

Probabilistic electoral methods, representative probability, and maximum entropy

Roger Sewell,
David MacKay,
Iain McLean

Abstract

A probabilistic electoral system is described in a context accessible to readers not familiar with social choice theory. This system satisfies axioms of: identical treatment of each voter and of each candidate; universal domain; fair representation of the pairwise preferences of the electorate; independence of irrelevant alternatives; and clarity of voting for pairwise outcomes; and hence Arrow's other axioms (weak Pareto and no dictator) are also satisfied. It produces in an information-theoretic sense the least surprising outcome given any candidate-symmetric prior beliefs on the voters' preferences, and is shown to be able to compromise appropriately in situations where a Condorcet winner would not be elected top under many other systems. However, difficulties can arise with this system in situations where one political party is permitted to flood the candidate list with large numbers of their own candidates.

The empirical properties of this system are explored and compared with the systems known as "Majority (or Plurality) Rule" and "Random Dictator".

We also make the case for using a probabilistic system even in the simple 2-candidate case.

1 Introduction

We offer a solution to a classic unsolved problem of democratic theory, viz., how to reconcile democracy with rights protection in a deeply divided society, as illustrated by one in which 60% of citizens are Tall and 40% are Short, and in which Talls and Shorts are in zero-sum competition over public goods. We

will refer to this solution as the "Maximum Entropy Voting System".

The problem has been recognised since Aristotle. Madison, Tocqueville, and J.S. Mill all discussed it extensively. Madison's solution is federalism. His classic expositions in the *Federalist* #10 and 51 are different, and arguably inconsistent, but both appeal to the concept of an extended republic. In *Federalist* #10 Madison argues that the extended republic, as a matter of sociological fact, will be sufficiently large that there will be no republic-wide faction capable of imposing its will on the minority by majority rule. In *Democracy in America*, Tocqueville confirms this sociological fact for the USA as he observed it in 1835. In *Federalist* #51 Madison argues that ambition must be made to counteract ambition, so that checks and balances, both vertical and horizontal, restrain full-throated majoritarianism¹.

J.S. Mill's approach is different, and in principle it applies to democracies of any size and constitutional structure, not merely to federal states. Chapter VII of his *Considerations on Representative Government* has the self-explanatory, if tendentious title 'Of True and False Democracy: Representation of All, and Representation of the Majority Only'. Mill here confronts the Aristotelian and Victorian nightmare that a monolithic working class might (soon) come to power and pass confiscatory legislation by majority rule. He discusses various schemes for proportional representation (PR), focusing mostly on the (wildly impracticable) scheme due to Thomas Hare. The Hare scheme is the ancestor of Single Transferable Vote as applied for national elections in both parts of Ireland, in the Australian upper house, and in many clubs and societies. Hare's original scheme was wildly impracticable because it treated the whole nation as a single district; vot-

¹ Aristotle, *Politics passim*, especially 1319b-1320a; J. Madison in *The Federalist* #10 and 51; A. de Tocqueville, *Democracy in America* especially Vol. I chs III, IX-XVI; J. S. Mill, *Considerations on Representative Government*, especially chapter VII.

For this publication, see www.votingmatters.org.uk

ers would have had to rank impossibly large numbers of candidates. The Australian and Irish implementations made Hare practicable by reducing district size. Both Irish implementations (north and south) were imposed by the British government before Irish independence in 1920-21. The Northern Ireland implementation was designed to protect the minority Catholic community there, and the Irish Free State implementation to protect the minority Protestant community there. The latter remains in the constitution of the Republic of Ireland, although the Protestant minority has dwindled to below the size that can be protected by the PR quota in use in the Republic (and has never been systematically persecuted).

The scheme below starts from a point that is well known, but little explored, in social choice. Satterthwaite (1975) proved that a direct implication of Arrow's Theorem was that all deterministic choice functions are either dictatorial or manipulable. Therefore, if you want a function that is neither, you should take probabilistic schemes² seriously. The best-known probabilistic scheme is the one called Random Dictator (RD) below. The idea goes back to ancient Greece, but has more recently been strongly advocated by Amar (1984). A version was proposed by Burnheim (1985) in ignorance of the social choice implications. We take both its merits and its demerits seriously and use it as a base for advance.

Gibbard (1977) considered probabilistic decision schemes (which ultimately output a top candidate with no ordering on the runners-up), and showed that given symmetry on candidates and voters, the combination of strategy-proofness and the weak Pareto property is enough to ensure that the scheme must indeed be RD. Moreover no probabilistic decision scheme (not even RD) can guarantee to provide an output distribution over the candidates that cannot be simultaneously bettered in the opinion of every single voter. McLennan (1980) extended these results to probabilistic social welfare functions (whose ultimate output is a strict total ordering over the candidates rather than just the identity of the top candidate) to show that if the symmetry axioms, strategy-proofness, and weak Pareto are met, then

² As we talk both of probabilistic voting and maximum entropy, it is useful to specify two traditions to which this paper does *not* belong. It is not about probabilistic voting theory in the sense used by Coughlin (1992), where the research question is the optimal strategy for a candidate who does not know for certain which voters are of which type. Nor is it about maximum entropy modelling in the sense used in many papers by R. J. Johnston and collaborators (e.g., Pattie et al. 1994), who use it as a technique to complete a flow-of-the-vote matrix with some unknown cells.

the induced decision scheme (which ignores the ordering other than for its top place) must be RD.

The authors believe that the symmetry axioms are the most fundamental³. Given these, we therefore cannot have both strategy-proofness and weak Pareto without confining ourselves to RD, whose weaknesses will be discussed below. Further, most would consider that failure to meet weak Pareto is more serious than failure to be strategy-proof. The approach we take therefore is to choose axioms weaker than strategy-proofness in its place, while retaining the symmetry axioms and the weak Pareto property.

However, most of the probabilistic systems that will be discussed coincide in the two-candidate case, and one of the first key points we want to make is that even in the two-candidate case, probabilistic schemes have very significant advantages over majority rule.

2 A tutorial exposition

(The reader who prefers mathematical precision will find it in the appendix section 9.)

The reader may first ask why there is any motivation to replace the simple and apparently easy to understand system of majority rule (MR)⁴. Our motivation is most easily seen by means of examples. The general setting will always be that each voter expresses their preference by placing the candidates in order, from first (most preferred) to last (least preferred), and the electoral system then gives an outcome, which also places the candidates in order from first to last⁵. The number of candidates elected will depend on the particular election; in some a single candidate is elected, in others several are elected (who occupy the top few places in the outcome). By taking this approach we tacitly assume that the ordering of the runners-up (if there is more than one runner-up) is an important part of the outcome of the election.

2.1 The problems with majority rule

Although majority rule has been in widespread use for many years, it has some important drawbacks.

³ Weighted voters are an easy modification of all the schemes considered, should one be so inclined.

⁴ Or plurality rule – for more than two candidates – the candidate with most votes wins. We will refer to this system throughout as MR (majority rule) for simplicity.

⁵ This type of system is known as a “(probabilistic) social welfare function” to distinguish it from a “(probabilistic) decision scheme” which only outputs the identity of the top candidate and ignores the ordering of the runners-up.

2.1.1 Majority rules – and in some places, always

By way of hypothetical example, let’s consider the little known country of Transmogria. It is inhabited by two peoples, the Talls and the Shorts. Talls make up 60% of the population, Shorts the remaining 40%.

Talls have had it their way for centuries. In consequence they are generally wealthier than the Shorts, and not surprisingly prefer policies of low taxation, low public spending, no provision for the poor, and no restrictions on employers in how they choose their employees or how they deal with them.

The Shorts differ from the Talls on a huge variety of issues: they want Fridays not Sundays as their regular day off, anti-discrimination laws and employee protection, a publicly funded health service, and a better choice of housing. Most of all, they want a say in how the country is governed – because under majority rule the Talls win the vote on every single issue all the time.

The result: for as long as anyone can remember, Transmogria has been in a state of civil unrest; the Talls claim that the Shorts are criminal political activists and protesters who continually resort to violence to achieve ends which “democracy” has ruled out, while the Shorts see themselves as oppressed by the Tall majority, and believe that their only recourse is to the armed struggle.

Let us suppose that Talls would consider themselves to be at 1.0 on a zero-to-one scale of satisfaction with the current situation, but that they would be at 0.0 if the Shorts somehow got into power. Likewise the other way round for the Shorts, currently at a satisfaction of 0.0. This all means that the average satisfaction level under majority rule is 0.6, but that the standard deviation across the population is 0.49. Surely there’s a better and fairer way to organise things than this.

2.1.2 No compromise

Recently a few brave people have migrated into Transmogria from the neighbouring country of Centralia. Appalled by what they found, they set up a small political party, the Compromisers, who, while they have the good of the whole population at heart, still only form 5% of the population. At significant cost to themselves, they have put forward a manifesto of tolerance and co-operation.

However, in every constituency only 20% of the Talls and 20% of the Shorts are prepared to vote for the Compromisers over their own party. These vot-

ers are divided randomly between the Talls and the Shorts in proportion to their occurrence in the population. The rest put their own party first, although they would happily put the Compromisers second over the opposition. The Compromisers vote for themselves first (believing they are a worthy cause with good intent) and equally for each of the other two second. The vote therefore splits as shown below – in the table we show not only each voter’s first preference but also his second and third:

Percentage of voters:

	45.6	13.9	30.4	10.1	0	0
voting:						
1st:	T	C	S	C	S	T
2nd:	C	T	C	S	T	S
3rd:	S	S	T	T	C	C

Table 1: The votes cast in an election between Tall, Short, and Compromiser. Each column shows the percentage voting for a particular order.

Since 45.6% of the voters placed the Tall candidate first (but only 30.4% placed the Short first and 24.0% placed the Compromiser first), in an MR election the Tall candidate would win.

However, one way to look at these votes is to examine which candidate would win in a head-to-head contest between any *two* candidates; if it should be the case that one candidate beats any other candidate in a head-to-head fight, it would be reasonable to hold that that candidate should be elected top. Let us therefore examine the table of preferences between pairs of candidates, which looks as follows:

Percentage of the population preferring c_1 to c_2 :

		c_2		
		T	S	C
	T	–	59.5	45.6
c_1 :	S	40.5	–	30.4
	C	54.4	69.6	–

Table 2: The pairwise preference table for the election of Table 1.

Thus we see that C would beat each of the other two parties in a straight two-candidate fight (as 54.4% of the voters prefer him to T and 69.6% prefer him to S) – such a candidate is known as a “Condorcet winner” - but under majority rule T always wins, with the result just as if the Compromisers had never existed. The only way C can win under MR is if tactical voting occurs – but Transmogrians would like to be straightforward and honest, and not have to engage in practices that require guessing the behaviour of the rest of the population.

In conclusion T wins under MR, even though more than half the population would have preferred C.

