GRESHAM COLLEGE

12TH NOVEMBER 2019

MATHS AND VOTING

PROFESSOR CHRIS BUDD OBE

Introduction

As I write this lecture we live in interesting times for politics. This won't be a lecture about Brexit in any way. But whatever Brexit has done for good or bad, it has certainly raised all of our awareness of the parliamentary process. We live in a democracy (for which I am very grateful) and the basis for our parliamentary system is the voting process. The purpose of voting is to allow the general public, to express their views, either directly as in a referendum, or through electing representatives, as in a general election. What could be simpler? However, the process of translating the 'will of the people' into a fair representation of that will in parliament, is very far from perfect. There are multitude types of voting process, all of which try to do this as well as possible. However, as we will see, all of these processes have faults, and in fact there is simply no 'perfect' voting process. It is hard enough, for example, to decide (as we do in a referendum) on a single issue if we only have one choice to make. However, things start to get complicated when we have to judge between multiple options with one vote. Then things get more complicated still when we can express multiple preferences for multiple options. At this point it is very hard to know what the best voting system at all is. In fact, as we will show from the start, there is, in a sense, no perfect voting system at all. It follows that every voting system must therefore be some sort of compromise between different choices of what might be conceived as being fair. In addition to this are various harder to define concepts such as the need to produce decisive government, and also to provide a clear link between the government itself and the people who have elected that government. Of course, voting isn't just used to elect a government, we also see it in action on the Internet, sport, selecting a candidate for a job, and in Strictly Come Dancing.

Given this complexity, it turns out that voting is not only very complex but is best understood in terms of mathematics. Indeed, many mathematicians have worked on producing different voting systems. One of these systems was derived by Charles Dodgson, also known as Lewis Carroll, and I will describe this in some detail.

In this lecture we will take a look at the different ways that we can take vote and the pros and cons of each. We will also have a look at the way opinion polls work as a way of predicting the result of an election, and ask the question of how reliable they are? Rather than apply our newfound knowledge to anything as trivial as Brexit, we will look at the far more serious business of the voting patterns in the Eurovision Song Contest. Hopefully we won't get Null Points as a result.

Why no voting system can ever be perfect.
Perhaps the key contribution to our understanding of voting was made in the 1950s by the economist Kenneth Arrow (see below), who then won the 1972 Nobel Prize in Economics for his work in this field. Arrow's seminal contribution was to show that there is **no perfect voting system**.

He did this by considering what might be reasonable expectations of any voting system, and then showing that there was no one system which could satisfy all of them. Arrow assumed that there was an election in which there were multiple candidates, and that each voter could express a different range of opinions on each candidate. For example, if there were candidates A, B, C and D, then each voter could express a preferential view on each candidate. For example, Voter number one could rank than as A, B, C, D in order of preference, whilst another might rank them as B, D, C, A. These votes would then be expressed anonymously at a ballot box in an election. The purpose of the voting system would then be to elect a single candidate as the output of the voting process. The question is, can this be done fairly.

Arrow's original conditions were as follows:

1. **(Dictator)** The system should reflect the wishes of more than one voter, so there is no dictator.

2. **(Unanimity)** If all voters prefer candidate A to candidate B then A should come out ahead of B in the final vote.

3. **(Universality)** The voting system should always return one clear result.

4. **(Independence of irrelevant alternatives)** In the final result, the ranking of A above B should only depend on how individual votes ranked A compared to B and not how they ranked them when compared to a third alternative C.

Other conditions can also be added; indeed, many have been since Arrow's work.

One is the **majority condition (M)**. In this a candidate who is top choice for a majority of the voters should get elected.

Another is the **monotonicity criterion**. A voting system is monotonic if it is neither possible to prevent the election of a candidate by ranking them higher on some of the ballots, nor is it possible to elect an otherwise unelected candidate by ranking them lower on some of the ballots.

There is also a final condition, that it the one of **anonymity (A)** so that it should not be possible to tell from the vote, who voted for which candidate.

All of these conditions seem pretty reasonable. But now here is the bad news, which is Arrow's key conclusion and for which he won the Nobel Prize and is described in [1]. See also [2] for a very readable account of voting procedures and the relevance of Arrow's conditions.

> It is mathematically impossible for a voting system involving three or more candidates to satisfy Arrow's conditions above.

Oh dear! Does this mean the end of democracy? Well not quite. What is shows is that the conditions laid down above are really very demanding, and that in a realistic voting method we have to make certain compromises, balancing different views of what it means to be fair. These inevitably lead to anomalies in the voting process, and we have to work these out and then be aware of them. Therein lies the science (and mathematics) of voting.

To illustrate all of this, let's see how well a simple voting system fairs against these conditions.

A commonly used system, which is often used when a committee ranks candidates for a job, is the **Borda system**. The Borda system is named after the 18th-century French mathematician Jean-Charles de Borda who devised the system in 1770, and it has been (re)discovered many times both before and since. It is very simple to use and has certain advantages over several other voting methods. In the simple Borda system each voter ranks the candidates in order. If there were N candidates, then they would then give a number between N-1 (best) and 0 (worst) for each candidate (and they are allowed to give the same ranking to two different candidates). Having done this the votes for each candidate are added up. The candidate with the largest number of votes then wins. Variants of the Borda system use different numbers for the rankings such as 1, ½, 1/3 etc… In Strictly Come Dancing the judges use a version of the Borda method where they give a mark between 1 and 10 for each dance pair.

Let's see how the Borda system works with three candidates A, B, C and three voters V1, V2, V3.

