candidate A as above). This means our majority vote may be defined as follows.

Definition 3 (Majority Vote)**.** *A social choice function $f : \{-1, 1\}^n \to \{-1, 0, 1\}$ is called* the majority vote function *if*

$$f(\vec{x}) = 1 \iff \mu(\vec{x}) > \frac{n}{2} \quad \textit{and} \quad f(\vec{x}) = 0 \iff \mu(\vec{x}) = \frac{n}{2}\,.$$

The first condition above states that a candidate wins precicely when they recieve the majority of the available votes. No doubt, this is the way most of us instinctively feel that votes ought to be counted when deciding between two options.

We also assign a name to a less than ideal way of counting votes, which we call a *dictatorship*.

Definition 4 (Dictatorship)**.** *A social choice function $f : \{-1, 1\}^n \to \{-1, 0, 1\}$ is called a dictatorship if there exists a voter $j \in \{1, \ldots, n\}$ such that*

$$f\big((x_1, x_2, \ldots, x_n)\big) = x_j\,.$$

This condition $f\big((x_1, x_2, \ldots, x_n)\big) = x_j$ means that we ignore the preferences of all of the voters except voter j. This voter is called the *dictator*, as their ballot paper completely decides the outcome of the election.

2.2 Properties of social choice functions

There are two natural ways to study social choice theory. One option is to design a social choice function, determine what properties that function has, and if these properties are desirable. Another approach, which is the one we shall soon see in action, is to look at all possible social choice functions, request that they exhibit certain features, and observe what functions we have left over. The three properties that we choose are called *anonymity*, *neutrality* and *positive responsiveness*[3].

Definition 5 (Anonymity). *A social choice function f is called* anonymous *if, for all $\vec{x}, \vec{y} \in \{-1, 1\}^n$, $f(\vec{x}) = f(\vec{y})$ whenever $\vec{x}$ is a permutation of $\vec{y}$.*

By permutation, we mean that $\vec{y}$ may be obtained from $\vec{x}$ by changing the order of the elements. For instance, $(1, -1, 1, -1)$ is a permutation of $(1, -1, -1, 1)$ since we may swap the final two elements to obtain one from the other.

Anonymity is also equivalent to the claim that if $\mu(\vec{x}) = \mu(\vec{y})$, then $f(\vec{x}) = f(\vec{y})$, since μ is invariant under permutation. As such, the idea of anonymity is exactly as the name suggests; the social choice function treats all voters anonymously. That is, it doesn't care about *who* voted for a given candidate, just *how many* people did.

Definition 6 (Neutrality). *A social choice function f is called* neutral *if, for all $\vec{x} \in \{-1, 1\}^n$, $-f(\vec{x}) = f(-\vec{x})$.*

The change from $\vec{x}$ to $-\vec{x}$ corresponds with every voter changing their vote; that is, a vote for candidate A is exchanged with a vote for candidate B, and vice versa.

Neutrality encompasses the idea that if every voter changes their vote, then the outcome of the election ought to change, unless of course the outcome is already a tie. In this sense, while anonymity encapsulates the idea that we treat all voters equally, neutrality encapsulates the idea that we treat both candidates fairly.

This condition is important as it means our arbitrary choice of which candidate to label with 1 and which candidate to label with -1 won't influence the outcome of the election.

Definition 7 (Positive Responsiveness). *A social choice function f is called* positively responsive *if $f(\vec{x}) \geq f(\vec{y})$ for any $\vec{x} = (x_1, x_2, \ldots, x_n)$ and $\vec{y} = (y_1, y_2, \ldots, y_n)$ satisfying $x_k \geq y_k$ for all $k \in \{1, \ldots, n\}$.*

Here, $x_k \geq y_k$ tells us that each vote in the profile $\vec{x}$ is at least as favourable for candidate 1 as the corresponding vote in $\vec{y}$. Thus, every voter voting in a way that is no worse for candidate 1 should lead to an outcome of the election which is no worse for candidate 1.

Some authors choose to also include a tie-breaking condition in this definition, where changing a single vote from a tied election changes the outcome of the election according to that vote. This will slightly change how we state the main theorem to follow, but the idea is mostly the same.

[3] This condition is sometimes referred to as *monotonicity*.

