CDMTCS
Research
Report
Series

New measures of the difficulty of manipulation of voting rules

Reyhaneh Reyhani
Department of Computer Science, University of Auckland

Geoffrey Pritchard
Department of Statistics, University of Auckland

Mark C. Wilson
Department of Computer Science, University of Auckland

CDMTCS-357
March 2009

Centre for Discrete Mathematics and Theoretical Computer Science

New measures of the difficulty of manipulation of voting rules

Reyhaneh Reyhani
Department of Computer Science, University of Auckland
Private Bag 92019, Auckland, New Zealand
rrey015@aucklanduni.ac.nz

Geoffrey Pritchard
Department of Statistics, University of Auckland
Private Bag 92019, Auckland, New Zealand
geoff@stat.auckland.ac.nz

Mark C. Wilson
Department of Computer Science, University of Auckland
Private Bag 92019, Auckland, New Zealand
mcw@cs.auckland.ac.nz

March 24, 2009

Abstract

We introduce new measures of manipulability of anonymous voting rules and argue for their superiority over some commonly used measures. We give a simple common framework that describes these measures and connects them to recent literature. We discuss their computation and present numerical results that allow for comparison of various common voting rules.

1 Introduction

From the perspective of mechanism design, it is generally regarded as desirable to minimize the occurrence of manipulation of voting rules. Many authors have tried to quantify the probability of such an event, under various assumptions on the distribution of voter opinions (see Section 4 for detailed discussion of relevant literature). More recently the idea of using measures based on computational complexity has arisen.

Most existing measures of manipulability of voting rules appear to us to be rather crude and unrealistic. Successful manipulation of an election, even in the case considered in the present article when the manipulators are opposed only by naive, sincere voters, requires considerable computational effort. To assemble a manipulating coalition, we must discover the preference rankings of voters, convince them to join the coalition, compute their strategy, and enforce their implementation of that strategy. The probabilistic measures that have been used mostly ignore the size of the manipulating coalition altogether. On the other hand, approaches based on computational complexity have not been particularly fruitful. Manipulation of a voting situation typically requires only the evaluation of the result of a single alternative election; thus it is only as difficult as determining the winner of the election, and this is a polynomial time computation (in n for fixed m) for most rules. This makes it less interesting from the point of view of traditional complexity theory. Furthermore, results from both streams of research show that there is a phase transition in the probability of manipulability as the coalition size grows, yet say nothing about how to compare rules near that threshold.

1.1 Our contributions

We introduce (Section 2.2) several new general probabilistic measures of susceptibility to manipulation, and argue that they are superior to existing measures. We investigate their values (Section 3) for members of an important class of voting rules, namely the scoring rules. This is done for each of two standard probability models for voter preferences. We also investigate the relationship between these different measures and put them in a unified framework, relating them to previous work. We discuss the computation of these measures in terms of efficient algorithms (Section 6.1) and constant factor speedups (Section 6.2).

Our results suggest strongly that for the purpose of computing the relative manipulability of rules, for small coalitions it matters which measure we use, but for large coalitions any of them can be used. As a result of our analysis so far, we make several conjectures about the behaviour of these measures (Section 5) and outline possible future work.

1.2 Basic terminology

Consider a set $V = \{v_1, \ldots, v_n\}$ of agents (the *voters*) choosing from a given set $C = \{c_1, \ldots, c_m\}$ (the *candidates*). Each voter has an *opinion* or preference ranking (a complete strict linear ordering of the candidates); the list of these $(R_1, \ldots, R_n)$ forms the *sincere profile*. Each voter submits a linear ranking R'_i, which may or may not be the same as his sincere opinion, and this gives the *expressed profile*. For example, consider a set of 5 voters $V = \{v_1, \ldots, v_5\}$ and a set of 3 candidates $C = \{a, b, c\}$ with sincere profile $(abc, abc, bac, cab, bac)$. The expressed profile $(abc, abc, bca, cab, bac)$ results from the voter v_3 not voting sincerely, while the other voters do.

A (resolute) voting rule is an algorithm that takes an expressed profile as input, aggregates these inputs and returns a single candidate (the winner) as an output. We may need to use a tiebreaking protocol. A key feature of most voting rules is *anonymity*: the function value is unchanged if voters are permuted, so the rule treats voters equally. In this case the profile information can be represented more succinctly as a *voting situation*, where we simply list the numbers of voters with each of the possible opinions. For example, for three candidates (a, b, c), with the standard order $abc, acb, bac, bca, cab, cba$ of opinions, the 6-tuple $\sigma = (n_1, \ldots, n_6)$ represents a voting situation with n_1 voters having preference order abc, etc. In the example above, this succinct input for the expressed profile would be $\sigma = (2, 0, 0, 1, 1, 0)$.

A voter v may try to manipulate the election result by submitting an expressed order that differs from his sincere opinion, so as to gain an outcome that v prefers to the sincere outcome. The fundamental result of Gibbard [8] and Satterthwaite [14] implies that for anonymous rules, provided $m \geq 3$ some voter will have incentive to act in this way in some voting situation.

A common classes of voting rules is the *scoring rules*. A scoring rule is defined by a weight vector $w_1, \ldots, w_m$ with $w_1 \geq w_2 \geq \cdots \geq w_m$, and voter i gives score w_j to candidate k precisely when j is the kth highest in i's preference order. The candidate with highest total score wins. Scoring rules are anonymous provided the tiebreaking procedure is symmetric. The most commonly used scoring rules are listed below.

- Plurality rule, defined by the weight vector $(1, 0, 0, \ldots, 0)$;
- Borda's rule, defined by the weight vector $(m-1, m-2, \ldots, 1, , 0)$;
- Antiplurality rule, defined by the weight vector $(1, 1, \ldots, 1, 0)$.

2 Definition of the measures

We define four types of measures of manipulability of voting rules. All our measures are probabilistic and depend on a probability model for the distribution of opinions in the voter population. We consider in our numerical results only the two most obvious distributions: the uniform distribution on profiles (known as the Impartial Culture hypothesis) and the uniform distribution on voting situations (known as the Impartial Anonymous Culture hypothesis). However the definitions make sense for any distribution, and the method of computation is the same.

2.1 The model of manipulation

Fix a voting rule. To simplify analysis and for reasons of *neutrality* (symmetry between candidates) we use random tie-breaking, even though this strictly speaking does not define a resolute voting rule. If several candidates rare tied as winners, we choose one uniformly at random.

We define manipulability of a voting situation in stepwise fashion as in [11]. We make the assumption of risk-averse voters, who will prefer to vote sincerely unless they can improve the probability of success of some candidate without damaging the chances of any candidate whom they rank higher.

(i) Let π be a profile. We say that π' is *preferred* to π by voter v if for each $k = 1 \cdots m$ the probability of electing one of v's most-favored k candidates under π' is no less than under π. (If $\pi' \neq \pi$ the condition implies that this probability will be strictly greater for some k.)

(ii) A subset $X \subset V$ is a *manipulating coalition* at the profile π if and only if there is a profile $\pi' \neq \pi$ which agrees with π on $V \setminus X$ and is preferred to π by all members of X;

(iii) A rule is *manipulable at the profile* π if and only if there exists a manipulating coalition at this profile;

(iv) A rule is *manipulable at a voting situation* σ if and only if there exists a profile π giving rise to σ, at which the rule is manipulable.