2.2 Arrow's theorem

Now suppose that, in the light of these problems, the Transmogrians decide to replace majority rule by some other more sophisticated system that takes into account not only the first preference of each voter but also the rest of their orderings of the candidates. A number of options, such as Single Transferable Vote and other forms of proportional representation come to mind; they are determined to pick a fair system which also has no incentive for tactical voting.

Unfortunately, they immediately hit a brick wall in the form of Arrow's theorem (named after Kenneth Arrow who proved it in the 1940s – see Arrow (1963)), which roughly says that no such system exists (a precise statement follows shortly).

Arrow's theorem, however, deals only with *deterministic* electoral systems. In these systems each voter votes by placing the candidates in order of preference, and the system then provides an output ordering of the candidates in which no two are ranked equal, as in all the systems we are considering; however where the system is deterministic, the output ordering is determined purely by the votes – if identical votes are cast in two elections, the output ordering will be the same in both.

His theorem proves that there is no deterministic electoral system which has even the following 4 minimal desirable properties, known as Arrow's axioms:

'No Dictator' (ND): there is No Dictator. In any system it would be a disaster if some voters were treated preferentially to others; one of the worst possible situations would be if the system treated one particular voter D as a 'Dictator', meaning that what D votes is automatically the result of the election.

'Universal Domain' (UD): if each individual voter votes legally, the system will output a valid election result. Thus for example UD would exclude a system that insisted on annulling the election if no candidate had an overall majority. It would also exclude a system that limited the number of candidates to 2.

'Irrelevant Alternatives' (IIA): whether the system outputs candidate c_1 above candidate c_2 depends only on how the voters ordered candidates c_1 and c_2 , not about where they placed any other candidate c_3 (i.e. the output takes no account of Irrelevant Alternatives).

'Weak Pareto' (WP): if everybody in the population prefers candidate c_1 to candidate c_2 , then so should the output of the system (this is known as the Weak Pareto condition).

Even though these properties are relatively simple and obviously desirable, Arrow's theorem tells us that we cannot have them in a deterministic electoral system. In particular it tells us that Single Transferable Vote, numerous other forms of Proportional Representation, and of course the British majority-rule system, cannot meet these simple requirements.

Fortunately, Arrow's theorem applies only to deterministic electoral systems, i.e. ones in which a particular set of votes always results in a particular outcome. In order to achieve a good electoral system, therefore, we have to 'think out of the box' and move to a system in which some other factor, as well as the votes, influences the results of the election. That other factor must be one that carries no bias, and allows the system to meet an appropriate set of axioms that should ideally include Arrow's four (ND, UD, IIA, and WP), but which should also include other much stronger axioms (such as SV (Symmetry among Voters) which requires that all voters are treated equally).

The factor that the Transmogrians are looking for is *randomness*. We will introduce this again by way of an example.

2.3 A simple alternative method for the two-candidate situation

We bring in this example first as a two-candidate situation, and later expand it.

2.3.1 The two-candidate probabilistic election

Returning to our two-candidate situation in Transmogria, consider the following simple but perhaps unexpected electoral system.

As before, we have two candidates in each constituency - a Tall and a Short. The Talls vote for the Tall candidate, while the Shorts vote for the Short candidate. Thus the vote splits:

Percentage:		60.0	40.0
1st:	T	S	
2nd:	S	T	

Table 3: The votes cast in an election between Tall and Short.

Now we draw a *random* ordering of the candidates, with probability 0.6 of picking $T > S$, and 0.4

of picking $S > T$, according to the fractions of the voters voting each way. Therefore T is elected with probability $3/5$, and S with probability $2/5$. Thus, supposing an election is to be held every year, in roughly three years out of five the constituency will be represented by a Tall, and in two out of five, it will be represented by a Short.

Thus the fraction of the time that the constituency is represented by a Tall will be equal to the fraction of Tall voters, and the fraction of the time it is represented by a Short will be equal to the fraction of Short voters. This abolishes the permanent rule by a majority over a large minority, while still being 'fair', in that an outcome occurs with probability proportional to the fraction of voters that favour it. As an illustration of the idea that this is fairer than MR, consider the spread of expected satisfaction across the population. Under the MR system the Shorts always have a satisfaction of 0.0 while the Talls always have a satisfaction of 1.0 giving a standard deviation of expected satisfaction across the population of 0.49; under this non-deterministic system the Shorts have an expected (average) satisfaction of 0.4 while the Talls have an expected satisfaction of 0.6, giving a standard deviation across the population of only 0.1 – thus satisfaction is being dealt out more evenly across the population. True, the overall average satisfaction has gone down from 0.6 to 0.52 – but this is a small price to pay for making the results fairer.

Before moving to the three-candidate situation, let us address two worries that are likely to occur to many readers.

2.3.2 Interlude to address the worries of how a random number can be chosen without abuse

Some people will immediately be worried about how such a random choice can be made without abuse; after all, we all know how difficult it can be to get two children to accept a coin toss as a decision between their preferences.

A suggestion is that we could adopt something like the following procedure.

The UK lottery machine (which is carefully arranged so that the number of balls can be checked at the beginning) is used to draw a random sequence of five balls out of a hundred. Each has a number between 0 and 99. The result is a ten-digit number. This is used to seed a pseudo-random number generator in a computer program which everybody in the country can inspect, replicate, and run. A random number x is then drawn from the generator in a

prescribed way such that it is uniformly distributed between 0 and 1 (i.e. $0 < x \leq 1$). If $0 < x \leq 0.6$ then T is elected, otherwise S is elected. Everybody in the country is able to check the result, as they have all watched the draw on television. Specifically the candidates (and an audience) can be present at the draw to verify that the procedure was carried out rather than the television transmission synthesised to deceive. Since everybody can inspect the software, everybody can check that it is fair.

This procedure is capable of setting up sequences of random numbers as well as individual ones, so that any computer software requiring random numbers can be initialised in this manner.

2.3.3 Interlude to address how things work in parliament

If we employ the above method in a two-party situation with many constituencies, each electing a candidate to represent it in parliament, then we have an ongoing problem when votes are taken in parliament. If decisions in parliament are still taken by majority rule, we will probably fail in our desire to reduce differences in satisfaction between different parts of the community.

To see this, consider Transmogria, voting as above with only the Talls and the Shorts present. Suppose there are 600 seats in parliament, elected using the probabilistic system described in section 2.3.1 above. In most years, we will see roughly 360 Tall and 240 Short members of parliament, varying by roughly 30 seats either way. Only once in every few thousand years will there be a Short majority in parliament. Therefore if the Talls want to pass a law that door handles should always be mounted six feet off the ground, they will succeed.

However, if parliament also passes or rejects bills in the same random way that MPs are elected, the fact that the Talls nearly always have a majority is less of a problem. The high-door-handle bill will be passed with probability 0.6 rather than 1.0.

However there are further considerations. If a few years later a contrary bill is introduced, insisting that door-handles are always mounted six inches off the ground (and repealing the old law), it will pass with probability 0.4; in this case this sequence of events is reasonably fair, though a right nuisance for builders and carpenters and those who have to pay for the door-handles to be moved.

The situation could however be much worse: the Talls might propose a bill to demolish the 5000-year-old historic Palace of the Shorts, or bomb the neighbouring country of Dwarfland. If the bill is

rejected (which it will be with probability 0.4), the Talls may reintroduce it, again and again, until it is passed once – after which it is too late to redress the situation. Indeed the probability that it will eventually be passed converges to 1 as the number of reintroductions approaches infinity.

Therefore, procedure, in particular the process whereby bills are introduced for consideration by parliament, also needs to be regulated, if probabilistic methods are also to be used in parliament. Alternatively, Amar (1984) believed that majority rule could be retained in parliament without losing the value of having probabilistic election of the representatives, while Wichmann (2009) believes that this issue is better dealt with by Human Rights legislation.

The issue of how parliamentary votes are conducted is an issue we will not address further in this paper, but which needs further thought. If stability of legislation is to be achieved, participants will have to achieve greater degrees of consensus than occur at present in parliamentary democracies.

2.4 Desirable axioms

We will now expand our horizons to take in elections with more than two candidates, and elections in which we may be electing more than one of the candidates. In all cases we will be interested in the whole of the outcome ordering of the election, even though not all candidates are being elected.

First, however, we need to consider what axioms we want our new probabilistic electoral system to satisfy.

The following axioms are potential candidates. All are more precisely defined in the appendix section 9.

‘Symmetry among Voters’ (SV): Each voter is treated identically; if the views of two voters are swapped, the probability of any given result should be unchanged.

‘Symmetry among Candidates’ (SC): Each candidate is treated identically. If a set of votes V leads to election result Q with probability p and $V_{c_1 \leftrightarrow c_2}$ denotes those votes with every voter’s views on candidates c_1 and c_2 swapped, and $Q_{c_1 \leftrightarrow c_2}$ denotes the result Q with the positions of candidates c_1 and c_2 swapped, then if the voting is $V_{c_1 \leftrightarrow c_2}$ the probability of getting result $Q_{c_1 \leftrightarrow c_2}$ should be p .

‘Universal Domain’ (UD): If each voter has voted legally, then the collection of all voters’ votes is legal and the electoral system will output a valid election result.

‘Clarity of Voting (Pairwise)’ (CVP): The best way for a voter to achieve candidate $c_1 >$ candidate c_2 is to vote $c_1 > c_2$ (i.e. the probability of the output ordering placing candidate $c_1 >$ candidate c_2 should be equally maximised by any vote that places candidate $c_1 >$ candidate c_2).

‘Representative Probability’ (RP): the probability of the outcome putting candidate c_1 above candidate c_2 should be the same as that of a randomly chosen voter preferring candidate c_1 to candidate c_2 . In other words, the probability of the outcome putting one candidate above another should be the same as the fraction of the voters preferring the one to the other.

We believe that these axioms cover most of what is required of an electoral system, but not quite all, as we shall see later. They are in particular sufficient to imply Arrow’s axioms WP, ND, and IIA, where we restate the last as:

‘Independence from Irrelevant Alternatives’ (IIA): The probability that the system places candidate c_1 above candidate c_2 should depend only on the voters’ orderings of candidates c_1 and c_2 and not on where they place any other candidates.

We should particularly note the subtle differences between CVP and two related axioms CVT and CVO, which we state adjacent to each other here for easy comparison. In each case the phrase “The best way to achieve X is Y” means that the probability that X occurs in the output ordering is (equally) maximised by any vote that satisfies Y.

‘Clarity of Voting (whole Ordering)’ (CVO): The best way for a voter to achieve a particular ordering of the candidates in the result is to vote for that ordering of the candidates.

‘Clarity of Voting (Pairwise)’ (CVP): The best way for a voter to achieve candidate $c_1 >$ candidate c_2 is to vote $c_1 > c_2$.