2.3 May's Theorem

This naturally leads to the following question: which social choice functions are neutral, anonymous and positively responsive. Quick inspection of the relevant definitions shows that the majority vote indeed satisfies these properties, and being the most obvious way of counting votes, this is somewhat reassuring. However, is it the only way to count votes which satisfies these properties, or do we have other equally reasonable choices for how to count votes? This question is answered by May's Theorem, the main result of [6] (stated slightly differently here).

Theorem 8 (May's Theorem)**.** *Suppose that $f : \{-1,1\}^n \to \{-1,0,1\}$ is an anonymous, neutral and positively responsive social choice function.*

(a) If f is chosen to minimise the number of possible ties, then f is the majority vote.

(b) If n is odd and f cannot result in a tie, then f must be the majority vote.

(c) If n is even and f cannot result in a tie, then no such f exists.

While we will not present a proof of this theorem here, writing such a proof is not too difficult. For the interested reader, we provide some guidance through the following exercises.

Exercise 9. *Show that the majority vote is anonymous, neutral and positively responsive.*

Exercise 10. *Prove that for all $\vec{x} \in \{-1,1\}$,*

(a) $\mu(\vec{x}) + \mu(-\vec{x}) = n\,;$

(b) $\mu(\vec{x}) > \frac{n}{2} \iff \mu(-\vec{x}) < \frac{n}{2}\,;$

(c) $\mu(\vec{x}) > \mu(-\vec{x}) \iff \mu(\vec{x}) > \frac{n}{2}$ *and* $\mu(\vec{x}) < \mu(-\vec{x}) \iff \mu(\vec{x}) < \frac{n}{2}\,;$

(d) $\mu(\vec{x}) = \frac{1}{2}\sum_{i=1}^{n} x_i + \frac{n}{2}\,.$

Exercise 11. *Suppose that $f : \{-1,1\}^n \to \{-1,0,1\}$ is an anonymous and positively responsive social choice function. If $\mu(\vec{x}) \geq \mu(\vec{y})$ for some $\vec{x}, \vec{y} \in \{-1,1\}^n$, then $f(\vec{x}) \geq f(\vec{y})$.*

Exercise 12. *Suppose that $f : \{-1,1\}^n \to \{-1,0,1\}$ is an anonymous, neutral and positively responsive social choice function. For all $\vec{x} \in \{-1,1\}^n$,*

(a) if $\mu(\vec{x}) > \frac{n}{2}$, then $f(\vec{x}) \in \{0,1\}$;

(b) if $\mu(\vec{x}) < \frac{n}{2}$, then $f(\vec{x}) \in \{-1,0\}$;

(c) if $\mu(\vec{x}) = \frac{n}{2}$, then $f(\vec{x}) = 0$.

Exercise 13. *Prove Theorem 8.*

3 Social welfare with three or more alternatives

We now turn our attention to the situation where there are more than two candidates to choose from. In this scenario, the question of what the voter profile should look like is a little more complex, where we could allow our voters to select one candidate only, provide an ordering of the candidates, or possibly select some subset of the candidates with which they approve. Any of these scenarios could be studied separately; however, we will first consider what are called *social welfare functions*, which are functions of the form

$$f : S_m^n \to S_m \,.$$

The notation S_m comes from group theory, and refers to what is called the *symmetric group* on m symbols. For our purposes, elements of S_m are complete rankings of m candidates. It may be thought of in mathematical terms as a function that takes in the candidates $\{1, \ldots, m\}$ and, for each of them, assigns a unique number from $\{1, \ldots, m\}$ which is its corresponding ranking, or as a total order $\prec$ on the set $\{1, \ldots, m\}$.

As with May's Theorem, let's now impose some reasonable constraints on a social welfare function, and see what functions satisfy these properties. We begin with what is called *unanimity*.

Definition 14 (Unanimity). *A social welfare function $f : S_m^n \to S_m$ is called unanimous if, for all $x \in S_m$,*

$$f\big((x, x, \ldots, x)\big) = x \,.$$

If everyone agrees on the complete ordering of their preferences, then it doesn't make sense for the outcome of the election to be anything other than the ordering with which everyone agrees.

The other condition that we will impose is called *independence of irrelevant alternatives*.