	A	B	C
V1	2	1	0
V2	2	0	1
V3	1	0	2
Total	5	1	3

Adding these up we see that

A gets 5 B gets 1 C gets 3

So, in this system A is the clear winner. We also note that A is the first choice for two out of the three voters.

Of course, we might instead have something like

	A	B	C
V1	2	1	0
V2	0	2	1
V3	2	1	0
Total	4	4	1

In this case A and B tie for first place with 4 votes each (even though A is the first choice for 2 out of the three voters). If in such a case the system delivers the same grade to the two top candidates, then the 'election' is repeated. Hopefully on the second run a winner emerges. Of course, this is easily done in a committee (which is one reason that this system is used in this case) but is much harder to do in a real election. Another problem with a real election is that it requires votes to be added up. Again, this is easy for a small election such as Strictly Come Dancing, but much harder if you are considering tens of thousands of voters.

But apart from these practical considerations, how well does the Borda system measure up against Arrow's conditions. It is clearly anonymous, which is a good start. It is also monotonic as increasing the vote on any one candidate can only increase their score overall.

Dictator the Borda system is a consensus system in that it tends to elect candidates who are broadly supported by the electorate. However, if there are only a few voters it is possible for one to be a dictator by giving an otherwise popular candidate a very low score or vice versa.

Unanimity is satisfied. If all voters prefer A to B, then A will always get a higher ranking than B. It follows the sum of the rankings of A will be higher than the sum of the rankings of B. Or mathematically if $X > Y$ and $W > Z$ then $X+W > Y + Z$.

Universality fails. We have seen this in the election above where A and B tie for first place. Rerunning the vote may well sort this out, but this cannot be guaranteed.

Independence also fails. This is more subtle and is a weakness of the Borda system which can be exploited in tactical voting and has meant that it is not widely used.

The *majority rule* also fails.

We can see both of the latter cases in the following example. Imagine that we have four candidates A, B, C, D and five voters V1, V2, V3, V4, V5.

The votes are recorded in the following table

	A	B	C	D
V1	3	2	1	0
V2	3	2	1	0
V3	3	2	1	0
V4	0	2	3	0
V5	0	2	0	3
Total	9	10	6	3

We see in this table that B is the winner of the election, and that A has come second.

However, this is a bit odd. In the election A is the first choice for a majority of 3 voters. Also, in this election, A beats B on *the majority of the occasions*. Thus, the election has failed the majority test. It is also clear that although B has won the election overall, they have *never been the first choice of any voter*. They can best be described as the completely safe candidate, who is never going to be exceptional. B is a safe pair of hands, but no super star.

This election has also failed the independence test. In this vote it has happened that both V4 and V5 really dislike A (wrong colour hair). As a result, they have deliberately put A at the bottom of the list and have given C and D high marks. These 'irrelevant alternatives' have lost A the election over B. This is a clear example of tactical voting.

As a similar example, we consider the modified Borda method used in Strictly Come Dancing. Suppose that the four judges are voting for two dancers. In the first case they give marks of 9,9,9 and 1. In the second case they give marks of 8,8,8,8.

Adding up the marks we see that the first dancer has 28 points and the second 32. So that the second dancer is the clear winner. However, the first dancer is the first choice for three out of the four judges. So, we have a problem. The last judge has in fact become a dictator, swinging the vote in their favour by a very poor mark.

Given these problems we can see that it is hard to find a voting method, which even comes close to satisfying reasonable assumptions. So, what are we to do? We will now look at a couple of widely used, but highly methods, before looking at some better systems more heavily based on mathematical reasoning.

First past the post

Single constituency voting

This is the simplest of all systems. In its purest form the voters are presented with a list of candidates or policies. They express a **single preference** (including abstaining) and the candidate or policy with most votes wins. This is used in each parliamentary constituency in a UK general election. It is also used in voting in the UK parliament in a division of the assembly. In this case, when the division bell is rung, the MPs express their vote by going through one of two doors (the Ayes and the Nays), where the votes are counted by Tellers.

Credit: ©UK Parliament/Jessica Taylor/ Stephen Pike CC BY 3.0

A very similar process was used in the Brexit referendum, where the voters were asked simply whether they wanted to leave or remain.

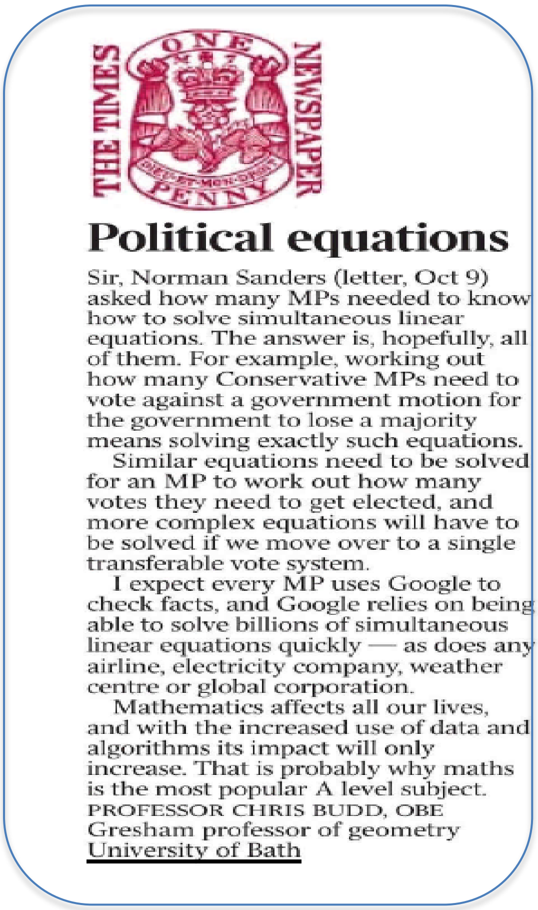

THE TIMES NEWSPAPER

Political equations

Sir, Norman Sanders (letter, Oct 9) asked how many MPs needed to know how to solve simultaneous linear equations. The answer is, hopefully, all of them. For example, working out how many Conservative MPs need to vote against a government motion for the government to lose a majority means solving exactly such equations.