‘Clarity of Voting (Top)’ (CVT): The best way for a voter to achieve candidate c_1 being placed top in the output ordering is to place him top in that voter’s vote.

These axioms turn out **not** to be equivalent. The system known as ‘Random Dictator’ (described below in section 2.6) satisfies all three of the CV axioms, while ‘Maximum Entropy Voting’ (described below in section 2.9) satisfies CVP but not CVO or CVT. ‘Sequential Random Dictator’ (described below in section 2.7) obeys CVO and CVT but not CVP. Unfortunately, there turn out to be significant disadvantages associated with the *known* methods of complying with all three of the CV axioms.

We will keep as a **‘Standard List of Axioms’ (SLA)** that a system should obey the following: SV,

SC, UD, and RP. A system obeying SLA then also obeys WP, ND, IIA, and CVP by virtue of obeying RP.

2.5 More than two candidates – preamble

Consider again the situation in Transmogria after the immigration of some Centralians, previously considered in section 2.1.2. The voting pattern we are dealing with is:

Percentage of voters:	45.6	13.9	30.4	10.1	0	0
voting:						
1st:	T	C	S	C	S	T
2nd:	C	T	C	S	T	S
3rd:	S	S	T	T	C	C

Table 4: The votes cast in the same election as in Table 1.

The question is how our new electoral system should set the probabilities with which it outputs each of the six possible orderings of the candidates.

The key difference between the two-candidate situation and those where there are more candidates is the following. In the two-candidate situation there is only one system that obeys RP (namely the one described in section 2.3.1). It turns out that when there are more candidates, there are an infinite number of systems that obey SLA (our Standard List of Axioms, which of course includes RP). We will now consider a few such systems, and one that doesn't obey SLA.

2.6 The 'Random Dictator' system

Considering the election of Table 4, we ask again how the probabilities of the different outcome orderings should be set. The first possible answer (though not necessarily the best) is to set the probabilities of the various orderings to be the same as the fractions of the voters voting for each ordering. This is sometimes known as the 'Random Dictator' (RD) system, as it is equivalent to the following procedure:

Everybody casts their votes; then

A voter is picked at random and the output ordering of the election is set to be the ordering given by that voter.

Since the voter who is picked gets his own views output by the electoral system, he is known as the 'dictator' (the 'random dictator' since he was chosen at random). Note that this system *does* obey the No Dictator (ND) axiom – there is no (fixed) voter

who can dictate the outcome (the 'random dictator' is chosen at random each time an election is held).

This procedure has the benefit of extreme simplicity. There is also a total absence of any computational difficulty beyond choosing a random member of the electorate. RD is very easy to understand.

RD also satisfies all the axioms so far considered.

This procedure, however, has some important disadvantages.

2.6.1 No compromise

The RD system can elect to *top* position only candidates whom some voter has put in top position. There is no possibility of placing a candidate who is everybody's second choice at the top, even though they may be preferred to any other candidate by a majority of the electorate (i.e. be a Condorcet winner).

2.6.2 No moderation – or 'It never rains but it pours'

Suppose a population consists of 50% Tall voters and 50% Shorts. Suppose, moreover, that there are ten candidates from each of these parties (a total of 20 candidates) of whom a total of eight will be elected (the eight at the top of list). Each voter places all the candidates from his party in some order at the top of the list, followed by those of the other main party.

In this situation the RD system as it stands will elect either eight Talls, or eight Shorts – but *never* a mixture of the two. This characteristic is the opposite of moderation.

The situation could be even worse if there is a small minority of "Exclusive Talls" who want to shoot all the Shorts, and who also field 10 candidates. With a 1% proportion of Exclusive Talls there would be a 0.01 probability that *every single elected candidate* would be an Exclusive Tall.

2.7 The 'Sequential Random Dictator' (SRD) system

One approach which might at first sight ameliorate the 'No Moderation' defect of the Random Dictator system is the 'Sequential Random Dictator' (SRD) system.

In this system the candidate to be placed top in the output ordering is selected according to the same technique as employed by the RD system. However, rather than then taking the dictator's views on second, third and subsequent placings, the candidate

placed top is removed from everybody’s votes, and a new voter is chosen at random to be dictator 2. The candidate at the top of dictator 2’s ordering (which has already had the candidate placed overall top removed from it) is then elected in second place. We continue selecting a new dictator at random until all places in the ordering are filled. This system avoids the lack of moderation described in section 2.6.2 above by changing the dictator after each place has been filled.

Moreover, this modification to RD is easily seen to result in SRD still obeying SV, SC, UD, ND, WP, and CVT. Let us consider whether it obeys RP and/or CVO and/or CVP; in other words roughly “Is it fair?” and “Is there an incentive for tactical voting?”.

2.7.1 Does Sequential Random Dictator obey Representative Probability?

Recall that ‘Representative Probability’ (RP) states that the probability of the output distribution yielding candidate $c_1 >$ candidate c_2 should be equal to the fraction of the population so voting. It is obeyed by RD, so it is perhaps slightly surprising that it is *not* obeyed by SRD. This is easily seen from the following table of voting on three candidates A, B, and C, and the probabilities of the output giving each ordering underneath, under the RD system and under the SRD system:

	Percentage of voters giving each order:					
	50.0	0	0	0	0	50.0
1st:	A	A	B	B	C	C
2nd:	B	C	A	C	A	B
3rd:	C	B	C	A	B	A
RD prob:	0.50	0	0	0	0	0.50
SRD prob:	0.25	0.25	0	0	0.25	0.25

Table 5: The votes cast in an election between A, B, and C and the outcome distributions under Random Dictator and Sequential Random Dictator systems.

The resulting pairwise preference table for the voters is:

Percentage of the population preferring c_1 to c_2 :

	c_2		
	A	B	C
c_1 :	A	–	50.0
	B	50.0	–
	C	50.0	50.0

Table 6: The pairwise preference table for the election of Table 5.

However, the SRD output distribution pairwise preference probability table is:

Probability that output prefers c_1 to c_2 :

	c_2		
	A	B	C
c_1 :	A	–	0.75
	B	0.25	–
	C	0.5	0.75

Table 7: The pairwise preference table for the outcome distribution under the SRD system in the election of Table 5.

Thus, although half the population voted A>B, the probability of the output ordering under SRD giving A>B is three-quarters. This is contrary to RP and gives a severe disadvantage to B; it also illustrates how some simple modifications of systems that obey all the axioms can fail to obey even basic ones.

2.7.2 Does SRD obey CVO and CVP?

It is relatively easy to show that SRD does obey CVO, but not CVP. We omit the proofs for brevity, given that SRD has already been shown to be wanting by failing to obey RP.

2.8 A conjecture

We conjecture that any probabilistic social welfare function satisfying SLA and CVO and CVT induces RD as the induced probabilistic decision scheme on any subset of the candidates chosen after the votes are cast.

Note that Gibbard (1977) and McLennan (1980) have together shown that SLA and SP2 (defined in the appendix section 9) are sufficient to ensure that RD is indeed induced.

2.9 The ‘Maximum Entropy Voting System’ (MEV0)

2.9.1 Description

Axioms to be complied with

Returning now to our Standard List of Axioms, we will restrict our attention to those systems that do satisfy SLA; we will not require adherence to CVO or CVT, since we have not been able to find a system that obeys these also without suffering the disadvantages of the Random Dictator system. There will of course be some disadvantages to not obeying CVO and CVT, in the form of susceptibility of some

properties of the output distribution to some forms of tactical voting; however no tactical voting will be able to influence the probability that one candidate is preferred by the system to another, because as we have seen SLA implies CVP.

The principle of minimising information

The basic idea and motivation are as follows. The problems with RD, notably the lack of moderation noted in section 2.6.2, stem from taking *too much* information from the votes – with RD, the outcome distribution matches the voter distribution too precisely. We want to reduce the information taken from the votes to a precise set of variables, the minimum set needed to ensure that RP is satisfied. Further, as it will turn out that restricting the information taken from the votes to precisely that set of variables is not sufficient to specify the system uniquely, we will turn our attention to that system in which the votes also give us the minimum amount of information about the ordering that will actually be chosen.

Two senses of the word information

There is quite a subtle distinction here between two uses of the word information, which we will dwell on briefly as it is important to what follows.

If the system's output depends only on the table of the fractions of the voters who prefer each candidate to each other candidate (for N candidates this is a total of $N(N-1)/2$ independent numbers), and if no two distinct such tables lead to the same output distribution, then in the first sense of the word information we have defined precisely which information we have taken from the votes. However, there are many systems that could be based on taking only this information from the votes (and that satisfy SLA), hence precisely defining the information taken from the votes is insufficient to uniquely specify the electoral system.

There is however, a second sense of the word information. If I tell you that there has not been an earthquake in London today, I am telling you little information, but if I tell you that there has been one, I am telling you a lot. Equally if I tell you at least one African died today, nobody will be surprised (because little information has been conveyed), while if I tell you that not a single African has died today, it will be newsworthy because a lot of information has been conveyed. In this sense the amount of information conveyed is greater if after receiving it our knowledge is very different from what it was before.

Now, let us consider an example. Suppose that there are six Tall candidates and six Short. Suppose also that we believe in advance that half the population prefers all the Tall candidates to all the Shorts, while the other half prefers all the Shorts to all the Talls, while within each group (Tall or Short) all voters are equally likely to have any preference. Suppose moreover that the voters do so vote. Whereas RD can only output highly polarised orderings with all Shorts above all Talls, or vice versa, there are other probability distributions over the output orderings which also satisfy RP: for example, the uniform distribution over all possible output orderings. RP simply requires (for this *particular* voting pattern) that for any pair of candidates, the output is equally likely to place one above the other as the other way round. If under this uniform distribution we were to find all the Talls above all the Shorts, this would be a considerable surprise, and the occurrence of this event would be newsworthy, i.e. carry a lot of information in comparison with finding one of a number of nondescript orderings. As we will see, MEVO will indeed output all the possible orderings with equal probability – minimising the surprise, and the information content of the output ordering about the votes, while still adhering to RP.

However, information content does depend on prior belief. Reverting to a two candidate election with one Tall and one Short candidate, suppose, strictly hypothetically, that we were to believe in advance that the Tall candidate is almost certain to be elected. Suppose then the votes, combined with the electoral system in use, confirm that the Tall candidate has been elected. Then in the second sense of the word information (the information-theoretic sense), the votes have supplied little information (that we didn't already know) about the result of the election. If however the Short candidate is elected, the votes have supplied comparatively more information. Therefore in the following paragraph we confine our attention to prior beliefs that are symmetric among the candidates.