Definition 15 (Independence of Irrelevant Alternatives). *Suppose that $f : S_m^n \to S_m$ is a social welfare function and that $A, B \in \{1, \ldots, m\}$ are candidates. Suppose furthermore that $\vec{x}, \vec{y} \in S_m^n$ are voter profiles where the relative ordering of candidates A and B is the same for all voters $j \in \{1, \ldots, n\}$. This means that, if voter j ranks A higher than B in the profile $\vec{x}$, then likewise voter j ranks A higher than B in profile $\vec{y}$.*

Then f is said to satisfy the independence of irrelevant alternatives *if, for all such profiles and candidates, the relative ordering of A and B is the same in $f(\vec{x})$ and $f(\vec{y})$.*

This definition is a little trickier than the those we considered for the two voter case, so let's consider an example to illustrate what this means. Suppose that we have the following voter profile $\vec{x}$, with each column representing a voter and their corresponding ballot paper, and the ranked preference order going down from top to bottom.

Ranking				**Ballots**			
1.	A	A	C	B	A	C	A
2.	B	C	A	C	C	B	C
3.	C	B	B	A	B	A	B

Consider also the profile $\vec{y}$, given below.

Ranking	Ballots						
1.	A	C	A	B	C	B	C
2.	C	A	C	A	A	A	A
3.	B	B	B	C	B	C	B

Notice that although the votes change, if we compare the ordering of A and B in corresponding votes, it is always the same. That is, if A is ranked higher than B in the first profile, this is also true in the second profile, and vice versa.

The independence of irrelevant alternatives tells us that when this occurs, the relative ordering of A and B in the output also shouldn't change. That is, if A comes out ahead of B in an election with voter profile $\vec{x}$, then A should likewise come out ahead of B in voter profile $\vec{y}$, and vice versa. What this represents is that the existence of additional options shouldn't change how a society feels about the existing candidates.

As it turns out, the independence of irrelevant alternatives is actually a very strong property. Accordingly, we have the following result of [1], which describes what happens when we impose only the conditions above.

Theorem 16 (Arrow's Theorem). *Suppose that $m \geq 3$. If $f : S_m^n \to S_m$ is a social welfare function which is unanimous and satisfies the independence of irrelevant alternatives, then f is a dictatorship.*

Even though we only ask our social welfare function to satisfy two very reasonable criteria, it turns out that our only choices which satisfy these criteria are dictatorship functions. Thus, if we are to conduct an election with a social welfare function, then we have to give up on either unanimity, independence of irrelevant alternatives, or the property of not being a dictatorship.

As with May's Theorem, while we will not prove this result here, we provide a little guidance for any reader that wishes to fill in the details themself.

(a) Begin by proving that, if every voter prefers candidate A over candidate B, then the social ranking prefers A over B.

(b) Prove that, if every voter ranks candidate A at either the top or bottom of their voter profile, then so does the social ranking.

(c) Prove that there exists a voter profile where a given candidate A is at the bottom of the social ranking, and there exists a single voter $j \in \{1, \ldots, n\}$ which can change their preference to move A to the top of the social ranking.

(d) Prove that voter j decides the order between any two candidates, neither of which is A.

(e) Prove that voter j decides the order between A and any other candidate.

The final two parts of the above lead to the conclusion that voter j is the dictator. The idea for this proof is due to [4] and is not the technique used in Arrow's original proof.

4 Social choice with three or more alternatives

When we started considering elections with three or more alternatives, we asked for something quite strong, which is a complete ranking of the alternatives as the output. Arrow's Theorem tells us that we can't have such a function which is unanimous, satisfies the independence of irrelevant alternatives, and is not a dictatorship. However, we could possibly recover something by only requesting that our election determines one elected candidate, rather than ranking all of them. That is, what happens if we instead consider functions of the form

$$f : S_m^n \to S$$

where $S = \{1, \ldots, m\}$. This is a more general notion of a social choice function that what we considered with the two candidate case.

One well documented issue with many electoral systems is their susceptibility to what is called *strategic voting*. Consider for instance the following results in a *plurality vote* (also known as *first past the post*), which allows voters to choose any one candidate, and the winner is the candidate that receives the most votes.