Similar equations need to be solved for an MP to work out how many votes they need to get elected, and more complex equations will have to be solved if we move over to a single transferable vote system.

I expect every MP uses Google to check facts, and Google relies on being able to solve billions of simultaneous linear equations quickly — as does any airline, electricity company, weather centre or global corporation.

Mathematics affects all our lives, and with the increased use of data and algorithms its impact will only increase. That is probably why maths is the most popular A level subject.

PROFESSOR CHRIS BUDD, OBE
Gresham professor of geometry
University of Bath

On the left you can see a letter which I wrote to the Times in response to a correspondent who had stated in a letter the day before that mathematics was of no great practical use and challenged the Times readers to think of any time that an MP needed to solve a set of simultaneous linear equations.

My response, as you can see, is that to operate a first past the post system an MP needs to solve such an equation every time that they vote.

For example, consider a vote in an imaginary parliament with two parties: A with 310 members and B with 250 members. You are in the larger party and know that 10 members are likely to abstain. Assuming that all members of the other party will vote against the motion, how many members of the larger party will have to vote for a motion to ensure that it passes?

We will solve this using algebra and two simultaneous linear equations. Let x be the number of members of A that vote for the motion and y that vote against. It follows from considering the party members that

$$x + y + 10 = 310$$

If the motion is to pass by one vote then

$$x = 250 + y + 1$$

Adding these two simultaneous linear equations we get:

$$2x + y + 10 = 561 + y$$

Hence, subtracting $y + 10$ from both sides we get

$2x = 551$ so that $x = 275.5$

Thus at least 276 members of Party A must vote to ensure that the motion passes

The first past the post method has been used for a long time, and we are very used to it. It has the significant merits of being easy and transparent to use, and of producing a clear answer. However even in this simple case it is severely flawed.

One obvious problem is that is an all or nothing method. Voters are asked to make a clear choice between several candidates and are unable to express any form of preference between them. Later in this lecture we will look at a variant called the IRV method (or AV), which modifies first past the post to allow for preferences.

A second, and related, problem is that similar candidates can split the vote. I experienced this directly during a vote conducted in the 1990s in Bristol about the future funding of education. The voters were asked to give a single vote to one of three options:

1. Increase Secondary School funding,
2. Increase Primary School funding,
3. Give no extra funding to schools.

As I recall the voting went something like: 25% for 1, 35% for 2, and 40% for 3.

On the basis of the Bristol Council concluded that the winner was 3 or 'no extra funding'. And that is what they did. This was despite the 'obvious' fact that 60% of the people in Bristol had in fact voted for an increase in funding for some form of education. A preferential vote would have avoided this.

A third problem is the possibility of error. Crucial decisions can be decided on the basis of a single vote majority. This is fair if everyone votes, but what happens if they don't. As an example, in the 2016 Brexit referendum the result was 52%-48% in favour of Leave. On this basis, the statement that "the majority of the UK chose to leave the EU" has been made. If all of the voters had voted then that would be true, but this did not actually happen. The turnout rate was about 72%, which means that only 37% of the eligible British electorate actually chose to vote Leave. Any statement about the whole population made on the basis of such a limited sample is at best uncertain. The question must arise as to whether a 52% majority on a 72% sample is strong evidence for a greater than 50% majority from the whole sample. This is a subtle question in statistics which is a matter of hot debate, see [3] for example. This uncertainty is a good reason for insisting on a margin of error in such referendums such as 60:40 before a decision based on that referendum is made.

Multiple constituency voting

The problems with first-past-the-post are exacerbated when there are multiple constituencies. This is the method used in both the UK and US elections, in which the voters in each constituency vote for a representative from a political party (e.g. an MP) using the first past the post system. The party with the largest number of representatives (MPs) then wins the election overall.

This method has two big advantages. Firstly, it gives a clear result and (generally) leads to a majority government. Secondly it clearly links the voters to their representative. Neither advantage is to be disregarded lightly.

However, it is easy to see theoretically why this system is flawed. Imagine that we have three constituencies with respectively 10,000, 3 and 5 voters and two parties A and B. In the first constituency everyone votes for A. In the other one everyone votes for B. Thus, party A gets one MP and 10,000 votes, whereas the other gets two MPs and 3 votes. Hence party B wins the election even though it got far fewer votes than party A.

This problem also arises in practice. In the last US presidential election, Donald Trump and the Republicans got 304 electoral votes and 46.1% of the vote, whereas Hilary Clinton and the Democrats got 227 electoral votes and 48.2% of the vote. On this basis the Republicans won, even though the Democrats won the 'popular vote'.

As an even more extreme example, suppose that there are three parties A, B, C and three constituencies with the votes recorded as follows

	A	B	C
Constituency 1	10,000	9,000	1,000
Constituency 2	10,000	9,000	1,000
Constituency 3	1,000	9,000	10,000

In this election party A wins two seats with 21,000 votes, and party C wins one seat with 12,000 votes. Thus, A wins the election. However, party B has won no seats at all, despite coming second in all of the constituencies and getting, at 27,000, the greatest number of votes.