Example 1 *(incentive to manipulate) Suppose that in profile π the sincere outcome is that a,c tie and in profile π' the sincere outcome is that b is the absolute winner. We want to see whether an abc prefers π or π'. As in profile π and π', for $k = 1$ the probability of a winning equals 0.5 and zero respectively, so π' is not preferred to π and also for $k = 2$ the probability of a or b winning equals 0.5 and 1 respectively, so π is not preferred to π' either.*

Example 2 *(manipulation) Consider the Borda rule, given by the weight vector $(2, 1, 0)$, and the voting situation with 2 abc, 2 bac, 2 bca, 3 cab voters. If one of the cab voters votes strategically as acb, then a and b tie. The new outcome is preferred by that voter because the winning probability distribution on the candidates has changed from $(0, 1, 0)$ to $(1/2, 1/2, 0)$.*

Example 3 *(manipulation in favour of bottom-ranked candidate) Consider the plurality rule, given by the weight vector (1, 0, 0), and a voting situation having 4 abc, 3 bca and 2 cab voters. The sincere winner is then a. By considering all possible combinations for manipulation, we find no manipulating coalition in favour of b since the only voters preferring b to a are already contributing the maximum score to b and the minimum to a. However, if the bca voters all vote strategically as cba, then c wins.*

Example 4 *(manipulation possible in more than one way) Consider an election with 3 abc, 2 cba, 2 bca voters, and scoring rule Plurality. The sincere scores are $s = (3, 2, 2)$. If both cba voters change their votes to bca in favour of b, we have a manipulating coalition with size 2 in favour of b. Also, we can consider a manipulating coalition with size 2 in favour of c. If both bca voters change their votes to cba, then the winner is c.*

2.2 Measures of manipulability

We first fix a number n of voters and a probability distribution on the types of voter opinions.

The first class of measures is obtained by considering the logical possibility of coalitional manipulation. Let I be the indicator random variable that takes 1 on a voting situation σ if and only if σ is manipulable, and 0 otherwise. Its expectation is simply the probability that a randomly chosen voting situation can be manipulated, which we denote by P. This simple measure has been used extensively in the social choice literature. It is relatively simple to compute, but fails to measure the computational effort required to assemble and coordinate a large coalition. It is entirely possible *a priori* that two rules may have the same value of P, yet one requires much more effort than the other in every voting situation.

The second class of measure is obtained by considering the minimum size of a manipulating coalition. For a voting situation σ that is manipulable, define M to be the minimum possible value of the size of a manipulating coalition. We would like to consider its expected value $E[M]$. However this does not make sense because certain situations are not manipulable by any coalition, and so the random variable M is *defective*. We could therefore consider $E[M \mid M < \infty]$. However *a priori* this may be rather small for rules that are almost never manipulable, and larger for rules that are often manipulable. More information is obtained by considering the distribution function of M. For each k with $1 \leq k \leq n$, we consider $\Pr(M \leq k)$. This is precisely the probability that a randomly chosen voting situation can be manipulated by k or fewer voters.

These measures take into account the increasing work required by larger coalitions. For example, the communication cost between coalition members may grow as M^2, if secret negotiations must be individually carried out. However, it is *a priori* possible that two rules may have the same value of M for every voting situation, yet one has very few small coalitions and one has many.

We also need to consider *minimal manipulating coalitions*. A manipulating coalition is minimal if no proper subset of the members can manipulate. In a minimal manipulating coalition, no member votes sincerely and all of them must act together in order to manipulate. Clearly every minimum size coalition is minimal, but the converse is not true in general. For example, consider the scoring rule with weights $(10, 9, 0)$ and a voting situation with 10 abc, 6 bac, 5 cab and 5 cba voters. The sincere result has the scores of a, b, c respectively being $s = (199, 195, 100)$. Consider manipulation in favour of b. If one of the bac voters changes to bca, the new scoreboard will be $s' = (190, 195, 109)$. So it is a minimal manipulating coalition of size 1, and clearly also minimum. By contrast, if 4 cba voters change their votes to bca, the new result will be $s' = (199, 199, 96)$. This is also a minimal manipulating coalition, yet not one of minimum size.

The third class of measure is obtained by considering the number of manipulating coalitions of a given size. For each voting situation σ and each k with $1 \leq k \leq n$, we define $N_\sigma(k)$ to be the number of subsets of size k that contain a manipulating coalition. The expectation over σ we denote by $N(k)$. A refinement is to consider the number of subsets $\overline{N}_\sigma(k)$ of size k, taken from the subset of voters that are unsatisfied with the sincere outcome (those who have an incentive to manipulate), and the corresponding expectation $\overline{N}(k)$.

The fourth and final class of measure is obtained by considering the informational effort required to discover a manipulating coalition. We assume that anonymity means that although a potential instigator of manipulation knows the distribution of opinions (in other words, the sincere voting situation), he does not know which voters hold which opinions. We assume that such a person must simply interview voters one by one at random, until he has enough people to form a manipulating coalition. For each voting situation, this gives a random variable equal to the number of such queries. Its expectation over the space of all query sequences for that voting situation, chosen uniformly at random, we denote by Q. A variant of this is the case where we only ask questions of unsatisfied voters, yielding a random variable $\overline{Q}$. These random variables are also defective, and are defined to be $+\infty$ if no manipulation is possible for the given voting situation.

We illustrate these definitions using the following example.

Example 5 *(computation of Q and $\overline{Q}$) Consider a setup with 2 voters, scoring rule Borda and 3 candidates. We want to calculate these two new measures Q and $\overline{Q}$ for both cultures.*

First, suppose we have no information about each voter's preference. We want to calculate the probability that a randomly chosen subset of k voters in a randomly chosen voting situation contains a manipulating coalition. We have 21 different voting situations, but by using symmetry, we need only consider those with $\mid a \mid \geq \mid b \mid \geq \mid c \mid$. It can be seen as in Section 6.1 that of these, only the voting situation $(1, 0, 1, 0, 0, 0)$ is manipulable. The sincere scoreboard is

$(3,3,0)$ *but if the* bac *voter changes his vote to* bca *then the result becomes* $(2,3,1)$. *Similarly, the* abc *can change to* acb. *Thus for this voting situation, we make 1 query with probability 1, and at most 2 queries with probability 1. Now we weight this voting situation according to the culture. For culture IAC, by using symmetry, the probability of voting situation will be* $3/21$. *So the values* $\Pr(Q \leq 1)$ *and* $\Pr(Q \leq 2)$ *are each* $1/7$. *For IC culture, the profile corresponds to* 6 *profiles by symmetry, so the weight is* $6/36$, *and the values of* $\Pr(Q \leq 1)$ *and* $\Pr(Q \leq 2)$ *are each* $1/6$.

Second, suppose we have partial information and know who is unsatisfied without knowing each voter's preference. We want to calculate the probability that a randomly chosen subset of k *or fewer unsatisfied voters can manipulate. Consider the above voting situation. Here, we have just the single* bac *voter unsatisfied, so we make 1 query with probability* 1. *Thus for IAC culture* $\Pr(\overline{Q} \leq 1)$ *and* $\Pr(\overline{Q} \leq 2)$ *are both* $1/7$ *and for IC both are* $1/3$.

2.3 Relations between the measures

A closer look reveals relationships between these measures, some of which have occurred implicitly in the literature. We present a common framework in which to view them.