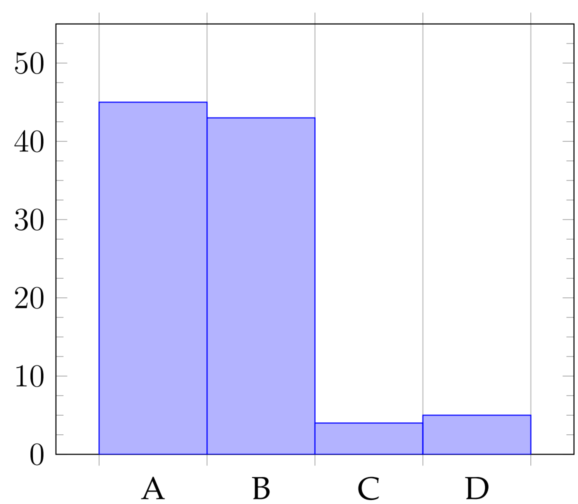

As a voter who prefers candidate C, if you saw polling data that looked like this, then you may feel somewhat discouraged. Why vote for candidate C when they have such a small chance of winning? If you have any preference between candidates A and B, then you may as well use your vote to express this preference, rather than letting it go to waste. This action is an instance of what is called *strategic voting*, where a voter may vote against their true preference to get an outcome that they prefer. More precisely, we have the following definition.

Definition 17 (Strategy Proofness). *A social choice function $f : S_m^n \to S$ is called* manipulable *if there exists a voter profile $\vec{x} = (x_1, \ldots, x_n) \in S_m^n$ and voter $j \in \{1, \ldots, n\}$ who can change their preferences to x_j' to get an outcome they prefer (as defined by x_j).*

A social choice function which is not manipulable is called strategy proof.

Ideally, we would hope for our elections to be strategy proof, so that everyone can turn up to the polling booth and express their preferences, without having to think more deeply about how others may be voting, or read about polling data ahead of time to determine their vote.

However, as it turns out, strategy proofness is once again one of these unexpectedly strong properties. Accordingly, we have the following theorem proven independently in [5] and [7].

Theorem 18 (Gibbard-Satterthwaite Theorem). *Suppose that $f : S_m^n \to S$ is a social choice function satisfying*

(a) $\#f(S_m^n) \geq 3$*;*

(b) f *is strategy proof.*

Then, f *must be a dictatorship.*

We require the condition $\#f(S_m^n) \geq 3$; that is, that at least three candidates have the possibility of winning, to prevent everything from collapsing into the two candidate case with the majority vote as per May's Theorem.

This leads us once again to an unfortunate conclusion. Strategy proofness is such a strong assumption that we cannot hope for our social choice function to hold this property, unless we accept an obviously undesirable way of counting votes.

5 Conclusion

One may despair at some of the conclusions we have reached within this article, particularly the Gibbard-Satterthwaite Theorem. However, strategy proofness need not be the be-all and end-all of voting. For instance, we could instead study manipulability from a probabilistic setting, where we try to minimise the probability of strategic voting without compromising on other desirable properties. There are also other ways we may weaken our assumptions to recover a somewhat reasonable mechanism of voting, such as allowing the possibility of ties which we excluded by definition in our three or more candidate social choice functions.

These impossibility theorems and paradoxes frequently tell us that our ideal voting system may not exist, but that doesn't mean that a reasonable and workable system doesn't, and there is plenty more to study to understand the other pitfalls of social choice, and where we can find acceptable compromises.

Beyond what was covered in this article, there are many other paradoxes and puzzles of the mathematical theory of voting, and the interested reader may want to read the *best-is-worst paradox*, *cloning paradox* and the *centre squeeze*.

References

[1] K.J. Arrow, Social Choice and Individual Values, *Wiley*, New York, 1951.

[2] J. de Borda, Mémoire sur les élections au scrutin, *Histoire de l'Académie Royale des Sciences* **102** (1781), 657–665.

[3] M. de Condorcet, Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix, *Imprimerie Royale*, Paris, 1785.

[4] J. Geanakoplos, Three brief proofs of Arrow's impossibility theorem, *Economic Theory* **26** (2005), 211–215.

[5] A. Gibbard, Manipulation of voting schemes: a general result, *Econometrica* **41** (1973), 587–601.

[6] K.O. May, A set of independent necessary and sufficient conditions for simple majority decisions, *Econometrica* **20** (1952), 680–684.

[7] M.A. Satterthwaite, Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions, *Journal of Economic Theory* **10** (1975), 187–217.