It clearly matters hugely in first past the post elections exactly where your voters are located. Indeed, small shifts in the boundaries between constituencies can have big effects in the results of an election. For example, moving a group of voters from a safe set to a marginal will not affect the safe seat, but could swing the marginal. The practice of *gerrymandering* arises when the process of moving boundaries is exploited by unscrupulous individuals with the aim of winning elections. It is named after the governor of Massachusetts, Elbridge Gerry, who two centuries ago approved of a new and oddly-shaped voting district (illustrated) that was long, thin and curvy, with the explicit outcome of (successfully) trying to rig the election.

A lot of mathematics and statistics is used to determine whether there is evidence of gerrymandering in practice. See [4] for a description of this in the US elections.

It would seem obvious that a better means of voting is needed!

Proportional representation

Proportional representation (PR) is widely held to be a fairer voting system than first past the post in the case of each voter having one preference. PR is used in many places, most notably in the elections to the European Parliament. In the simplest form of this the number of representatives of each party is taken to be directly proportional to the number of votes that they get in an election in which the voters express only one preference. Proportional representation systems aim to allocate seats to parties approximately in proportion to the number of votes received. For example, if a party wins one-third of the votes then it should gain one-third of the seats. Simple!

Sadly not. There are a number of reasons which make a purely proportional system difficult to use in practice. Most notably amongst these is that usually these divisions produce fractional numbers of seats. For example, if there are ten seats and each party get a third of the votes then the best thing we could do would be to divide the seats up as 3,3,4. But why should one party get more than the others? Another problem is that some parties, or countries in the European elections, may get no representation at all. This is avoided by giving each country a minimum number of seats, which again breaks pure proportionality.

A number of methods have been devised to deal with the issue of fractional seats by trying to minimise different kinds of disproportionality. Widely used methods are the d'Hondt method and the Webster/Sainte-Laguë method. The former minimises the number of votes that need to be left aside so that the remaining votes are represented exactly proportionally and slightly favours large parties and coalitions over scattered small parties. In comparison, the latter reduces the reward to large parties, and it generally has benefited middle-size parties at the expense of both large and small parties.

The d'Hondt method works by getting parties to 'buy seats' with their votes until there are no seats left. The price of a seat starts too high and is gradually reduced as the seat allocation continues. It works as follows. After all the votes have been counted the party with the largest proportion wins **one seat.** Then the number of votes V for each party is divided by the number of seats s it has plus one to give the number N (which is how much it can afford for a seat) where:

$$N = V/(s + 1).$$

The second seat is then given to the party with the largest value of N. The value of N for that party is then adjusted, and the party with the next highest value of N gets the next seat. This process then continues until all of the seats have been allocated. Exactly this method was used in the European elections on May 22, 2014. It is simple to operate this method. Indeed, here is a simple Matlab code, which does the whole calculation.

```
s = 0*[1:partynum];
      for k = 1:seatnumber
            N=V./(1 + s);
            [m,j] = max(N);
            s(j) = s(j)+1;
      end
```

As an example if you have 7 seats to allocate to 4 parties, A,B,C,D with votes

A = 100 000, B = 80 000, C = 30 000, D = 20 000

Round 1: A gets one seat
Round 2: N(A) = 50 000, N(B) = 80 000, N(C) = 30 000, N(D) = 20 000: B gets one seat
Round 3: N(A) = 50 000, N(B) = 40 000, N(C) = 30 000, N(D) = 20 000: A gets one seat
Round 4: N(A) = 33 333, N(B) = 40 000, N(C) = 30 000, N(D) = 20 000: B gets one seat
Round 5: N(A) = 33 333, N(B) = 26 666, N(C) = 30 000, N(D) = 20 000: A gets one seat
Round 6: N(A) = 25 000, N(B) = 26 666, N(C) = 30 000, N(D) = 20 000: C gets one seat
Round 7: N(A) = 25 000, N(B) = 26 666, N(C) = 15 000, N(D) = 20 000: B gets one seat

So the seat allocation is **A:3, B:3, C:1, D:0**

The vote proportion is: 43%, 35%, 13%, 9%, and the seat proportion is 43%,43%,14%,0%. So, B has done rather well, and D has done badly.

The role of mathematics in voting was brought into sharp relief with the recent EU elections for 751 seats amongst 28 member states (with concerns about the effect of Brexit on the voting process). Because of the wide disparity of size of the states, since the 2009 Lisbon Treaty, the larger states have agreed to be to be underrepresented so that the smaller states can be better represented. The treaty also requires that no country should have fewer than 6, or more than 96 seats. All of this has to be taken into account when dividing up the votes proportionally. To address this a committee of mathematicians from around Europe was formed, chaired by Geoffrey Grimmett at Cambridge. The committee came up with two mathematically based solutions, the *Cambridge Compromise* and the *Power Compromise.* Sadly, despite the transparency of the process, the mathematically based methods were rejected by the European Constitutional Affairs Committee (AFCO). You can learn more details about why there is no maths for Europe in [5].

Condorcet methods

Having looked at two practical and simple, but flawed, methods, we will now have a look at more complex methods. These are more firmly based on mathematical principles with the aim of producing a fairer election.

Condorcet voting methods are named for the 18th-century French mathematician Marie Jean Antoine Nicolas Caritat, the Marquis de Condorcet (1743-1794) (illustrated below) who championed what he (and many others) saw as the best possible voting systems. In a pure Condorcet method, the choices of each voter are compared against everyone else in a series of 'tournaments'. If one candidate wins all of the tournaments, then they win overall.