If we normalize the measure $N(k)$ and $\overline{N}(k)$ by dividing by the number of possible k-subsets chosen from all voters (respectively, all unsatisfied voters) we are computing the probability that a randomly chosen subset of voters contains a manipulating coalition, in a randomly chosen voting situation. These quantities can be interpreted, respectively, as the distribution functions $\Pr(Q \leq k)$ and $\Pr(\overline{Q} \leq k)$. To see this for Q, note that we can consider all possible query sequences (permutations of voters) to be fixed in advance. Each subset of size k occurs with probability $\binom{n}{k}^{-1}$ as the initial subset of queries of a query sequence. Then $Q > k$ means precisely that this initial subset does not contain a manipulating coalition. Thus $\Pr(Q \leq k) = N(k)/\binom{n}{k}$. A similar argument works for $\overline{Q}$.

Finally, we can also consider a trivial query model for M. We assume that in this case we know all the voters' preferences, in other words the sincere profile, and our "query" consists of simply approaching a voter and inviting him to join a coalition (we assume that our invitations are never refused). We would make precisely M queries in order to minimize effort.

Thus the values of $\Pr(Q \leq k), \Pr(\overline{Q} \leq k)$ and $\Pr(M \leq k)$ correspond the probability that we can form a coalition after k queries in the case of no extra information (only the voting situation), partial extra information (who is unsatisfied), and full information (the complete profile), respectively.

3 Numerical results

We present numerical results on the above measures for various values of n and k, for both IC and IAC, and for a selection of voting rules. For details of the algorithms and implementation, see Section 6.

Tables 1 and 2 give values of $\Pr(Q \leq k)$ and $\Pr(\overline{Q} \leq k)$ for $1 \leq n$ (we include all values of k for which the next value of k gives a different answer, for some rule, for that value of n). Tables 3 and 4 present the analogous data for IAC. Similar tables for M can be found in [11].

Table 1: Values of $\Pr(Q \leq k)$ under IC.

n	k	Plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	Anti-plurality
2	1	0.0000	0.3333	0.1667	0.1667	0.1667	0.3333
2	2	0.0000	0.5000	0.1667	0.1667	0.1667	0.3333
3	1	0.0000	0.0000	0.1111	0.1667	0.1389	0.1111
3	2	0.0000	0.0000	0.1944	0.2222	0.2222	0.1111
3	3	0.0000	0.0000	0.2500	0.2500	0.2500	0.1111
4	1	0.1111	0.2083	0.1528	0.1759	0.1852	0.1481
4	2	0.2037	0.3519	0.2176	0.2917	0.3009	0.2222
4	3	0.2778	0.4583	0.2639	0.3657	0.3704	0.2685
4	4	0.3333	0.5417	0.2917	0.4028	0.4028	0.2963
5	1	0.0741	0.1173	0.1296	0.1620	0.1119	0.2099
5	2	0.1481	0.2099	0.2122	0.2662	0.1767	0.3148
5	3	0.2222	0.2901	0.2793	0.3465	0.2191	0.3580
5	4	0.2963	0.3611	0.3472	0.4120	0.2531	0.3688
5	5	0.3704	0.4167	0.4167	0.4630	0.2816	0.3750
6	1	0.0412	0.1260	0.1376	0.1472	0.1229	0.1070
6	2	0.0905	0.2168	0.2155	0.2423	0.2252	0.1523
6	3	0.1451	0.2946	0.2760	0.3230	0.3169	0.1677
6	4	0.2016	0.3629	0.3283	0.3969	0.3956	0.1718
6	5	0.2572	0.4218	0.3774	0.4623	0.4594	0.1741
6	6	0.3086	0.4707	0.4237	0.5163	0.5071	0.1754
7	1	0.1097	0.1301	0.1286	0.1411	0.1300	0.1331
7	2	0.1989	0.2201	0.2092	0.2349	0.2300	0.2202
7	3	0.2716	0.2940	0.2749	0.3088	0.3098	0.2805
7	4	0.3314	0.3608	0.3361	0.3711	0.3747	0.3242
7	5	0.3815	0.4236	0.3947	0.4257	0.4287	0.3563
7	6	0.4244	0.4829	0.4509	0.4746	0.4748	0.3791
7	7	0.4621	0.5379	0.5034	0.5187	0.5154	0.3948
8	1	0.0800	0.1197	0.1279	0.1388	0.0992	0.1669
8	2	0.1543	0.2114	0.2065	0.2369	0.1629	0.2668
8	3	0.2235	0.2869	0.2694	0.3181	0.2096	0.3248
8	4	0.2877	0.3507	0.3258	0.3905	0.2457	0.3572
8	5	0.3467	0.4051	0.3790	0.4562	0.2746	0.3751
8	6	0.3998	0.4520	0.4295	0.5154	0.2989	0.3855
8	7	0.4461	0.4922	0.4781	0.5674	0.3206	0.3925
8	8	0.4841	0.5250	0.5251	0.6106	0.3418	0.3970
9	1	0.0533	0.1107	0.1232	0.1258	0.1004	0.0995
9	2	0.1120	0.1889	0.2012	0.2089	0.1866	0.1536
9	3	0.1715	0.2555	0.2648	0.2772	0.2653	0.1819
9	4	0.2284	0.3167	0.3229	0.3389	0.3371	0.1965
9	5	0.2808	0.3744	0.3775	0.3964	0.4015	0.2043
9	6	0.3275	0.4289	0.4289	0.4507	0.4578	0.2089
9	7	0.3681	0.4804	0.4765	0.5021	0.5053	0.2119
9	8	0.4028	0.5296	0.5194	0.5502	0.5434	0.2138
9	9	0.4321	0.5785	0.5568	0.5941	0.5715	0.2150
10	1	0.1024	0.1156	0.1206	0.1233	0.1064	0.1194
10	2	0.1849	0.1975	0.1971	0.2082	0.1907	0.2037
10	3	0.2522	0.2648	0.2591	0.2781	0.2616	0.2662
10	4	0.3081	0.3240	0.3153	0.3402	0.3226	0.3142
10	5	0.3556	0.3773	0.3681	0.3969	0.3756	0.3519
10	6	0.3970	0.4259	0.4182	0.4492	0.4221	0.3815
10	7	0.4339	0.4704	0.4659	0.4976	0.4630	0.4044

Table 2: Values of $\Pr(\overline{Q} \leq k)$ under IC.