Minimal information means minimal surprise

Let us now suppose that in advance our beliefs about the votes are symmetric among the candidates; i.e. we may believe that it is more likely that the voters will align into two camps than that there will be an equal number voting for each candidate, but if so, we believe that it is equally likely that the two camps are favouring candidates A and B, as that they favour A and C, or B and C. *Under these circumstances we also aim to minimise the informa-*

tion about the result ordering given by the votes. (Not to require the symmetry among candidates in our prior beliefs about the votes risks violating SC in the resulting system.) In other words we want the outcome of the election to be as *unsurprising* as possible, given whatever candidate-symmetric prior knowledge about voting patterns we may have had, while of course maintaining compliance with SLA.

What we give up when we go for minimal surprise

We should note here that in deliberately choosing to go for minimal information, and hence minimal surprise, we are deliberately saying we want one general sort of outcome rather than another.

Suppose there is a religious minority in Transmorgia, the Narrows, who form 1% of the population. Suppose the Narrows will only be happy if their 10 candidates occupy *all* the top 10 positions in the outcome ordering; getting a mere 9 candidates in the top 10 positions is something they would regard as an outcome tainted by heresy, and no better than having all of their candidates come bottom.

Under RD the Narrows will be happy 1% of the time – just as they form 1% of the population – because 1% of the time, under RD, Narrow candidates will occupy all the top 10 places in the outcome ordering. However, the Narrows will not like a system that minimises information conveyed, because it is extremely unlikely to yield the very surprising outcome that a party with a tiny minority of support gets *all* its candidates in the top positions.

Introducing MEV0 is an action of people who do not want such surprising outcomes; it must be realised that introducing MEV0 will reduce the possibility of minorities such as the Narrows ever being happy.

A somewhat more mathematical point of view

From a mathematical point of view, for any given system and for any set of votes, the system gives a set of probabilities on the set of orderings of the candidates, which are non-negative and which sum to 1. If there are N candidates, there are $N!$ orderings, and the possible probability distributions may be represented as points in $N!$ -dimensional real space $R^{N!}$; in fact they all lie in a $(N! - 1)$ -dimensional simplex that lies obliquely across the corner of the positive ‘quadrant’ of this space. Let U denote this simplex.

Now, given a particular set of votes, adherence to RP implies that the set of points that could be occu-

ried by the outputs of systems adhering to SLA lies within a hyperplane satisfying $N(N - 1)/2$ linear constraints; let U_0 denote the intersection of this hyperplane and U .

There are many ways to place the output of such a system within U_0 and still ensure that it satisfies SLA. Some we might consider are the mean of U_0 , the point in U_0 that is closest to the origin, etc – indeed almost any point that can be distinguished without specifically referring to any voter or any candidate. So how are we going to choose one?

Now, there is a quantity called *entropy* (of a probability distribution over e.g. a finite set T) that measures the uncertainty we have about a choice of elements of T . If the distribution puts probability 1 on one element of T and none on the others, the distribution has zero entropy; the uniform distribution on T will have the maximum amount of entropy possible. Entropy is in an important sense the opposite of information (in its second sense): when we acquire information about a quantity, on average we reduce the entropy of the distribution that describes what we now know about that quantity.

Thus if we want to choose a distribution in U_0 that minimises the amount of information about the output ordering we are supplying, we should choose the distribution in U_0 that has maximum entropy. Fortunately it turns out that this specifies a unique distribution.

Another way of looking at this is to note that RD is not at all moderate (as noted in section 2.6.2). So we may ask what is the most moderate distribution we can find in U_0 ? One might argue that the most moderate distribution is the one that mixes in as many different output orderings as possible, while still adhering to RP. This again leads us to the distribution in U_0 that has maximum entropy. Let us call that distribution u_1 .

Therefore we choose to define the **Maximum Entropy Voting** system (MEV0) (the 0 (zero) is introduced because we will in a future paper define variants of MEV) as that system which outputs an ordering chosen at random from the distribution u_1 in U_0 that has maximum entropy. That distribution, of all those in U_0 , ensures that the votes give us least information about the actual ordering that will finally be output by the system when the random draw from u_1 is made, ensures that that ordering will be as unsurprising as possible, ensures that we know exactly what properties of the votes are being extracted and used, and in an important sense is the most moderate distribution consistent with obeying RP.

A more formal definition of MEV0 is given in an appendix (section 9) and a discussion of implemen-

tation will follow in section 5 below.

2.9.2 Example

Taking the three-candidate problem of section 2.1.2 as an example, let's see how the Compromiser party fares under MEVO. The voting pattern is, as before, to within rounding error, given in Table 8 along with the outcome probability distribution under MEVO:

	Percentage voting each ordering:					
	45.6	13.9	30.4	10.1	0	0
voting:						
1st:	T	C	S	C	S	T
2nd:	C	T	C	S	T	S
3rd:	S	S	T	T	C	C
MEVO Prob:	0.247	0.238	0.095	0.212	0.099	0.111

Table 8: The voting pattern as in Table 1, along with the outcome distribution under the Maximum Entropy Voting system.

In this result we see that the probability of the Compromiser being elected top is not zero (as it was with Majority Rule), or 0.24 (as it would be with RD), but 0.45 (i.e. $0.238 + 0.212$). That this higher value is more appropriate is seen from the pairwise preferences table for the population:

Percentage of the population preferring c_1 to c_2 :

		c_2		
		T	S	C
c_1 :	T	—	59.5	45.6
	S	40.5	—	30.4
	C	54.4	69.6	—

Table 9: The pairwise preference table for the voters in the election of Table 8.

which shows that Compromisers are preferred by a majority of the population to any other single candidate (and they are the only candidate with this status). In contrast, the Talls are elected top with probability 0.358 and the Shorts with probability 0.194. One can of course also verify that RP is being met, by calculating the pairwise preference table for the outcome distribution and showing that it is identical to that for the votes.

2.10 The spectrum from 'Random Dictator' to 'Maximum Entropy Voting'

Now, when we defined MEVO, we specified essentially two things. First, a set of constraints that the MEVO output distribution must satisfy (leading to a set U_0 of potentially satisfactory distributions), and

second, a rule by which to choose which distribution in U_0 we should draw the output ordering from, namely the maximum entropy rule.

There are many ways we could specify the set of constraints. At one end of the spectrum we could say that the probability of the output distribution giving each particular ordering should be the same as the fraction of the voters giving that ordering – this would define the RD system, as there would then be only one distribution in U_0 . At the other end of the spectrum we could require adherence to RP only, yielding the MEVO system. In between there are a variety of other sets of properties in which we could require the output distribution to match the votes distribution.

For example, one could specify that the output system should also give the same probabilities of ordering all subsets of three candidates in each of the six possible ways for each such subset as the voters did. (For a three-candidate election, that would in fact force the system to be RD, but for more candidates such a system would be distinct from RD). Alternatively, one can allow the voters to express their preferences not just as a set of pairwise preferences (which is all that is taken from the orderings by MEVO), but also by *optionally* stating combinations of pairwise preferences that they want to occur together (e.g. $c_1 > c_2$ and $c_1 > c_3$).

In each case there are two technical constraints that we must ensure are satisfied, namely non-emptiness of the set U_0 of potentially satisfactory distributions, and convexity of that set. Provable non-emptiness is required because otherwise we cannot guarantee to meet UD (there may be some voting patterns for which there is no possible output distribution), and we choose to require convexity because otherwise we may not be able to prove that the maximum entropy rule chooses a unique distribution.

Now, non-emptiness of the set of potentially satisfactory distributions is guaranteed for any such system by the fact that the RD output distribution, in all senses equal to the vote distribution, matches the vote distribution in all the properties we might consider incorporating. Convexity will be guaranteed providing the constraints specified are of the form

$$P(\text{outcome ordering has property } X) = P(\text{the ordering of a randomly chosen voter has property } X)$$

where the equality can be replaced by a non-strict inequality in either direction.

3 Measures of satisfaction under MEVO

In the Transmogrian two-candidate election discussed in section 2.1.1 we noted that, under majority rule, the mean satisfaction of the population was 0.6 and the standard deviation of satisfaction level across the population was 0.49. Now, it may easily be seen that if any probabilistic system obeying SLA is used (and there is in fact only one such system in a two-candidate election, namely that introduced in 2.3.1, which coincides with both RD and MEVO in this setting), the mean satisfaction of the population will be 0.52 while the standard deviation of expected satisfaction level will be 0.1. Though there is some reduction in *average* satisfaction, satisfaction is much more fairly distributed through the population.

In situations where there are more than two candidates, we now ask whether similar improvements can be obtained from probabilistic systems such as RD and MEVO.

In order to get an empirical measure of the benefits of probabilistic systems we simulated elections on four candidates, and considered various ways in which the voters' opinions on individual candidates might combine to give an overall satisfaction with an outcome ordering. We ran 400 different elections (different sets of votes) and drew 500 random samples from the output distributions of each election under each system (for majority rule all 500 random samples were of course the same, since majority rule is a deterministic system).

For each election, we started off by simulating the opinions of the voters. The details of how this was done are in an appendix (section 10 below). In each election, the voters were clustered in 8 different broad clusters in their opinions, with the positions of the clusters being randomly distributed with a tendency to avoid neutral opinions. This resulted in each voter having a score (the "input score") between zero and one for each candidate, indicating how much they liked that candidate.

We then deduced from these scores the order that the voters would place the candidates in when voting (assuming that each voter votes his true opinions).

We then applied each electoral method to the votes, and deduced the output ordering distribution. We drew 500 sample results from the distribution for each election, each of which is an ordering of the candidates.

It was then necessary to consider how satisfied an individual voter would be with any specific out-

put ordering. We considered three different ways in which the output ordering might be combined with the voter's scores on the candidates to give an overall satisfaction rating for each voter.

In the following definitions, the words 'rank', 'score', and 'correlation' will have the following meanings.

'Rank' means 1 for the top candidate, 2 for the candidate in 2nd place, etc.

'Score' refers to either input score as described above, or to an output score derived from the output ordering by drawing a uniform-random score vector from those score vectors that would place the candidates in the chosen output ordering; in the latter case score has the intuitive meaning that it is the degree of satisfaction the system's output ordering had with that candidate, just as input score is the individual voter's degree of satisfaction with the candidate. More formally, if there are N candidates, then the unit hypercube I_N in R^N may be identified with the possible sets of scores on the candidates, and each point in it thus mapped to a particular ordering of the candidates. Given a particular output ordering of the candidates, the output score was then drawn uniform-randomly from that subset of I_N that is mapped to the given output ordering.

'Correlation' between two vectors means the cosine of the angle between the two vectors. Thus if the two vectors point in the same direction (e.g. two rankings of the candidates are identical) the correlation will be +1. If they point in exactly opposite directions in R^N it will be -1.

The three methods by which opinions on individual candidates were combined to give a voter's overall satisfaction with the result of the election were then as follows:

1. Correlation of output rank with voted rank (**RankCorrel**) (which might be expected to give the advantage to MEVO or RD);
2. Correlation of output score with voted score (**ScoreCorrel**) (where the output score is as explained above);
3. Input score of the candidate most preferred by the result ordering (**Winner'sScore**) (which might be expected to give the advantage to MR).