Condorcet methods are in a sense the best possible voting system in elections for which there are many candidates and many voters, who all express views on the candidates. Let's suppose that each of the voters ranks the candidates in order of preference. Imagine that we have candidate A and candidate B. Ten of the voters place A ahead of B, and 6 place B ahead of A. Thus, in a head to head competition between A and B we can see that A wins over B. Now, let's imagine that one candidate wins all of the head to head competitions. We would not unreasonably think that that candidate had won. This candidate is then the *Condorcet winner*. If a Condorcet winner exists, then we would like it if our voting method selected them. A voting system which does this, is called a *Condorcet method*. It seems intuitively correct that such a voting method is fair.

A few examples will help to see what is going on. We will consider an election with three candidates and 30 voters. The preferences shown by the voters are as follows:

Number of voters	Preferences
10	A > B > C
1	A > C > B
5	C > A > B
0	C > B > A
9	B > C > A
5	B > A > C

In this election:

A beats B on 16 occasions and B beats A on 14, so A is 2 ahead of B on a head to head.

A beats C on 16 occasions and C beats A on 14, so again A is 2 ahead of C on a head to head

B beats C on 24 occasions and C beats B on 6, so B is 18 ahead of C on a head to head

This is illustrated in the graph below.

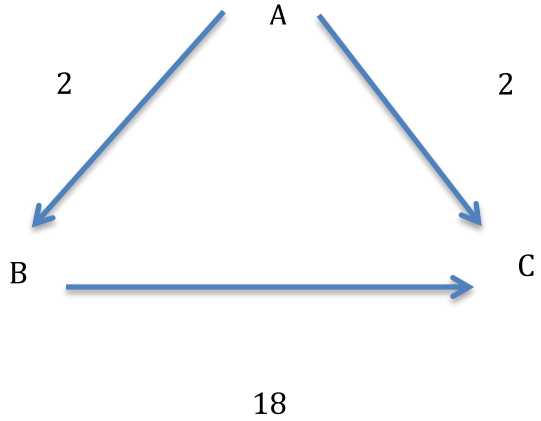

We can see that A is the Condorcet Winner. Hooray.

A first past the post method would give A 11 votes, B 14 votes and C 5 votes. So this would elect B. This is therefore not a Concordet Method

Earlier we looked at the Borda method. We can ask whether this will find the Condorcet winner. If we rank as before then we get the following table:

Number of voters	A	B	C
10	2	1	0
1	2	0	1
5	1	0	2
0	0	1	2
9	0	2	1
5	1	2	0
Total	32	38	20

This the Borda voting system also elects B (by some margin) over A. It has not found the Condorcet winner and is therefore also not a Condorcet method.

The procedure of finding a Condorcet Winner is often viewed as a gold standard in voting. However, things are not as simple as they might seem. There may be a case when in a series of votes there is always a clear winner between two options, however in the round there is no overall winner. A classic example of this is the game of rock-scissors-paper

In this game, rock beats scissors, paper beats rock, and scissors beats paper. So who wins? Well no one does. This is called a *cyclic situation* and there is no Condorcet winner. It is completely symmetric, so there is no way to distinguish between any of the 'candidates'.

A more complex example is given in the following table:

Number of votes	Preferences
3	A > B > C > D
1	D > B > A > C
1	D > C > A > B
1	C > D > B > A
1	B > D > C > A

The graph showing the results of head to head contests is given below. In this the arrow shows, for example, that B has defeated C in a head to head by a margin of 3.

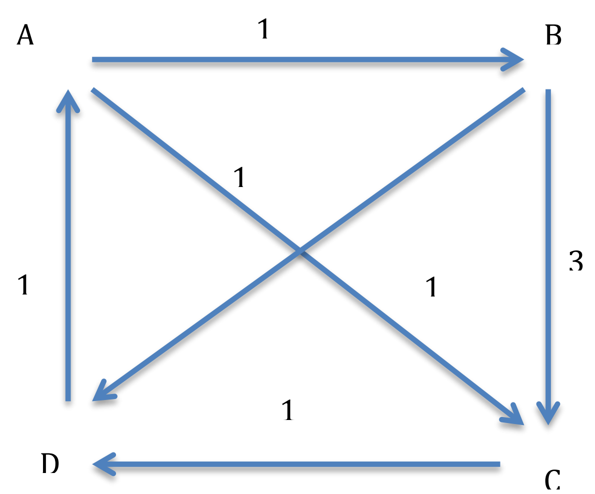

In this case A and B have both won two head to head contests, and C and D have each won one. So, there is no Condorcet winner.

It is easy to see that in a first past the post election, Party A would get 3 votes, D would get 2 votes and B, C would get one vote, so A would win overall. A longer calculation shows that in a Borda vote B would just win (with 12 points) over A (with 11 points). So, who has won the election?

A number of voting strategies have been devised which can deal with the cyclic cases and deliver a winner, which will also give the Condorcet winner if one actually exists. Perhaps the earliest of these was devised in 1299 by Ramon Llull and has been reinvented as **Copeland's method**. This method elects the candidate who wins the most head to head contests. In our example above, A and B tie for first place in Copeland's method. This is not unusual and means that the method is not widely used in practice, although a version of it is used in Premier League football.