n	k	Plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	Anti-plurality
2	1	0.0000	0.5000	0.1667	0.1667	0.1667	0.3333
2	2	0.0000	0.5000	0.1667	0.1667	0.1667	0.3333
3	1	0.0000	0.0000	0.2222	0.2500	0.2222	0.1111
3	2	0.0000	0.0000	0.2500	0.2500	0.2500	0.1111
3	3	0.0000	0.0000	0.2500	0.2500	0.2500	0.1111
4	1	0.2222	0.3333	0.2083	0.2685	0.2685	0.2454
4	2	0.3333	0.5324	0.2824	0.3843	0.3843	0.2963
4	3	0.3333	0.5417	0.2917	0.4028	0.4028	0.2963
4	4	0.3333	0.5417	0.2917	0.4028	0.4028	0.2963
5	1	0.0741	0.1605	0.1944	0.2438	0.1752	0.2562
5	2	0.1481	0.3349	0.3866	0.4035	0.2816	0.3503
5	3	0.2222	0.3958	0.4144	0.4329	0.2816	0.3750
5	4	0.2963	0.4074	0.4167	0.4537	0.2816	0.3750
5	5	0.3704	0.4167	0.4167	0.4630	0.2816	0.3750
6	1	0.0823	0.1955	0.1916	0.2149	0.1912	0.1269
6	2	0.2058	0.3513	0.3220	0.3944	0.3616	0.1672
6	3	0.3086	0.4498	0.4099	0.4928	0.4743	0.1754
6	4	0.3086	0.4699	0.4230	0.5109	0.4999	0.1754
6	5	0.3086	0.4707	0.4237	0.5163	0.5058	0.1754
6	6	0.3086	0.4707	0.4237	0.5163	0.5071	0.1754
7	1	0.1715	0.1840	0.1805	0.1991	0.1903	0.1965
7	2	0.2709	0.3313	0.3278	0.3477	0.3526	0.3244
7	3	0.3344	0.4815	0.4528	0.4653	0.4708	0.3824
7	4	0.3798	0.5338	0.4982	0.5102	0.5146	0.3948
7	5	0.4072	0.5374	0.5032	0.5148	0.5154	0.3948
7	6	0.4347	0.5379	0.5034	0.5167	0.5154	0.3948
7	7	0.4621	0.5379	0.5034	0.5187	0.5154	0.3948
8	1	0.0960	0.1717	0.1758	0.2020	0.1505	0.2045
8	2	0.1989	0.3230	0.3072	0.3557	0.2476	0.3241
8	3	0.3166	0.4317	0.4204	0.4944	0.3079	0.3780
8	4	0.4292	0.4940	0.4975	0.5794	0.3417	0.3943
8	5	0.4567	0.5218	0.5229	0.6042	0.3418	0.3970
8	6	0.4750	0.5241	0.5251	0.6089	0.3418	0.3970
8	7	0.4841	0.5250	0.5251	0.6106	0.3418	0.3970
8	8	0.4841	0.5250	0.5251	0.6106	0.3418	0.3970
9	1	0.0875	0.1594	0.1686	0.1767	0.1522	0.1200
9	2	0.1942	0.2842	0.2997	0.3108	0.2894	0.1825
9	3	0.2817	0.4063	0.4129	0.4228	0.4095	0.2075
9	4	0.3451	0.5247	0.5088	0.5383	0.4957	0.2141
9	5	0.3894	0.5751	0.5423	0.5773	0.5432	0.2150
9	6	0.4001	0.5783	0.5560	0.5911	0.5621	0.2150
9	7	0.4108	0.5784	0.5568	0.5941	0.5678	0.2150
9	8	0.4214	0.5785	0.5568	0.5941	0.5705	0.2150
9	9	0.4321	0.5785	0.5568	0.5941	0.5715	0.2150
10	1	0.1515	0.1618	0.1637	0.1712	0.1553	0.1689
10	2	0.2525	0.2942	0.2883	0.3040	0.2867	0.2897
10	3	0.3329	0.4008	0.3956	0.4208	0.3996	0.3731
10	4	0.4108	0.4879	0.4947	0.5221	0.4866	0.4215
10	5	0.4872	0.5533	0.5738	0.5902	0.5442	0.4404
10	6	0.5048	0.5721	0.5904	0.6155	0.5626	0.4439
10	7	0.5155	0.5733	0.5972	0.6224	0.5627	0.4439

Table 3: Values of $\Pr(Q \leq k)$ under IAC.

n	k	Plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	Anti-plurality
2	1	0.0000	0.2857	0.1429	0.1429	0.1429	0.4286
2	2	0.0000	0.4286	0.1429	0.1429	0.1429	0.4286
3	1	0.0000	0.0000	0.1429	0.2143	0.1786	0.2143
3	2	0.0000	0.0000	0.2500	0.2857	0.2857	0.2143
3	3	0.0000	0.0000	0.3214	0.3214	0.3214	0.2143
4	1	0.0714	0.1548	0.1190	0.1667	0.1905	0.2619
4	2	0.1310	0.2540	0.1548	0.2619	0.2857	0.3413
4	3	0.1786	0.3333	0.1786	0.3214	0.3333	0.3810
4	4	0.2143	0.4048	0.1905	0.3571	0.3571	0.4048
5	1	0.0429	0.0905	0.1381	0.1524	0.1286	0.2810
5	2	0.0857	0.1595	0.2214	0.2476	0.1952	0.3857
5	3	0.1286	0.2190	0.2810	0.3190	0.2333	0.4286
5	4	0.1714	0.2762	0.3333	0.3714	0.2619	0.4429
5	5	0.2143	0.3333	0.3810	0.4048	0.2857	0.4524
6	1	0.0260	0.0931	0.1126	0.1580	0.1385	0.1970
6	2	0.0589	0.1537	0.1684	0.2411	0.2433	0.2619
6	3	0.0961	0.2045	0.2156	0.3032	0.3286	0.2857
6	4	0.1351	0.2506	0.2602	0.3563	0.3935	0.2948
6	5	0.1732	0.2944	0.3052	0.4026	0.4394	0.3009
6	6	0.2078	0.3377	0.3506	0.4416	0.4740	0.3052
7	1	0.0563	0.0877	0.1126	0.1234	0.1277	0.2056
7	2	0.1017	0.1454	0.1797	0.1973	0.2226	0.3124
7	3	0.1390	0.1931	0.2299	0.2537	0.2922	0.3788
7	4	0.1701	0.2366	0.2719	0.2996	0.3420	0.4212
7	5	0.1970	0.2785	0.3074	0.3380	0.3781	0.4495
7	6	0.2208	0.3203	0.3377	0.3712	0.4069	0.4697
7	7	0.2424	0.3636	0.3636	0.4015	0.4318	0.4848
8	1	0.0396	0.0769	0.1043	0.1253	0.1055	0.2075
8	2	0.0778	0.1319	0.1606	0.2074	0.1782	0.3154
8	3	0.1152	0.1785	0.2070	0.2721	0.2299	0.3771
8	4	0.1523	0.2203	0.2497	0.3260	0.2662	0.4125
8	5	0.1890	0.2587	0.2898	0.3717	0.2930	0.4336
8	6	0.2244	0.2942	0.3268	0.4112	0.3144	0.4476
8	7	0.2576	0.3275	0.3596	0.4452	0.3328	0.4580
8	8	0.2867	0.3590	0.3869	0.4732	0.3497	0.4662
9	1	0.0283	0.0709	0.0979	0.1072	0.1029	0.1632
9	2	0.0599	0.1186	0.1566	0.1714	0.1897	0.2434
9	3	0.0917	0.1589	0.2029	0.2217	0.2640	0.2869
9	4	0.1215	0.1957	0.2438	0.2645	0.3262	0.3110
9	5	0.1484	0.2301	0.2811	0.3022	0.3766	0.3257
9	6	0.1720	0.2624	0.3161	0.3359	0.4159	0.3357
9	7	0.1923	0.2929	0.3496	0.3666	0.4461	0.3432
9	8	0.2098	0.3220	0.3826	0.3953	0.4699	0.3490
9	9	0.2248	0.3506	0.4166	0.4226	0.4885	0.3536
10	1	0.0450	0.0687	0.0913	0.1055	0.0987	0.1650
10	2	0.0818	0.1147	0.1436	0.1727	0.1782	0.2631
10	3	0.1131	0.1540	0.1866	0.2278	0.2445	0.3322
10	4	0.1409	0.1900	0.2252	0.2759	0.2985	0.3834
10	5	0.1669	0.2240	0.2604	0.3187	0.3413	0.4216
10	6	0.1922	0.2568	0.2924	0.3570	0.3749	0.4498
10	7	0.2175	0.2887	0.3212	0.3917	0.4018	0.4710

Table 4: Values of $\Pr(\overline{Q} \leq k)$ under IAC.