For each measure of satisfaction, a number of descriptive statistics were calculated and used to summarise the characteristics of how satisfaction was distributed among the voters and between elections.

We choose to name these statistics **unfairness, average satisfaction, macrovariation, microvariation, and immoderation**. Their precise definitions are given in the appendix section 11 below.

	Method of measuring satisfaction		
	RankCorrel	ScoreCorrel	Winner'sScore
unfairness (small is good):			
MEV0	0.0746	0.0186	0.155
RD	0.0762	0.0280	0.157
MR	0.256	0.0801	0.297

“Unfairness” captures the degree to which we can expect different members of the population to be disgruntled with the electoral system to differing extents – we would like this number to be small, indicating that everybody can expect to be similarly satisfied over the long term.

average satisfaction (large is good):			
MEV0	0.547	0.512	0.536
RD	0.548	0.521	0.537
MR	0.613	0.517	0.589

“Average satisfaction” is self-explanatory – it is an average over everything, and the bigger it is the better.

macrovariation (small is good):			
MEV0	0.0722	0.0259	0.0742
RD	0.0703	0.0257	0.0721
MR	0	0	0

“Immoderation” captures the degree to which the system is likely to produce extreme outcomes; for example, a system that is immoderate without being unfair is one which given a 50/50 split of the electorate either has all the elected candidates coming from one party, or all from the other, but never a mix – so we would like immoderation to be small.

microvariation (small is good):			
MEV0	0.274	0.0840	0.243
RD	0.287	0.0984	0.245
MR	0	0	0

Finally “macrovariation” and “microvariation” capture different aspects of how the system causes variable degrees of satisfaction as the random number generator seed ω changes; “microvariation” captures the variability with ω seen by an individual voter, while “macrovariation” captures the variability with ω of average satisfaction over the population. For a deterministic system these two quantities will of course be zero. Ideally we might wish these parameters to be small – but as we have seen, Arrow’s theorem prevents us meeting SLA and having zero values of macro- and micro-variation.

immoderation (small is good):			
MEV0	0.273	0.0805	0.288
RD	0.289	0.0992	0.293
MR	0.256	0.0801	0.297

Table 10: The various statistics of the three measures of satisfaction under the Maximum Entropy Voting (MEV0) system, the Random Dictator (RD) system, and the Majority Rule (MR) system. Explanations of the statistics (unfairness, ..., immoderation) are given in the preceding text, while their precise definitions are given in appendix section 11.

The level of estimated uncertainty in these statistics is mostly small compared with the differences between them.

It is interesting that MEV0 causes less unfairness than the other two systems *whichever* satisfaction measure was used (and a lot less unfairness when assessed by score correlation). Similarly MEV0 is less immoderate than RD, whichever satisfaction measure is used. To ‘pay’ for this reduction in unfairness, MEV0 loses only around 10% on average satisfaction (more like 1% if satisfaction is measured by score correlation) compared with MR (as the price of adhering to sensible axioms), but it does of course introduce micro- and macrovariation because of its non-deterministic nature; nonetheless the macrovariation is small.

The results for one particular set of 400 elections on 4 candidates, of which 500 samples each were examined, were as follows. Changing the parameters of the distributions generating the voters’ scores made only small differences to these results, and none to the relative magnitudes. A bar-chart is shown in Figure 1 (page 33).

Puzzling over why there was not more difference in the immoderation statistic between RD and MEV0, we experimented with other ways of dis-

tributing the voters' mean scores. It turned out that while we could never get the immoderation of MEVO to exceed that of RD, or be below that of MR (except when assessed on Winner'sScore only), there are scenarios where there are much bigger differences in immoderation. One such is shown in Figure 2 (page 34). Here, voters' opinions were arranged such that there was a high probability of voters either viewing candidates 1 and 3 as much better than candidates 2 and 4, or viewing candidates 2 and 4 as much better than candidates 1 and 3; such a situation occurs in real life where both candidates and voters are distributed on different ends of the political left to right axis. Again, the details of how voters' opinions were distributed are in the appendix section 10.

Thus we see that in both scenarios both MEVO and (to a lesser extent) RD cause vastly less unfairness than the Majority Rule system. In Transmoggia this should lead to less social unrest. The reductions in average satisfaction are very small compared with the large benefits obtained by reducing unfairness while avoiding immoderation.

4 Tactical voting under MEVO

MEVO does not obey CVO and CVT (proof not given). It is therefore possible that voters interested primarily in getting a particular ordering as the result of an election, or more likely, interested in getting a particular candidate top of the ordering, may be able to gain by voting other than their true opinion. What MEVO does guarantee is that they cannot by tactical voting increase the probability of their favoured candidate being above any other specified candidate; the scope for gain by tactical voting is therefore likely to be fairly limited.

For example, if somebody desiring $A > B > C$ as the ordering of candidates A, B, and C, and who especially desires that A should come top, knows that B is the most popular candidate, they could consider voting $A > C > B$ instead. This would slightly increase the probability of A coming first – but it would also make it more likely that C will come first, and C is this voter's least favoured candidate. What it will not do is make any difference to the probability that A will beat B or the probability that A will beat C.

The weaknesses of MEVO in this regard are likely to be more prominent in situations where only one candidate is being elected from a constituency, as opposed to all the candidates being elected in some order, or several candidates being elected. One

might think that elections where only one candidate is being elected would be better addressed by considering the election output to be non-strongly ordered with the top candidate $>$ all the others, and all the others equal to each other, and then requiring RP to apply. To get the benefits of MEVO one would however have to interpret the voters' orderings in the *original* way. This distinction unfortunately leads to the RD distribution not satisfying the new RP condition, and to elections where there is indeed *no* output distribution possible that satisfies this new version of RP.

5 Implementation of MEVO

So far we have discussed the theoretical basis of MEVO and its benefits and drawbacks in various situations. We next turn to how the necessary calculations can actually be carried out in practice. While for RD essentially the only issue is how to choose a voter uniform-randomly from the population of voters, with MEVO we have a significantly more difficult problem.

We suggest two usable approaches. Neither is perfect and there is plenty of scope for better methods of implementation to be developed. Both are presented as a rough verbal description rather than as precise mathematics. Software that carries out each of these implementations (in the Matlab language) can be downloaded from the directory <http://www.inference.phy.cam.ac.uk/sewell>.

As above, let N denote the number of candidates.

5.1 Implementation for low N

The first approach is to calculate the distribution u_1 explicitly. This means calculating the probability under u_1 of each possible output ordering of the candidates. Since there are $N!$ such orderings, this is a calculation that will necessarily take (at the very least) $N!$ operations. Since $20!$ is about 2,432,902,008,176,640,000, it can be seen that this approach will take rather a long time for a 20-candidate election. However, where the number of candidates is under about 7, such an approach is feasible. The definition of MEVO leads, via the Lagrange multiplier technique, to a set of non-linear simultaneous equations on the $N!$ probabilities to be determined, plus some non-negativity constraints. The non-negativity constraints are usually redundant in practice, as the value of u_1 is only zero on those orderings which give a pairwise comparison favoured by zero of the voters. These orderings can be eliminated at the start; elsewhere the

gradient of the entropy becomes infinite as any individual probability approaches zero, pointing towards that probability being positive.

The problem is then one of solving the set of non-linear simultaneous equations, and as such a range of techniques is available in the literature. One solution is embodied in `sim1.m` on the web site mentioned above.

Once we are at a point where the constraints are satisfied and the gradient of the entropy is normal to the set of points satisfying the constraints, that point is the output distribution. Since it is a discrete distribution we can take a random sample from it by inverting the cumulative distribution function, and using a uniform random variate. The random sample thus taken is an ordering on the candidates, which we deliver as the output ordering.

5.2 Implementation for larger N

As discussed in section 5.1, the above technique is excessively computationally intensive when the number of candidates rises above about 7. Under these circumstances we must resort to a different method.

By considering a Lagrange-multiplier solution to the relevant constrained maximisation problem, u_1 may be shown to be of the following form:

$$u_1(t) = Ke^{\left(\sum_{i,j:i>j} \lambda_{i,j} d_{i,j}(t) \right)}$$

where K is a constant and i and j index the candidates. Therefore, if we know the values of the $\lambda_{i,j}$, we may sample from u_1 using a Markov Chain Monte Carlo algorithm, (e.g. Metropolis-Hastings, using proposal distributions that simply interchange two candidates in the current sample of t).

A possible approach, then, is to initialise the $\lambda_{i,j}$ to random values, run such an MCMC algorithm yielding a pile of non-independent samples, assess for each pair (i, j) whether the current fraction of the samples in which $c_i > c_j$ is too big or too small, and adjust each $\lambda_{i,j}$ upward if the fraction needs increasing or downwards otherwise. Such an iterative approach will eventually converge approximately (under the assumptions that each component MCMC run is ‘sufficiently long’, the step size for adjusting the $\lambda_{i,j}$ is ‘sufficiently small’, etc). It is possible to assess how accurately the constraints are currently being met at each point in the run. Nonetheless, it would be better to find a non-iterative perfect sampling system (for example one

using Fill’s algorithm); we have so far not been able to.

It turns out that the choice of proposal distribution used is important. If there is inadequate mixing, such a scheme does not converge. Software implementing the best approach that we know of is available in `sim2.m` in the same directory on the web (see section 5 above).

The limiting factor that governs convergence of sampling in such an approach seems to be that one requires a large number of samples of the ordering at any one of values of the $\lambda_{i,j}$ to get accurate estimates of the fraction of samples preferring one candidate to another. However good a proposal distribution is used, it would seem that an MCMC approach with feedback to the $\lambda_{i,j}$ will always have running speed limited in this way.

In practice such a system has been developed and tested for up to 40 candidates. If such an approach were to be used in practice it would be necessary to set precise criteria for when convergence could be considered adequate.

6 How could one set about introducing such a system?

In the grand scheme, the MEVO system itself (and the RD system likewise) could be implemented in two parts: the election of local candidates to be members of parliament, and the application of MEVO (or RD) to parliamentary procedure. It would probably be best to introduce the election to parliament first, reserving the somewhat more difficult procedural issues until experience had been gained in electing the members.

However, it is clear that substantial education of the population on the benefits would be necessary, and before any public election, suitable trials on smaller and more restricted elections would be needed. Such smaller elections could be surveyed to assess the real satisfaction of voters with the different systems, which might help the public to accept the necessity of randomisation to achieve fair elections.

7 Discussion

We have thus seen that probabilistic voting systems (both “decision schemes” and “social welfare functions”) can reduce the unfairness to minorities that occurs with majority rule. We have seen how the impasse of Arrow’s theorem may be circumvented by such systems.