A much more sophisticated procedure is **Schulze's method** which was invented in 1997 [6]. This is now quite widely used. In the Schulze method we draw the same graph as above, but in the case of (say) A defeating B we put the number of times they have won (in this case 4) on the arrow

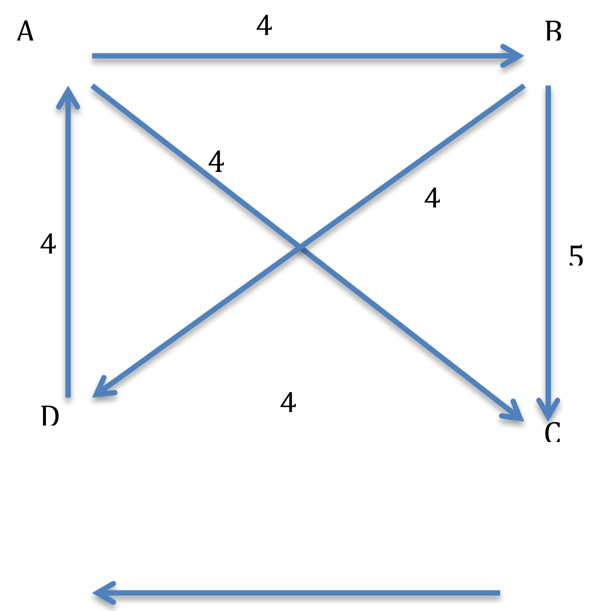

From each candidate A, B, C, D we then find a path from them to the other candidates e.g. from B to C. Usually there will be several such paths. The weakest link on that path is the smallest number on the graph above i.e. the smallest number of contests that has to be one. The strongest path is the one which has the least weak link. The value of this is the strength of the path. As an example, there is a path from B to C direct of strength 5, and a path BDC of strength 4. So, the strength of the path from B to C is defined by P (B, C) is 5. The strengths of the paths are shown below

	Path to A	B	C	D
Path from A	x	4	4	4
B	4	x	5	4
C	4	4	x	4
D	4	4	4	x

In this example we see that P (B, C) = 5 and all other paths have strength 4.

The winner of a Schulze election is the candidate X so that P (X, Y) > P (Y, X) for all possible Y. In other words, it takes more voters to fancy X over Y than Y over X.

In the election above **B is the Schulze winner**.

The Schulze method has many other advantages in a voting process, not least of which is that it is comparatively easy to compute. It fails Arrow's independence criterion, but satisfies many others, including the majority and the monotonicity conditions. Users of the Schulze method include Debian (2003), Gentoo (2005), Topcoder (2005), Wikimedia (2008), KDE (2008), the Pirate Party of Sweden (2009), and the Pirate Party of Germany (2010).

There are many other similar methods to find a Condorcet winner if one exists, or to find a winner otherwise, but my favourite amongst them all is **Dodgson's method**. This method of voting was invented in 1876 by

Charles Dodgson (illustrated below). Dodgson is otherwise known as Lewis Carroll, the author of Alice in Wonderland and Alice Through the Looking Glass.

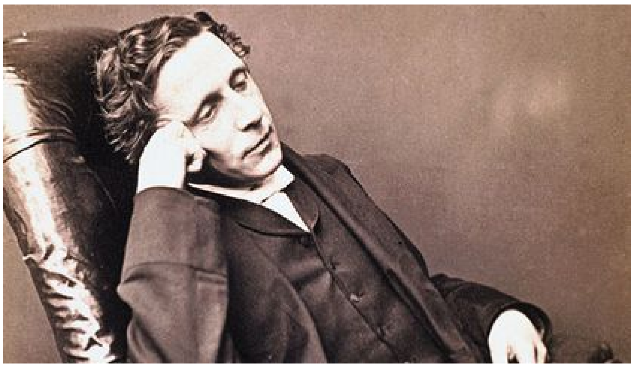

Dodgson was a mathematics don at Christchurch College Oxford, and whilst there he wrote a number of papers on voting. It is rumoured that he did so because a number of votes in the college had gone against him, including the design of the college belfry, which he referred to as a 'tintinnabulatory tea chest'. His (mathematical) life was celebrated in the recent Gresham lecture by Robin Wilson [7]. Amongst his contributions to the 'science' of voting we can include the papers on: *A Discussion of the various methods of procedure in conducting elections (*1873), *Suggestions as to the best method of taking votes, where more than two issues are to be voted on (*1874), *A method of taking votes on more than two issues (*1876), *Lawn tennis tournaments (*1883), *The principles of parliamentary representation (*1884).

The Dodgson method extends the Condorcet method by *swapping candidates*. The idea behind this is that in any election (such as the decision for who to appoint to a position), it is often a good idea to re run the vote to allow people to change their minds having seen how things have been developing. In the Dodgson method swaps are made between adjacent candidates until a Condorcet winner is found. The winner is the candidate which requires the **minimum number** of swaps to win. Dodgson proposed this voting scheme in his 1876 work *A method of taking votes on more than two issues* [8].

Let's see the Dodgson method in action for our problem above. The original table of preferences is:

Number of votes	Preferences
3	A > B > C > D
1	D > B > A > C
1	D > C > A > B
1	C > D > B > A
1	B > D > C > A

If **one** voter in the first row changes their preference to B > A > C > D (a single adjacency change) then B beats A 4 times and loses 3 times, so becomes the Condorcet winner. A lengthy calculation shows that for other candidates to win, more than one adjacency swaps are needed. So, **B becomes the Dodgson winner**.