n	k	Plurality	(3,1,0)	Borda	(3,2,0)	(10,9,0)	Anti-plurality
2	1	0.0000	0.4286	0.1429	0.1429	0.1429	0.4286
2	2	0.0000	0.4286	0.1429	0.1429	0.1429	0.4286
3	1	0.0000	0.0000	0.2857	0.3214	0.2857	0.2143
3	2	0.0000	0.0000	0.3214	0.3214	0.3214	0.2143
3	3	0.0000	0.0000	0.3214	0.3214	0.3214	0.2143
4	1	0.1429	0.2381	0.1548	0.2817	0.2817	0.3532
4	2	0.2143	0.4008	0.1865	0.3492	0.3492	0.4048
4	3	0.2143	0.4048	0.1905	0.3571	0.3571	0.4048
4	4	0.2143	0.4048	0.1905	0.3571	0.3571	0.4048
5	1	0.0429	0.1270	0.2175	0.2286	0.2067	0.3456
5	2	0.0857	0.2857	0.3579	0.3683	0.2857	0.4333
5	3	0.1286	0.3214	0.3786	0.3893	0.2857	0.4524
5	4	0.1714	0.3286	0.3810	0.4000	0.2857	0.4524
5	5	0.2143	0.3333	0.3810	0.4048	0.2857	0.4524
6	1	0.0519	0.1439	0.1537	0.2297	0.2141	0.2431
6	2	0.1385	0.2411	0.2866	0.3749	0.3697	0.2965
6	3	0.2078	0.3292	0.3461	0.4312	0.4487	0.3052
6	4	0.2078	0.3372	0.3502	0.4385	0.4671	0.3052
6	5	0.2078	0.3377	0.3506	0.4416	0.4719	0.3052
6	6	0.2078	0.3377	0.3506	0.4416	0.4740	0.3052
7	1	0.0855	0.1256	0.1656	0.1808	0.1950	0.2731
7	2	0.1331	0.2279	0.2750	0.3025	0.3378	0.4186
7	3	0.1656	0.3403	0.3396	0.3806	0.4148	0.4697
7	4	0.1905	0.3608	0.3595	0.3963	0.4316	0.4848
7	5	0.2078	0.3633	0.3633	0.3994	0.4318	0.4848
7	6	0.2251	0.3636	0.3636	0.4004	0.4318	0.4848
7	7	0.2424	0.3636	0.3636	0.4015	0.4318	0.4848
8	1	0.0513	0.1117	0.1430	0.1832	0.1622	0.2609
8	2	0.1102	0.2043	0.2527	0.3262	0.2737	0.3983
8	3	0.1885	0.2827	0.3444	0.4218	0.3302	0.4487
8	4	0.2627	0.3466	0.3772	0.4622	0.3496	0.4647
8	5	0.2747	0.3577	0.3859	0.4700	0.3497	0.4662
8	6	0.2827	0.3586	0.3869	0.4722	0.3497	0.4662
8	7	0.2867	0.3590	0.3869	0.4732	0.3497	0.4662
8	8	0.2867	0.3590	0.3869	0.4732	0.3497	0.4662
9	1	0.0443	0.1033	0.1389	0.1540	0.1584	0.2055
9	2	0.1006	0.1839	0.2376	0.2624	0.2918	0.3084
9	3	0.1419	0.2630	0.3283	0.3389	0.3933	0.3431
9	4	0.1712	0.3321	0.3916	0.4014	0.4458	0.3530
9	5	0.1915	0.3490	0.4103	0.4171	0.4728	0.3536
9	6	0.1998	0.3504	0.4155	0.4209	0.4826	0.3536
9	7	0.2081	0.3506	0.4165	0.4226	0.4863	0.3536
9	8	0.2165	0.3506	0.4166	0.4226	0.4882	0.3536
9	9	0.2248	0.3506	0.4166	0.4226	0.4885	0.3536
10	1	0.0669	0.0973	0.1247	0.1501	0.1485	0.2166
10	2	0.1144	0.1764	0.2194	0.2624	0.2726	0.3552
10	3	0.1598	0.2453	0.2939	0.3602	0.3711	0.4453
10	4	0.2165	0.3254	0.3539	0.4289	0.4277	0.4876
10	5	0.2754	0.3729	0.3802	0.4648	0.4547	0.5053
10	6	0.2837	0.3805	0.3857	0.4763	0.4615	0.5115
10	7	0.2887	0.3815	0.3883	0.4792	0.4615	0.5115

Table 5: Least manipulable rules under IC by various measures.

Number of voters	2	3	4	5	6	7	8	9	10
Q	P	P	None	None	None	None	None	None	None
$\overline{Q}$	P	P	B	None	None	None	None	None	None
M	P	P	B	(10,9,0)	A	None	(10,9,0)	A	None
$\Pr(M = 1)$	P	P	B	(10,9,0)	A	B	(10,9,0)	A	B
$\Pr(M < \infty)$	P	P	B	(10,9,0)	A	A	(10,9,0)	A	A

Table 6: Least manipulable rules under IAC by various measures.

Number of voters	2	3	4	5	6	7	8	9	10
Q	P	P	None	P	P	P	P	P	P
$\overline{Q}$	P	P	P	P	P	P	P	P	P
M	P	P	B	P	P	None	None	None	None
$\Pr(M = 1)$	P	P	B	P	P	(3,1,0)	(3,1,0)	(3,1,0)	(3,1,0)
$\Pr(M < \infty)$	P	P	B	P	P	P	P	P	P

In Tables 5 and 6 we present the best (least manipulable) of our chosen rules under each of the measures $M, \overline{Q}, Q$, for each value of n up to 10 (we write "None" to mean that there is no rule that is dominant over the entire range of k). We use the abbreviations P, B, A for the Plurality, Borda and Antiplurality rules respectively.

In Figures 1–3 we show the values of $M, \overline{Q}, Q$ respectively when $n = 25$, under IC. The corresponding results for IAC are shown in Figures 4–6.

In Figures 7 and 8 we compare the measures $M, \overline{Q}, Q$ for $n = 25$ and the Borda rule.

Finally, in Table 7 we display the expected values of $M, \overline{Q}, Q$, conditional on the voting situation being manipulable.

Comments on results The most basic feature of the results is that for each culture, the relative ordering of the rules (when considering moderately large coalitions) is the same for each of the measures $M, \overline{Q}, Q$. This demonstrates the robustness of these measures, defined using quite different assumptions. It shows that the amount of information in our query model does not play as large a role as we might have expected, *a priori*.

The numerical results, although only extending to $n = 25$, give ample scope for conjecture about the behaviour of these measures for large n. Some of these conjectures were described in [11] for M, and we can make the analogous ones for the other two measures. For example, it seems that under IC and for reasonably large n, Borda is the positional rule for which the distribution function converges fastest to its maximum as k increases.

4 Comparison with existing literature

Assumptions and models After the initial news that manipulability is inevitable [8, 14] much work has been done to minimize the perceived damage done by manipulation. One

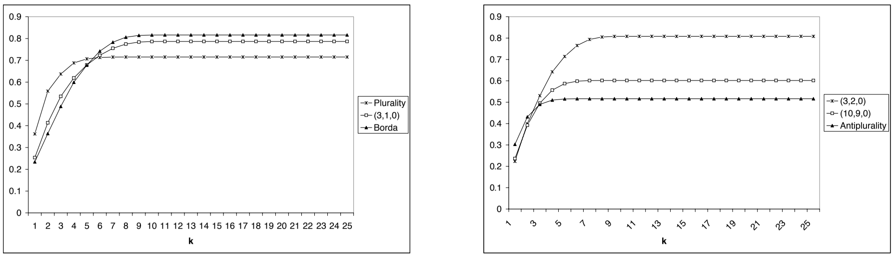

Figure 1: Values of $\Pr(M \leq k)$ when $n = 25$, under IC.