We assumed as obvious that we must have symmetry among both voters and candidates and universal domain, and that our system should respect unanimous opinions (i.e. obey the weak Pareto condition). We then saw (from Gibbard's and McLennan's results) that addition of strategy-proofness immediately restricts us to using the Random Dictator system (at least as far as the induced decision scheme goes), which has the serious drawbacks that it is completely unable to compromise and can be very immoderate in its results.

We therefore chose instead to add the much weaker axiom of representative probability (RP) and to output (ultimately) an ordering on the candidates rather than just the name of the top candidate, even if our real aim is to elect only one candidate. Implicit in doing so is the decision that some degree of satisfaction will be afforded to voters by their desired candidate coming e.g. 2nd rather than 3rd, even if only one candidate is being elected (even though one of the authors (IM) has argued strongly against this in the past). Given RP we are guaranteed also pairwise clarity of voting (CVP), but from Gibbard's results we know that this is nowhere near as strong an axiom as strategy-proofness, and we have also made it clear that this does not imply clarity of voting on the top position or on the output ordering.

In choosing between the many possible systems which obey just these axioms (SV, SC, UD, RP), we concentrated on the one that offers the least element of surprise in the results given any candidate-symmetric prior beliefs, and saw how this is the one that minimises the information taken from the votes and maximises the entropy of the output distribution. This system is MEVO, the basic "maximum entropy voting" system.

MEVO was shown in experimental simulations to provide very much less unfairness than majority rule while diminishing overall average satisfaction very little. Inevitably any probabilistic system must increase micro- and macro-variation compared with any deterministic system (which has none). Similarly a probabilistic system would be expected to increase immoderation when compared with majority rule (although in fact this is not true if "Winner'sScore" is used as the satisfaction measure); nonetheless MEVO causes much less increase in immoderation than Random Dictator does.

Neither RD nor MEVO guarantees to elect a Condorcet winner. We have however seen from examples that MEVO is usually much more likely than RD to do so where one exists. What both do guarantee is that to the extent that the Condorcet winner wins unanimously, to that extent also he will be

placed higher than the others in probability.

The one real weakness of MEVO that we are aware of is the ability of a political party to increase the chances of the top elected candidate belonging to itself by flooding the candidate list with lots of its own candidates – i.e. the difficulty of "candidate loading". Note, however, that this does not mean that any one of these candidates has any favour compared with any other candidate, from that party or otherwise. The difficulty of combating this problem lies largely in the difficulty of detection of "membership" of a party, as this may not be formal (e.g. membership of the "party" of those who have lots of spare time). There is also an argument to say that if there are more people of one persuasion willing to give up their time to politics then they should each be given their fair chance. The main reason that RD does not suffer from candidate loading is that it avoids all compromise – and we believe that avoiding compromise is bad. We hope to publish a future paper discussing in detail the possibilities of ameliorating the issue of candidate loading.

A major difficulty with MEVO is that it is hard for the average voter to understand. However, in general it is not necessary for the voter to understand more than that he should place the candidates in order according to his true beliefs. Indeed, the fact that the system is hard to understand should be a strong disincentive to tactical voting, as the effects of tactical voting will be very difficult to predict (and of course it cannot alter the pairwise outcome probabilities anyhow).

Discussion of alternative electoral systems should of course also consider some of the other long-standing attempts to do better than majority rule, for example Single Transferable Vote (STV). It is the authors' hope that consideration of the two candidate scenario, as in sections 2.1 and 2.3 above, will suffice to convince the reader that without recourse to a probabilistic system one cannot avoid the inherent unfairness of majority rule, whichever of the other deterministic systems one may adopt.

Others may object that MEVO is likely to elect middle-of-the-road candidates and avoid any firm leadership (as is also claimed against most other ways of avoiding majority rule). Only real experiment and time will show whether this is true, and any system which promotes compromise can be criticised in this way – if you don't like compromise, then use Random Dictator, and take a small risk of getting a few years of extremist rule! However it is our hope that use of MEVO would lead to a need to reach agreement by genuine discussion that considers the needs of *all* parties, before voting, to a

greater extent than under majority rule. Only then can one be reasonably sure of what the outcome of the vote will be.

Finally, since the readership of *Voting matters* are particularly familiar with STV systems, we comment on the relationship between our RP axiom and the concept of Droop Proportionality (Woodall 1994):

Droop proportionality criterion (DPC): If there are $|V|$ voters and an election is to elect M of the available N candidates then we define the Droop Quota to be $\frac{|V|}{M+1}$, and require that for any $k, m \in N$ with $0 < k \leq m$, and for any subset C_0 of the set of candidates C , if more than $k \frac{|V|}{M+1}$ voters place all members of C_0 in one of the top m places in their ranking, then at least k members of C_0 should be elected.

The DPC is a very different concept of proportionality to RP. To see this, consider the case $|V| = 2, m = 1, k = 1$ (i.e. a 2-candidate election to elect one candidate). Then any election satisfying the DPC must elect whomever gets more votes. In other words DPC requires fallback, in the 2-candidate situation, to the MR system – totally different from the RP requirement that the probability of electing each of the two candidates should be proportional to the number of voters preferring that candidate.

8 References

- Amar, A.R. (1984). *Choosing Representatives by Lottery Voting*. Yale Law Journal 93: sections 1283-1308.
- Arrow, K.J. (1951). *Social Choice and Individual Values*. New York; John Wiley and Sons; 2nd edition 1963.
- Barbera, S. (1979). *Majority and positional voting in a probabilistic framework*. Review of Economic Studies, 46: 379-389.
- Burnheim, John (1985). *Is Democracy Possible? The alternatives to electoral politics*. Cambridge: Polity Press.
- Coughlin, P.J. (1992). *Probabilistic Voting Theory*. Cambridge: Cambridge University Press.
- Gibbard, A (1977). *Manipulation of schemes that mix voting with chance*. Econometrica, 45: 665-681.
- Guinier, Lani (1994). *The tyranny of the majority: fundamental fairness in representative democracy*. New York; London: Free Press.
- Hare, Thomas (1873). *The Election of Representatives, parliamentary and municipal*. 4th ed. London: Longmans, Green.
- Horowitz, Donald L. (1991). *A democratic South Africa?: constitutional engineering in a divided society*. Berkeley; Oxford : University of California Press.
- Howard, J.V. (1992). 'A Social Choice Rule and Its Implementation in Perfect Equilibrium', *Journal of Economic Theory* 56(1): 142-59.
- McLennan, A. (1980). 'Randomized preference aggregation: additivity of power and strategy proofness'. *Journal of Economic Theory*, 22: 1-11.
- Pattie C.J., Johnston R.J., Fieldhouse E. (1994). 'Gaining on the swings — the changing geography of the flow-of-the-vote and government fortunes in British general-elections, 1979-1992', *Regional Studies* 28 (2): 141-154.
- Reilly, Ben (2002). 'Electoral systems for divided societies', *Journal of Democracy* 13 (2): 156-170.
- Satterthwaite, M.A. (1975). 'Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions', *Journal of Economic Theory* 10: 187-217.
- Wichmann, B.A. (2009). Personal communication.
- Woodall, D.R. (1994). 'Properties of preferential election rules', *Voting matters* 3: 8-15.

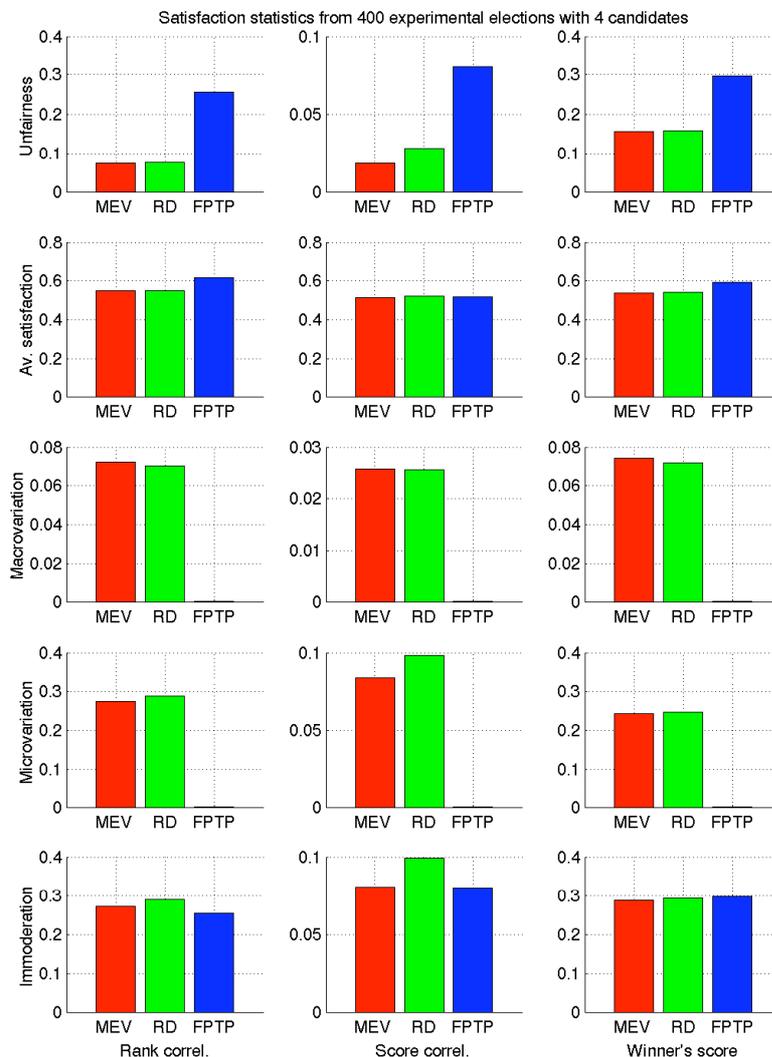


Figure 1: Satisfaction statistics in simulated elections. Of the three MR is most unfair and least variable, while MEV0 is least unfair, is less immoderate and microvariable than RD, and has very similar immoderation to MR (here labelled FPTP).