Whilst in many ways intuitively appealing, the Dodgson method has significant disadvantages. It fails the monotonicity condition and many other conditions for a voting system. It is also NP-hard, which means that it takes a long time (more than polynomial time) to find the winner as so many changes to the voting preferences have to be considered. Thus, it is almost never used in practice

Practical voting methods
The voting methods described above go a long way to being fair and are based on firm mathematical principles. They are, however, rather sophisticated in the way that they arrive at the winner. This is fine if there is either a relatively small number of votes, or where the voting, and hence the processing, can all be done electronically. This is why they have taken off in internet voting. However, electronic voting has not (yet) been considered for

serious political decisions such as national elections, mainly due to issues associated with fraud. Instead, national elections prefer to use a system in which voters cast ballots on paper. These then have to be processed by hand, and the results announced rapidly. This means that practical voting systems have to be used which are a simplification of those considered above. Of course, the simplest of all of these is the first past the post system, which only requires one count of the votes to reach a decision (unless a recount is needed if the voting is close). Given the problems with this, other systems have been devised which have more of the merits of the ideal systems above but have to make compromises that make them all flawed in some way.

We will now give examples of a couple of these

The Instant Run-Off (IRV) or Alternative Vote (AV) System

The IRV voting method is used to elect a single candidate in an election. Unlike first past the post it allows the different preferences by the voters for the candidates to be taken into account. This system (used for example in Australia and Ireland where it is called preferential voting) works as follows. We assume as usual that we want to elect one candidate from several alternatives.

1. The voters rank each of the candidates in order of preference. An example of an Australian ballot paper is given below.

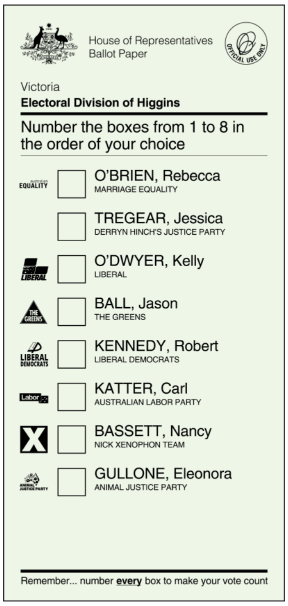

2. The ballots are initially counted for each voter's top choice and if a candidate has more than half of the vote based on first choices, that candidate wins.

3. If not, then the candidate with the fewest votes is eliminated.

4. The voters who selected the defeated candidate as a first choice have their votes for that candidate added to the totals of their next choice.

5. This process continues until a candidate has more than half of the votes.

6. When the field of candidate is reduced to two, it h becomes an instant run-off that allows a comparison of the top two candidates head-to-head.

IRV is used in several national elections. For example, it is used to elect members of the Australian House of Representatives and at least one house of all Australian state parliaments, the President of India and members of legislative council in India; the President of Ireland, members of Congress in Maine, in the United States; and the Parliament in Papua New Guinea. The method is also used in local elections around the world. The possible use of the AV system (as it is known in the UK) was put to a referendum of the British public in May 2011, and the motion was defeated.

A simple example (inspired by the Wikipedia article on IRV) is provided in the accompanying table. Three candidates are running for election, A, B and C. There are five voters, a, b, c, d, e. To win, a candidate must have a majority of votes that is, three or more. In Round 1, the first-choice rankings are that A and B both have two votes and C has one. No candidate has a majority, so a second "instant runoff" round is required. Since C is in bottom place, they are eliminated. Any voter who ranked C first (e.g. c) has their ballot modified as follows: the original 2nd choice candidate becomes their new 1st choice, and their original 3rd choice becomes their new 2nd choice. This results in the Round 2 votes as seen below. This gives B 3 votes, which is a majority. So the IRV winner is **B.**

	Round 1						**Round 2**					
Candidate	a	b	c	d	e	Votes	a	b	c	d	e	Votes
A	1	2	3	1	2	2	1	2	2	1	2	**2**
B	3	1	2	3	1	2	2	1	1	2	1	**3**
C	2	3	1	2	3	1						

As a separate calculation, the **Borda** winner in this case is **A.**

Note in this case that in the first round, in terms of voter preferences: B beats A, A beats C, and C beats B. So, this is a cyclic system with no Condorcet winner and no Copeland winner. As an exercise, work out the Shulze and Dodgson winners.

In fact, it can be shown that whilst IRV satisfies the majority condition, it fails both the monotone and Condorcet winner conditions. Thus, it is far from a perfect voting system.

<u>Single Transferable Vote (STV)</u>

This system, which was invented in 1819 by Thomas Hill is very similar to the IRV method, and is used if there are multiple candidates who can be elected to N posts. It is used, for example, for the election of officers to the London Mathematical Society (LMS) and in many different countries. It gives approximate proportionality when used correctly. There are many detailed variations of STV but essentially it has the same process as transferring votes as the IRV method. In each round the bottom candidate is removed, and their votes reallocated. The process continues until only N candidates remain.

STV has very much the same advantages and disadvantages as the IRV method. It is easy to use and approximately proportional but does not necessarily deliver the Condorcet winner.

Eurovision

The Eurovision Song Contest is run every year as a 'celebration' of music writing across Europe and beyond. Since its foundation in the 1950s it has grown into a major international institution. It takes place each year in the country that won it the previous year. That country, along with five other countries who put a lot of cash in (France, Germany, Spain, Italy and the United Kingdom) get automatic entry to the contest, and semi-finals are held in the week running up to decide which countries will make up the rest of the 26 entrants, whittled down from 40. As well as the songs it is a chance to celebrate ridiculous staging, hilarious costumes, cringe worthy announcers and sarcastic commentators. In my departments we have a sweepstake on the winner, with a special bonus award if your country gets Null Points. Occasionally there are even great songs, such as *Waterloo* from Abba in 1974. In 1994 the overall winner of the contest, then held in Ireland, was widely considered to be *Riverdance*, which happened to be the interval music. But as far as I'm concerned, it is the voting at the end which is by far the best part of the competition. It is here that we see the true conflict between a fair assessment of the value of each song, and outrageous tactical and political voting.