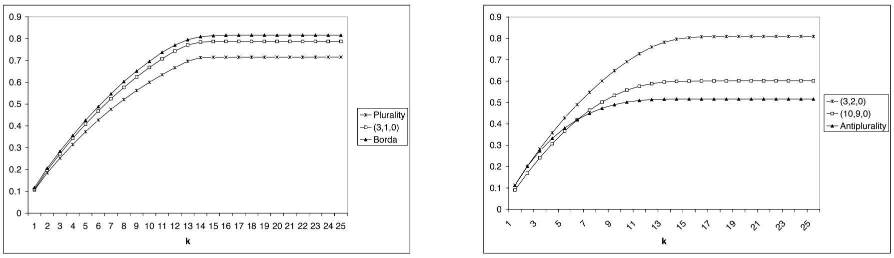

Figure 2: Values of $\Pr(\overline{Q} \leq k)$ when $n = 25$, under IC.

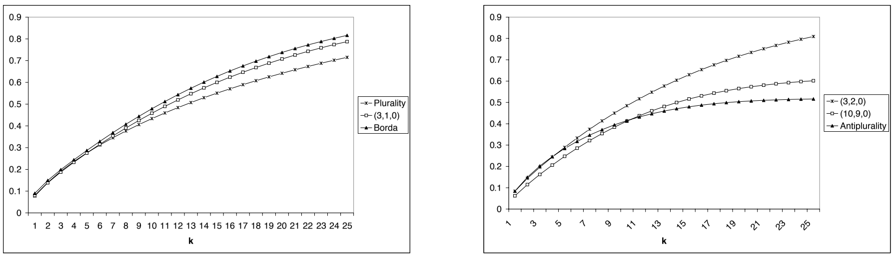

Figure 3: Values of $\Pr(Q \leq k)$ when $n = 25$, under IC.

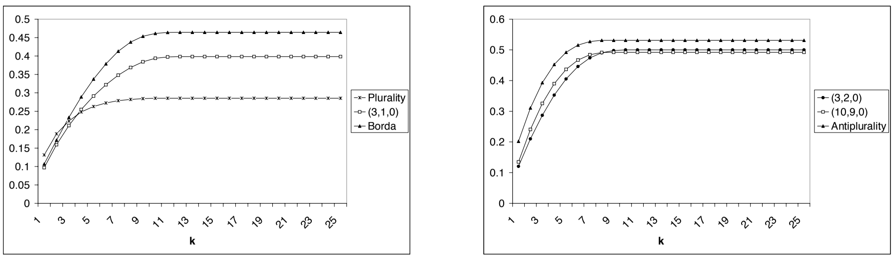

Figure 4: Values of $\Pr(M \leq k)$ when $n = 25$, under IAC.

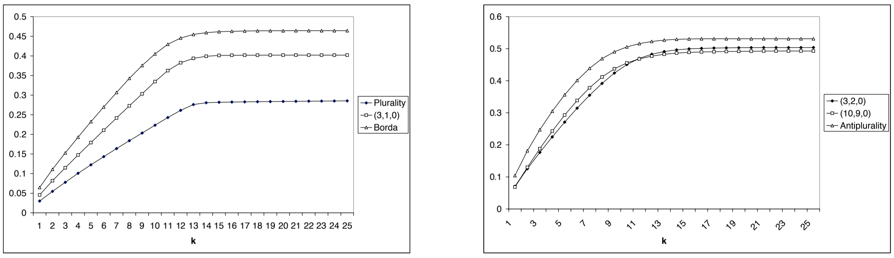

Figure 5: Values of $\Pr(\overline{Q} \leq k)$ when $n = 25$, under IAC.

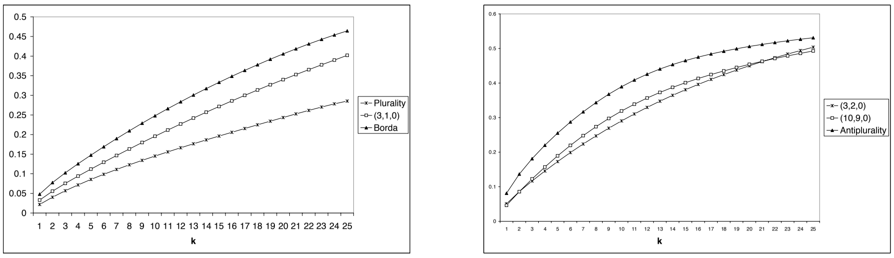

Figure 6: Values of $\Pr(Q \leq k)$ when $n = 25$, under IAC.

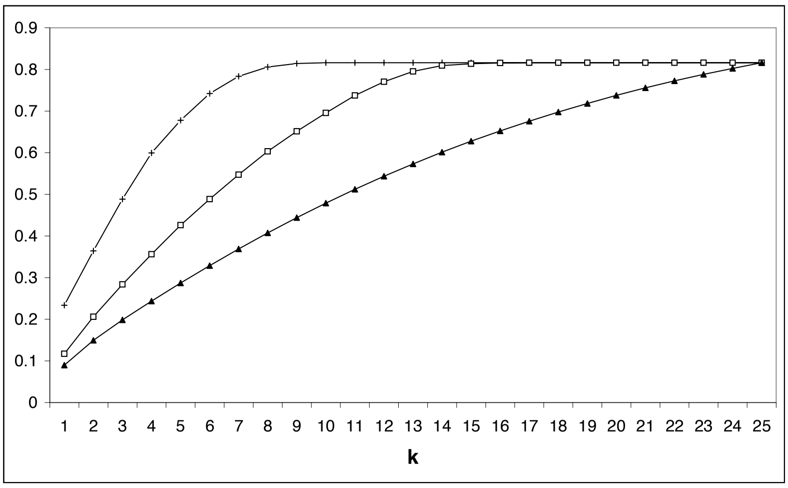

Figure 7: Values of M (top), $\overline{Q}$ (middle), Q (bottom) for Borda when $n = 25$, under IC.

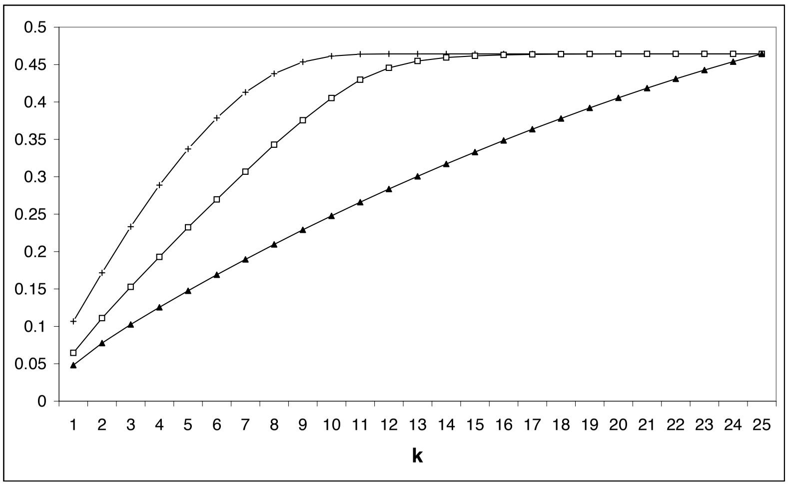

Figure 8: Values of M (top), $\overline{Q}$ (middle), Q (bottom) for Borda when $n = 25$, under IAC.

Table 7: Conditional expected values under IC/IAC for $n = 25$.