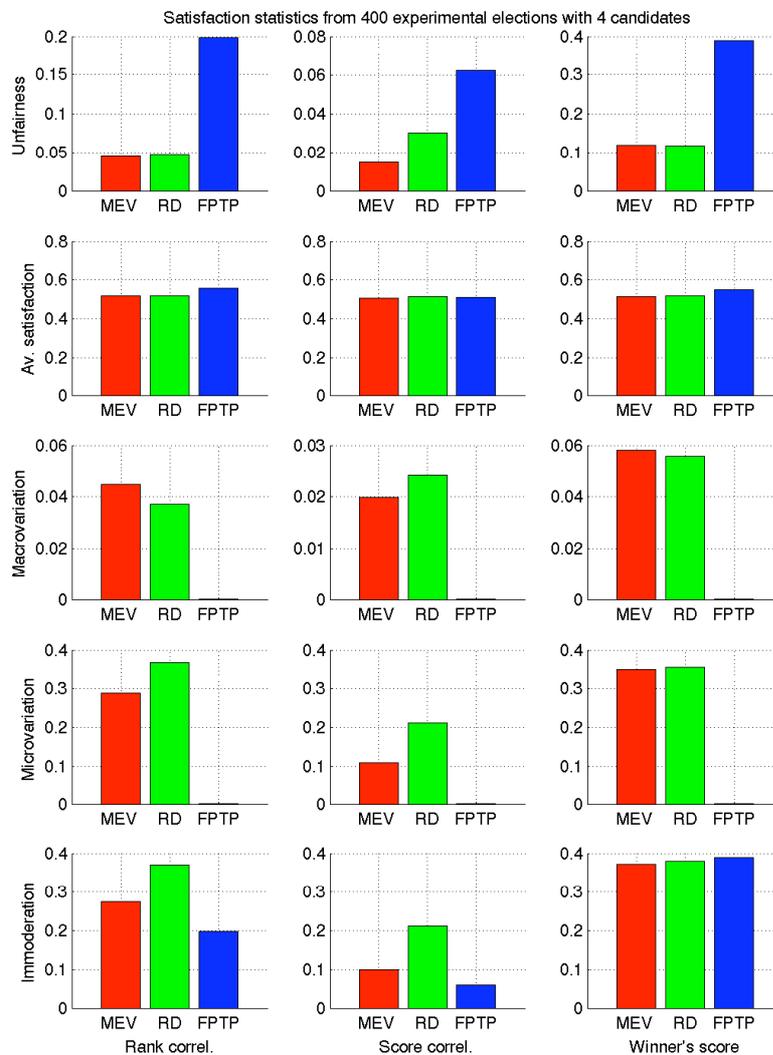


Figure 2: An election similar to that of Figure 1, but where there is a correlation between voter groups and candidates along an (e.g. left-right) axis.

9 Appendix - Formal description of Maximum Entropy Voting

In this appendix we describe the axioms and the MEV0 system formally.

9.1 Definitions of the axioms

Let C denote the (finite) set of Candidates, and N the number of candidates.

Let V denote the (finite) set of Voters, endowed with the uniform probability measure P_V .

Let T denote the set of strong total orderings⁶ on C , and W the set of total orderings on C . Small letters will denote members of the sets denoted by capitals.

Let X denote the set of possible sets of votes by the voters. Then $X = W^V$.

Let U denote the set of probability distributions on T , and D denote the set of probability distributions on C .

Let Ω denote the underlying probability space (think of this as the set of possible seeds for the random number generator, of very large uncountable size), and P_Ω the associated probability measure (which we will denote simply P when there is no ambiguity).

Let J denote the set of permutations of C , and $\rho : J \rightarrow W^W$ be the function such that for all $c_1, c_2 \in C$ and any $w \in W$, $c_1 \leq_{\rho(j)(w)} c_2$ iff $j(c_1) \leq_w j(c_2)$.

We define a probabilistic decision scheme (PDS) by abuse of notation as either a measurable function $f : X \rightarrow D$ or as a measurable function $f : X \times \Omega \rightarrow C$; it will always be clear which is meant by the number of arguments. We define a probabilistic social welfare function (PSWF) by abuse of notation as either a measurable function $f : X \rightarrow U$ or as a measurable function $f : X \times \Omega \rightarrow T$ (similarly). In both cases we will only be interested in schemes/functions that additionally meet certain axioms yet to be defined. (We assume in both cases that all subsets of X are considered measurable.)

We say that a PSWF $f : X \times \Omega \rightarrow T$ induces a PDS $f' : X \times \Omega \rightarrow C'$ on a subset C' of C

⁶ We will say that t is a "total" ordering iff

$$(\forall c_1, c_2 \in C) (c_1 \geq_t c_2 \text{ or } c_2 \geq_t c_1)$$

(i.e. for any pair of candidates either one is above the other or it is below the other, but can't be unrelated to the other; in contrast to a "partial" ordering).

We will say that t is a "strong" ordering iff

$$(\forall c_1, c_2 \in C) ((c_1 \geq_t c_2 \text{ and } c_2 \geq_t c_1) \Rightarrow (c_1 = c_2))$$

(i.e. the ordering can't rank two distinct candidates as equal).

iff for all $x \in X$, $\omega \in \Omega$, and $c' \in C'$, we have $f'(x, \omega) \succeq_{f(x, \omega)} c'$.

For the purpose of defining the SP axioms we define a utility function as a function from C, D, T , or U to R (the real line). A utility function $g : D \rightarrow R$ is defined to be risk-neutral if there exists a utility function $h : C \rightarrow R$ such that for all $d \in D$, $g(d) = E_{\omega \in \Omega, c \sim_d h} h(c) = \sum_{x \in C} h(c) d(c)$, i.e. if the value of $g(d)$ depends only on the mean of the utility of c under h when c is distributed according to d . If so, then g will be said to be the utility function on D induced by h . In an exactly similar way we define the concept of a risk-neutral utility function on U . We define all utility functions on C and T to be (vacuously) risk-neutral.

Given this background the axioms are then defined for a probabilistic social welfare function f as follows:

SV: For any permutation j of V and any $x \in X$, $f(x) = f(x \circ j)$.

SC: For any $j \in J$, any $x \in X$, and any $\omega \in \Omega$, $f(\rho(j) \circ x, \omega) = \rho(j) \circ f(x, \omega)$.

UD: (this is automatically met by any PSWF defined as above).

RP: For any $c_1, c_2 \in C$, any $x \in X$, then

$$P_\Omega (c_1 \leq_{f(x, \omega)} c_2) = P_V (c_1 <_{x(v)} c_2) + \frac{1}{2} P_V (c_1 =_{x(v)} c_2),$$

where $<_w$ and $=_w$ have the obvious meanings derived from the (not necessarily strong) ordering $w \in W$ written \leq_w .

IIA: For any $c_1, c_2 \in C$, and [for any $x_1, x_2 \in X$ such that for all $v \in V$, $c_1 \leq_{x_1(v)} c_2 \Leftrightarrow c_1 \leq_{x_2(v)} c_2$], then $P_\Omega (c_1 \leq_{f(x_1, \omega)} c_2) = P_\Omega (c_1 \leq_{f(x_2, \omega)} c_2)$.

WP: For any $c_1, c_2 \in C$ and any $x \in X$, $((\forall v \in V) (c_1 \leq_{x(v)} c_2) \Rightarrow P_\Omega (c_1 \leq_{f(x, \omega)} c_2) = 1)$.

CVP: For any $c_1, c_2 \in C, v \in V$, and [for any $x_1, x_2 \in X$ such that [for all $v' \in V \setminus \{v\}$ and all $c'_1, c'_2 \in C$, $c'_1 \leq_{x_1(v')} c'_2 \Leftrightarrow c'_1 \leq_{x_2(v')} c'_2$],

$c_1 \leq_{x_1(v)} c_2 \Rightarrow P_\Omega (c_1 \leq_{f(x_1, \omega)} c_2) \geq P_\Omega (c_1 \leq_{f(x_2, \omega)} c_2)$.

CVT: For any $c_1 \in C, v \in V$, and [for any $x_1, x_2 \in X$ such that [for all $v' \in V \setminus \{v\}$ and all $c'_1, c'_2 \in C$, $c'_1 \leq_{x_1(v')} c'_2 \Leftrightarrow c'_1 \leq_{x_2(v')} c'_2$],

$$\begin{aligned} & ((\forall c_2 \in C) (c_1 \geq_{x_1(v)} c_2)) \\ & \Rightarrow P_\Omega ((\forall c_2 \in C) (c_1 \geq_{f(x_1, \omega)} c_2)) \\ & \geq P_\Omega ((\forall c_2 \in C) (c_1 \geq_{f(x_2, \omega)} c_2)). \end{aligned}$$

CVO: For any $t \in T, v \in V$, and [for any $x_1, x_2 \in X$ such that [for all $v' \in V \setminus \{v\}$ and all $c'_1, c'_2 \in C$,

$c'_1 \leq_{x_1(v')} c'_2 \Leftrightarrow c'_1 \leq_{x_2(v')} c'_2$],

$x_1(v) = t \Rightarrow P_\Omega (f(x_1, \omega) = t) \geq P_\Omega (f(x_2, \omega) = t)$.

SP (for PDSs (also denoted f)): For any risk-neutral utility functions $g : D \rightarrow R, h : C \rightarrow R$ such that g is induced by h , and for any $v \in V$, and [for any $x_1, x_2 \in X$ such that [for all $v' \in V \setminus \{v\}$ and all $c_1, c_2 \in C, c_1 \leq_{x_1(v')} c_2 \Leftrightarrow c_1 \leq_{x_2(v')} c_2$]],
 $((\forall c_1, c_2 \in C) (c_1 \geq_{x_1(v)} c_2 \Leftrightarrow h(c_1) \geq h(c_2)))$
 $\Rightarrow g(f(x_1)) \geq g(f(x_2))$.

SP1 (for PSWFs): For any risk-neutral utility functions $g : U \rightarrow R, h : T \rightarrow R$ such that g is induced by h , and for any $v \in V$, and [for any $x_1, x_2 \in X$ such that [for all $v' \in V \setminus \{v\}$ and all $c_1, c_2 \in C, c_1 \leq_{x_1(v')} c_2 \Leftrightarrow c_1 \leq_{x_2(v')} c_2$]],

$$((\forall t \in T) (h(x_1(v)) \geq h(t))) \Rightarrow g(f(x_1)) \geq g(f(x_2)).$$

SP2 (for PSWFs): For any non-empty $C' \subseteq C$ the PDS f' induced by f on C' satisfies SP.

SLA: f is a PSWF that satisfies SV, SC, and RP.

9.2 Definition of Maximum Entropy Voting

We use the same notation as in section 9.1.

Each voter casts his vote by giving a total ordering (not necessarily strong) on C .

Let $s_{i,j}$ be the fraction of the votes that prefer c_i to c_j , where $c_i, c_j \in C$; a voter that rates two candidates equally preferable counts as half a vote in each direction, and where $i = j$ we will set $s_{i,j} = \frac{1}{2}$. Let S denote the matrix composed of the $s_{i,j}$.

Now let T be the set of strong total orderings on C .

Then for each $t \in T$ we define pairwise preference variables $d_{i,j}(t)$ which are 1 if c_i is preferred to c_j under t (i.e. if $c_i >_t c_j$), $\frac{1}{2}$ if $i = j$, and 0 otherwise.

Now let U be the set of probability distributions on T . Then for each $u \in U$, we define a matrix $D(u)$ whose elements $d_{i,j}(u)$ are the probability that under a random ordering t drawn from u candidate c_i is preferred to candidate c_j , and which are given by

$$d_{i,j}(u) = \sum u(t) d_{i,j}(t).$$

The map B that maps u to $D(u)$ is linear⁷.