Before 1975 the winner in Eurovision was decided by using a form of the Borda system. Songs were rated by a jury of experts from each country with the top songs getting the scores of 12 (Douze Points), 10, 8, 7, 6, 5, 4, 3, 2, 1. Many songs are awarded 0 or Null Points at each stage. (It is a mark of great distinction to get Null Points overall. The first country to achieve this was Norway in 1978. Norway managed to again achieve this great result in 1981, with the UK following in 2003.) A country could not vote for its own song. The scores from each country were then added up to give a rank ordering as in the Borda method. Since 2016 a twin Borda method has been used which is combined with proportional representation. Each country produces a ranked list of the first 10 songs, one by the judges on the juries and the other by Tele-Voting. The latter vote is determined simply by proportional voting expressed by a single vote from each person taking part in the Tele-Vote. Each ranked list is then given the scores 12 to 1 as above. In the contest itself, the voting is revealed in two phases. In the first phase the votes from the juries are combined country by country. This gives the first rank ordered list and a nice build up of tension. After this has been done the Borda votes from the tele-votes in each country are simply added to the scores of the juries. With a total of 43 voting countries (the maximum number of participating countries), the maximum number of points one can mathematically receive is now 1008 (42 countries giving 12 points in each of jury and popular votes). In practice of course, a single country will get less than this. Here are the top ten results from 2016 computed using this system.

Country	Jury score	Jury Rank	Tele-vote score	Tele-vote rank	Total score	Total rank
Sweden	365	1	272	3	625	1
Italy	171	6	356	1	527	2
Russia	234	3	286	2	520	3
Belgium	186	5	190	4	376	4
Australia	224	4	124	6	348	5
Latvia	249	2	88	8	337	6
Norway	163	7	37	10	200	7
Estonia	53	11	144	5	197	8
Israel	77	8	102	7	179	9
Georgia	62	10	51	9	113	10

We have already seen that the Borda system has the disadvantages of not being a Condorcet method, nor of necessarily electing the song favoured by the majority of the voters. It does, however, work quickly, allowing real time scoring and a build up of tension as the results are announced. (This could not happen with the Shulze method for example.)

We have seen already that the Borda method is very vulnerable to tactical voting. In Eurovision this is taken to great extremes! Although countries cannot vote for themselves, there is very clear statistical [9] (and simply observational - watch it yourself) evidence of the existence of voting blocks in which one country will support another. This can be due to ethnic diaspora voting, a tendency for culturally close countries to have similar musical tastes, or straight political bias. Several countries can be categorised (see [9]) as voting blocs, which regularly award one another high points: Greece, Cyprus and Bulgaria; Turkey, Azerbaijan and Russia; Australia, Malta, Ireland, United Kingdom, Cyprus; Austria, Germany and Switzerland; The Netherlands and Belgium; Andorra, Portugal and Spain; Albania and Italy; Italy and San Marino; Sweden, Norway, Finland, Denmark, and Iceland; Estonia, Latvia, and Lithuania; Romania and Moldova; The Balkan countries; the former Yugoslav countries: Serbia, Bosnia and Herzegovina, Slovenia, Montenegro, North Macedonia, and Croatia; the former Soviet Union countries of Belarus, Ukraine, Russia, Azerbaijan, Armenia, Georgia and Moldova; Hungary and Serbia. This is all neatly summarised in the graphic below from LyricsTranslations.com.

I very much recommend that you watch the next Eurovision Song Contest to see the voting in action. You can mute out the sound if you wish.

References

[1] Kenneth Arrow, (1951) *Social choice and individual values*, Jon Wiley and Sons.

[2] Marianne Freiberger and Rachel Thomas, (2019), *Understanding Numbers*, White Lion Publishing.

[3] Nicholas Donaldson, Nora Donaldson and Grace Yan, (January 2018), *Brexit: Britain has spoken, or has it?* Significance magazine.

[4] *The Mathematics behind gerrymandering*
https://sinews.siam.org/Details-Page/detecting-gerrymandering-with-mathematics

[5] Marianne Freiberger, (May 2019), *No maths for Europe*, Plus Maths Magasine
https://plus.maths.org/content/democratic-dilemmas

[6] Markus Schulze, (2011), *A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method*, Social Choice and Welfare, **36**, 267–303. (Preliminary version in *Voting Matters*, 17:9-19, 2003.)

[7] Robin Wilson, (October 2019), *The mathematical world of C. L. Dodgson*, Gresham lecture.

[8] Charles L Dodgson (1876), *A Method of Taking Votes on More Than Two Issues*,

[9] Derek Gatherer, (2005), *Comparison of Eurovision Song Contest Simulation with Actual Results Reveals Shifting Patterns of Collusive Voting Alliances*. Journal of Artificial Societies and Social Simulation vol. 9, no. 2.

[10] Alexander Mantzaris, Samuel Rein and Alexander Hopkins, (2018), *Examining Collusion and Voting Biases Between Countries During the Eurovision Song Contest Since 1957*, Journal of Artificial Societies and Social Simulation vol. 21, no. 1.