	IC			IAC		
scoring rule	$E(Q)$	$E(\overline{Q})$	$E(M)$	$E(Q)$	$E(\overline{Q})$	$E(M)$
Plurality	9.5098	5.8755	1.8756	11.1379	6.7792	2.390
(3,1,0)	9.8509	5.8572	2.9631	11.4501	6.3701	3.8973
Borda	9.7408	5.8306	3.2502	11.4501	6.3701	3.8973
(3,2,0)	9.5963	5.8318	2.9706	9.9792	5.6009	3.4336
(10,9,0)	8.2645	5.0210	2.2368	8.9986	5.1654	2.9661
Antiplurality	6.3244	3.9813	3.6414	7.4724	4.4646	2.5520

method is to find rules that minimize the number of situations in which manipulation can succeed. This leads to the measure P (and also to the idea of restricting user preferences, which we ignore in the present article). This approach has led to what we call the "traditional" social choice literature on manipulability of voting rules. Most authors in this area consider only IC or only IAC; they have various ways of dealing with tiebreaking (typically ignoring it or allowing a fixed candidate to prevail, thereby violating neutrality). For a more thorough discussion of the assumptions used in this literature, and a more extensive bibliography, see [11].

Another approach is to find rules for which even if manipulation is logically possible, it is practically too difficult to carry out. This leads to what we call the "computational" social choice literature. Although this line of thought leads to the measures we have introduced in this paper, most recent research in the area has instead concentrated on using the theory of computational complexity to find barriers to manipulation. Most authors have considered only IC. They have used the standard theory of NP-completeness as a measure of hardness, with few exceptions.

One important assumption common to most work in the area is that we have a coalition of manipulators, with all other voters acting sincerely. Relatively little work has been done on integrating game theory into this framework, and assuming, for example, full common knowledge and a sophisticated strategically voting electorate. Coalitional manipulation does not lead to the usual Nash equilibrium concept, which is appropriate to individual manipulation. A related assumption is that manipulation is an occurrence that should be minimized. Interesting recent work [9] argues that total social welfare is often in fact increased by encouraging all agents to vote strategically. However, we restrict to the standard assumptions in the present paper.

Methodology In the "computational" area, NP hardness results have been shown using reductions to standard NP-complete problems, for example in [4]. Polynomial time algorithms for manipulation are mostly rather obvious to derive and have been included in several papers, for example [4, 20].

Turning to the "traditional" area, we note that many authors have employed statistical sampling rather than exhaustive enumeration of voting situations, even when $m = 3$ (for example [17, 1]). Of course, for $m > 3$ this may be essential as we discussed in Section 6.1). Many other papers have derived analytic formulae for finite n under IAC (see [18, 10] for the best way to do this), and asymptotic formulae under IC.

Results In the "computational" area, it has gradually become clear that manipulation of anonymous voting rules, at least in our case of a few strategic players opposed only by naive, sincere voters, is no harder, in the sense of complexity theory, than computing the result of the election. This is a polynomial time procedure (in n, for fixed m) for all but a handful of commonly used rules (Kemeny's rule is a notable exception). Lately most of the research effort appears to have gone into weighted voting, where agents have different weights for their votes, and several NP-hardness results have been obtained, most notably [4]. The significance of the latter is that it shows that some problems are NP-hard for fixed m, whereas older results allowed m to be arbitrary (in the language of parametrized complexity, most problems considered in the area are fixed-parameter tractable).

In any case, the theory of NP-hardness gives only worst-case bounds, and it has been shown [3] that under weak conditions (satisfied by all commonly used rules) there does not exist a voting rule that is hard to manipulate on average. Among recent results showing that ordinary manipulation is easy, we mention [7] which shows that for $m = 3$, for rules that are sufficiently far from dictatorships, a single voter can manipulate with high probability by making a random change in his vote (this result should yield a lower bound on $\Pr(Q \leq 1)$). Explicit polynomial-time algorithms for coalitional manipulation of several voting rules are given in [20].

We note that the quantity M was introduced in [2]. It was investigated in detail for scoring rules in [11] and [12]. The measure $\Pr(M \leq k)$ has been used implicitly in [16], where it was shown that for IAC, and for all faithful scoring rules (those where all weights in the weight vector defining the rule are different) and runoff rules based on them, there is a constant C, depending on m and the rule, so that $P(M \leq k) \leq Ck/n$. Similar results for IC but with the bound $Ck/\sqrt{n}$ have been obtained [15]. A closely related result by [19] shows that for a wide class of (anonymous) rules that includes all commonly used rules, there is a threshold at around $\Theta(\sqrt{n})$ below which the probability under IC that a profile is manipulable by a fixed subset of voters is asymptotically zero, and above which it is asymptotically 1. By symmetry, this latter probability equals $\Pr(Q \leq k)$ for IC.

In the case of IC, our results clarify the conventional wisdom on the relative manipulability of scoring rules, and Borda's in particular. For example, Saari [13] claims that (when $m = 3$) the Borda rule is the scoring rule that is least susceptible to "micro manipulations" (only individual voters or small coalitions) but is quite poor among the positional voting methods for macro manipulations. His definition of "micro manipulation" deals with the case $k = o(\sqrt{n})$, where there are few manipulating coalitions for positional rules as we have seen. In [12] the last two authors have proved Saari's assertion in more generality. Also, it appears likely from our numerical results that Borda is the most manipulable scoring rule when k is of order $\sqrt{n}$ or greater, by all our measures.

The results that we and other authors have derived show a great sensitivity to the culture assumption. Sometimes the best rule under IC is precisely the worst under IAC by the same measure. However the results for the rules between plurality and Borda yield the same ordering for both IC and IAC, and we suspect that these results will be the same for all rules other than antiplurality. This robustness to distributional assumptions gives us more confidence in the use of these measures.

5 Extensions And Future Work

There are several obvious directions in which to extend the work of the present article. These include:

- different probability models for voter preferences;
- cases where voters have even less information about others' opinions (perhaps only a probability distribution);
- weighted voting;

- developing software to systematically analyse classes of voting rules as above, and allow optimization of them with respect to parameters;
- introducing game theory explicitly into the models;
- other classes of voting rules;
- asymptotic results for large n.

As specific examples for the last two points, we note that [12] describes the asymptotic growth of M for scoring rules in any number of variables, under IC, in great detail. We would like to do something similar for $\overline{Q}$ and Q. Also, preliminary experiments show that our measures behave rather differently for the Copeland rule, a well-known Condorcet consistent rule, than they do for the positional rules. In particular, under IC the minimum coalition size appears to grow faster than $\sqrt{n}$. This is perhaps related to findings in [6] on the complexity of manipulation for variants of this rule.

A more radical change is to follow the line of work in [5, 9] that argues that widespread strategic voting should be encouraged for reasons of utilitarian efficiency. This seems rather attractive to us: after all, a dictatorship is never manipulable, so minimizing the occurrence of manipulation should probably not be our primary goal. More general measures of social welfare should be explored and voting rules compared in that way.

Appendix

6 Computation of the measures

6.1 Algorithms

The computation of the measures we have discussed is easy from the point of view of computational complexity theory, but practically challenging. Let $t = m!$. The number of possible voting situations is the number of solutions in nonnegative integers to the equation $n_1 + n_2 + \cdots + n_t = n$. This equals $\binom{n+t-1}{n}$ and is therefore in $\Theta(n^{t-1})$. For each voting situation, a brute force approach to computing the distribution functions means that we must consider each type of subset of size k, for each k from 1 to n. The number of these types of subsets is $\Theta(k^{t-1})$ for Q and $\Theta(k^{t/2-1})$ for U and M. Once we have a voting situation and a subset, we must test to see whether a successful manipulation is possible. For the rules presented here (and many others), this can be described by the feasibility of an integer linear programming problem in $\Theta(t)$ variables with $\Theta(m)$ constraints (see [11] for details). This feasibility can be decided in polynomial time in n for fixed t and so hence can the entire computation.