Now let $B^{-1}(\{S\})$ denote the subset U_0 of U such that $u \in U_0$ iff $u \in U$ and $Bu = S$. U_0 is thus the set of distributions on T that satisfy the constraints imposed by RP. U_0 is convex since it is the

⁷ i.e. for any two non-negative real numbers α and β which sum to 1, and any $u_1, u_2 \in U$,

$$B(\alpha u_1 + \beta u_2) = D(\alpha u_1 + \beta u_2) = \alpha D(u_1) + \beta D(u_2) = \alpha B(u_1) + \beta B(u_2).$$

intersection of a convex set with the inverse image of a single-point set under a linear map.

We next want to establish that $B^{-1}(\{S\})$ is non-empty, and to that end will exhibit an element u_0 in it. Let u_0 denote the distribution on T defined as follows: $u_0(t)$ is the probability that a randomly chosen voter orders the candidates according to t . (Voters who vote with a non-strong ordering are divided up into several equally-weighted fractional voters who vote with each of the contributory strong orderings.) (Intuitively u_0 is the output distribution of the RD system.)

Then $Bu_0 = S$, therefore $U_0 \neq \emptyset$.

Now consider the function $f : U \rightarrow R$ defined by

$$f(u) = -\sum_{t \in T} u(t) \log(u(t))$$

f gives the entropy of a distribution u on T . Then f is readily verified to be a strictly concave function on U . Moreover the gradient of f is infinite in magnitude wherever any point value $u(t_1)$ of u approaches zero, and is then pointing away from the surface $u(t_1) = 0$. Therefore f has a single maximum on any convex subset of U , and in particular on U_0 .

Let

$$u_1 = \operatorname{argmax}_{u \in U_0} f(u),$$

i.e. u_1 is the distribution on T of maximum entropy which satisfies $Bu_1 = S$.

We will then define the distribution of outcomes from the MEV0 system to be u_1 ; to obtain the actual ordering that is the outcome, we draw a random sample t_1 from u_1 .

9.3 Axioms adhered to

MEV0 satisfies the following axioms:

SV, SC, UD, WP, IIA, ND, RP, CVP.

Moreover MEV0 uses only the information in S (the pairwise preference probability matrix) from the votes, and no other. Under MEV0, and any candidate-symmetric prior on the votes, the information about the output ordering in the votes is as low as is possible for a system consistent with RP. Moreover it is the unique method with this overall combination of properties.

Stated slightly differently, MEV0 is the unique system meeting SLA which uses only, and all of, the pairwise preference probabilities from the votes, and produces an output distribution on the orderings containing as little information about the output ordering consistent with the preceding properties.

9.4 Proofs

SV, SC, UD are obvious.

RP holds because $Bu_1 = S$.

RP implies CVP, since if we wish to maximise the probability that the result prefers candidate c_i to candidate c_j , we need to ensure that $(Bu_1)_{i,j} = s_{i,j}$ is maximised. But $s_{i,j}$ is the fraction of the voters that prefer c_i to c_j , and all that we can do to maximise that is ensure that we vote c_i above c_j . RP also implies IIA and ND trivially.

To prove WP, note that if all candidates prefer c_i to c_j , then $s_{i,j} = 1$, and hence almost all t_1 (i.e. all but some with total probability 0 under u_1) prefer c_i to c_j , so the probability of the result preferring c_i to c_j is 1.

It is clear that no information from the votes is used other than S , since the only information derived from the votes that enters the system is S . On the other hand, since S is deducible from u_1 , all the information in S is also in u_1 .

Nonetheless the number of bits of information that the votes contain about the ordering, $I(T; V|y)$, given the choice of voting system y , takes a value defined by the relevant probability distributions, and the equation

$$I(T; V|y) = \int P(t, v|y) \log_2 \frac{P(t|v, y)}{P(t|y)} d(t, v),$$

where we envisage the probabilistic model relating the variables that is shown in Figure 3. Moreover, despite the fact that all the information about T comes through S , $I(T; V|y)$ does vary with y .

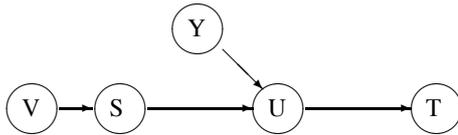


Figure 3: Relationships between the variables at an election.

Now, we have some prior distribution on the votes $P(v)$ which is candidate-symmetric, and some fixed choice of system y to be determined, from the set Y of systems that obey SLA and take only the pairwise probability matrix S from the votes. For any fixed v there is only one possible value of S , and for any fixed value of S and any choice of system $y \in Y$ there is only one possible value of u . We assert that MEV0 is the system \check{y} given by

$$\check{y} = \operatorname{argmin}_{y \in Y} I(T; V|y).$$

Indeed we have

$$\begin{aligned} \check{y} &= \operatorname{argmin}_{y \in Y} \int P(v) P(t|v, y) \log_2 \frac{P(t|v, y)}{P(t|y)} d(t, v) \\ \check{y} &= \operatorname{argmin}_{y \in Y} \left(\int P(v) P(t|v, y) \log_2 P(t|v, y) d(t, v) \right. \\ &\quad \left. - \int P(v) P(t|v, y) \log_2 P(t|y) d(t, v) \right) \end{aligned}$$

But since $P(v)$ has been assumed to be candidate-symmetric, and y obeys SLA, $P(t|y)$ is a constant varying neither with t, v , nor y . Moreover $P(t|v, y)$ integrates over t to 1 for any fixed v and y (as it is a probability distribution), and $P(v)$ integrates to 1 for the same reason. Therefore the second term above does not vary with y , and we have

$$\begin{aligned} \check{y} &= \operatorname{argmin}_{y \in Y} I(T; V|y) \\ &= \operatorname{argmin}_{y \in Y} \int P(v) P(t|v, y) \log_2 P(t|v, y) d(t, v) \\ &= \operatorname{argmax}_{y \in Y} \int P(v) H(T|v, y) dv \end{aligned}$$

where $H(T|v, y) = - \int P(t|v, y) \log_2 P(t|v, y) dt$, the entropy of the output distribution. But our choice of the MEV0 system for y uniquely maximises $H(T|v, y)$ for any fixed v , and hence uniquely maximises $\int P(v) H(T|v, y) dv$ no matter what candidate-symmetric distribution $P(v)$ may be.

That concludes the proofs of the formal properties of MEV0. Illustrations of how it differs from RD in practice are in sections 2.9, 3, and 4 above.

10 Appendix – Simulation of voters’ opinions

For the purpose of the simulated elections of section 5 above, we simulated voters’ opinions as follows. We describe first the set of elections whose results are shown in Table 10 and Figure 1, and then indicate how the second set of elections (of Figure 2) differs.

We created eight voter interest groups. For each group we drew a mean score x_c on each of the four candidates independently from a Beta distribution with parameters 0.5, 0.5. For each group we also drew independently a single strength of opinion value y from a Gamma distribution with parameters $m = 2, r = 0.1$.

We assigned each of 50 voters randomly to one of the eight voter interest groups. Each voter’s score on each candidate c was then drawn independently from a Beta distribution with parameters $yx_c, y(1 - x_c)$, with x_c and y being the parameters appropriate to that voter’s voter interest group. Thus the mean score for a voter on candidate c was the appropriate group’s x_c , and when y was large the score distribution was tighter around x_c than when y was small.

The overall effect is therefore that the voters' opinions are clustered in eight clusters of varying tightness, with both tightness and position of the clusters in opinion space varying from one simulated election to another. In this first election the position of the clusters, though tending to be nearer edges and corners than in the centre of the 4-d hypercube, are in a sense unprincipled – i.e. there is no link between the opinions of voters in different groups.

For the second set of elections, instead of using a Beta(0.5, 0.5) distribution for each x_c , we used either a Beta(2, 0.5) or a Beta(0.5, 2) distribution, according to whether the parity (oddness or evenness) of the candidate number matched or did not match respectively that of the voter group number⁸. Thus we may think, for example, of the odd numbered candidates as being “right-wing” and the even numbered candidates as being “left-wing”, and similarly for the voter groups; left wing voters are unlikely to think highly of right wing candidates (and vice versa). As we have seen, this set of elections brings out to a rather greater extent how MEVO ameliorates RD's tendency to immoderation.

11 Appendix – Statistics used to describe profile of satisfaction with results of simulated election

For each measure of satisfaction, the following statistics were calculated:

1. $E_{v \in V, m \in M, \omega \in \Omega} (a(v, m, \omega))$, where M is the set of elections, V the set of voters, Ω the set of seeds of the random number generator, and $a(v, m, \omega)$ the satisfaction of voter v in election m using random seed ω , and where E denotes expectation; we refer to this statistic as “**average satisfaction**”;
2. $E_{m \in M} (std_{v \in V} (E_{\omega \in \Omega} (a(v, m, \omega))))$, where also std denotes standard deviation; we refer to this statistic as “**unfairness**”;
3. $E_{m \in M} (E_{\omega \in \Omega} (std_{v \in V} (a(v, m, \omega))))$; we refer to this statistic as “**immoderation**”;
4. $E_{m \in M} (std_{\omega \in \Omega} (E_{v \in V} (a(v, m, \omega))))$; we refer to this statistic as “**macrovariation**”;
5. $E_{m \in M} (E_{v \in V} (std_{\omega \in \Omega} (a(v, m, \omega))))$; we refer to this statistic as “**microvariation**”.

⁸ The Beta(0.5, 0.5) distribution has two equal peaks near 0 and near 1 as well as a lesser amount of probability in the middle, while the Beta(0.5, 2) distribution is missing the peak near 1 and the Beta(2, 0.5) distribution is missing the peak near 0.

12 Biographical Sketches

Roger Sewell is at Cambridge Consultants Ltd, a technical design consultancy, and specialises in developing tractable signal processing solutions to complex problems. He is particularly interested in applications of probability, including Bayesian inference, to a wide variety of situations. He worked previously as a hospital doctor specialising in renal medicine, gaining his DM for work on mechanisms of charge modification in minimal change nephropathy.

David MacKay studied Natural Sciences at Trinity College, before heading off to the USA to complete a PhD in Computation and Neural Systems at Caltech. In 1992 he returned to Cambridge as a Royal Society research fellow at Darwin College. In 1995 he became a university lecturer in the Department of Physics, and in 2003 he was promoted to be Professor of Natural Philosophy. Prof. MacKay's research interests are centered on applications of information theory and Bayesian methods, including error-correcting codes and machine learning.

Iain McLean is Professor of Politics, Oxford University and a fellow of Nuffield College. He has held visiting appointments at Stanford, Yale and the Australian National universities. He has worked on the history of social choice and the properties of electoral systems. His relevant books include 'Condorcet: foundations of social choice and political theory' (with F. Hewitt, 1994) and 'Classics of Social Choice' (with A.B. Urken, 1995).