However even when $m = 4$ the number of voting situations exceeds 10^{24} and so exhaustive computation is practically impossible. In this paper we present computational results only for $m = 3$, so as not have have to resort to stochastic simulation. Even when $m = 3$ and $n = 100$, the number of possible voting situations is nearly 10^8 and small speedups can make the difference between practical and impractical computation. We discuss some of them in Section 6.2.

Table 8: Different types of manipulation for a three-candidate election.

Sincere outcome	Manipulated outcome	Possible?	Coalition member types
1.$\|a\| > \|b\| \geq \| c\|$	b wins	Yes	bac, cba
	a, b tie	Yes	bac, cba
	c wins	Yes	cab, bca
	a, c tie	Yes	cab, bca
	b, c tie	No	
	3-way tie	No	
2.$\|a\| = \|b\| \geq \| c\|$	a wins	Yes	abc, cab
	b wins	Yes	bac, cba
	c wins	No	

Manipulation of scoring rules We can describe manipulable voting situations by linear systems. We give details for one case and other cases are similar.

Suppose we have $|a| > |b| \geq | c|$ and we want to calculate the number of minimal manipulating coalitions in favour of b.

So the voters who can participate in this coalition are those who prefer b to a and want to change their votes. We only enumerate the voters that want to change their votes and whose participation in a manipulating coalition is useful. So, abc, acb, bca and cab are not included and we have only bac and cba who can change their votes to bac and bca. We denote the number of people that express the preference bac by y_1 and those who express the preference bca by y_2.

Now, we calculate the total score of a, b and c after manipulating:

$|a'| = |a| + (y_1 - x_3)w_2 + y_2w_3 - x_6w_3$

$|b'| = |b| + (y_1 - x_3)w_1 + y_2w_1 - x_6w_2$

$|c'| = |c| - (y_1 - x_3)w_3 + y_2w_2 - x_6w_1$

As the manipulation is in favor of b, we should have $|b'| \geq |a'|$ and $|b'| \geq| c'|$ and also $y_1 + y_2 = x_3 + x_6$. These conditions yield a linear system and we can exploit this method for the other types of manipulations that can be seen in Table 8.

6.2 Implementation details

We need only perform computation for those voting situations for which $|a| \geq| b| \geq |c|$, because of our tiebreaking convention. This means that each such voting situation is weighted by the size of its orbit under the symmetric group of permutations of the candidates. This size is $1, 3, 6$ according as there is a 3-way tie, 2-way tie, or single winner.

The probability of a given voting situation under IAC is given by $\frac{1}{|S|}$ where S is the set of all voting situations, namely the number of solutions to the equation $\sum_{i=1}^{6} n_i = n$ for non-negative integers v_i. Thus $|S|$ is equal to $\binom{n+5}{5}$.

Under IC, then all of the profiles are equally likely, meaning that the probability of the voting situation is given by

$$\frac{\text{Number of profiles yielding a given voting situation}}{\text{Total number of profiles}}.$$

The numerator is equal to $\frac{n!}{n_1!n_2!\cdots n_6!}$ and the denominator is equal to 6^n.

In order to compute $\Pr(M \leq k)$ we proceed as follows. For each voting situation, instead of examining each type of subset and seeing whether it contains a manipulating coalition, we use a case analysis to solve the integer programming problem whose objective function is the size of the manipulating coalition [11].

To compute $\Pr(Q \leq k)$ and $\Pr(\overline{Q} \leq k)$ we have used brute force enumeration of (types of) subsets of voters. It may be much easier to enumerate minimal manipulating coalitions and then compute the probability that a random subset contains at least one of them. However the inclusion-exclusion argument involved was sufficiently complicated that we have not yet carried it out.

We can readily compute the conditional expectations $E[M \mid M < \infty]$, etc, from the distribution functions as follows. We have

$$\sum_{k=0}^{n} \Pr(M > k) = \sum_{k=1}^{n} k \Pr(M = k) + (n+1) \Pr(M = \infty).$$

Reorganizing this equation yields

$$\begin{aligned} E[M \mid M < \infty] &= \frac{\sum_{k=1}^{n} k \Pr(M = k)}{\Pr(M < \infty)} \\ &= \frac{\sum_{k=0}^{n} [1 - \Pr(M \leq k)] - (n+1) \Pr(M = \infty)}{\Pr(M < \infty)} \\ &= n + 1 - \frac{\sum_{k=1}^{n} \Pr(M \leq k)}{\Pr(M \leq n)} \\ &= n - \frac{\sum_{k=1}^{n-1} \Pr(M \leq k)}{\Pr(M \leq n)}. \end{aligned}$$

C and C++ code used to generate the numerical results in this paper, and following the above outline, is available on request from the authors.

References

[1] FT Aleskerov and E Kurbanov. A degree of manipulability of known social choice procedures. In *Current trends in Economics: theory and applications*, pages 13–27, 1999.

[2] John R. Chamberlin. An investigation into the relative manipulability of four voting systems. *Behavioral Sci.*, 30(4):195–203, 1985.

[3] Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *AAAI*, 2006.

[4] Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *J. ACM*, 54(3):Art. 14, 33 pp. (electronic), 2007.

[5] Keith Dowding and Martin van Hees. In praise of manipulation. *British Journal of Political Science*, 38(01):1–15, 2008.

[6] Piotr Faliszewski, Edith Hemaspaandra, and Henning Schnoor. Copeland voting: ties matter. In *AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 983–990, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.

[7] Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. *Foundations of Computer Science, Annual IEEE Symposium on*, 0:243–249, 2008.

[8] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–601, 1973.

[9] Aki Lehtinen. The welfare consequences of strategic behaviour under approval and plurality voting. *European Journal of Political Economy*, 24(3):688–704, September 2008.

[10] Dominique Lepelley, Ahmed Louichi, and Hatem Smaoui. On ehrhart polynomials and probability calculations in voting theory. *Social Choice and Welfare*, 30(3):363–383, April 2008.

[11] Geoffrey Pritchard and Mark C. Wilson. Exact results on manipulability of positional voting rules. *Soc. Choice Welf.*, 29(3):487–513, 2007.

[12] Geoffrey Pritchard and Mark C. Wilson. Asymptotics of the minimum manipulating coalition size for positional rules under impartial-culture behaviour. *submitted for publication*, 2009.

[13] Donald G. Saari. Susceptibility to manipulation. *Public Choice*, 64:21–41, 1990.

[14] Mark Allen Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *J. Econom. Theory*, 10(2):187–217, 1975.

[15] Arkadii Slinko. How large should a coalition be to manipulate an election? *Math. Social Sci.*, 47(3):289–293, 2004.

[16] Arkadii Slinko. How the size of a coalition affects its chances to influence an election. *Soc. Choice Welf.*, 26(1):143–153, 2006.

[17] David A. Smith. Manipulability measures of common social choice functions. *Soc. Choice Welf.*, 16(4):639–661, 1999.

[18] Mark C. Wilson and Geoffrey Pritchard. Probability calculations under the IAC hypothesis. *Math. Social Sci.*, 54(3):244–256, 2007.

[19] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *ACM Conference on Electronic Commerce*, pages 109–118, 2008.

[20] Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artif. Intell.*, 173(2):392–412, 2009.