An Optimal Single-Winner Preferential Voting System Based on Game Theory

Ronald L. Rivest and Emily Shen

Abstract

We describe an *optimal* single-winner preferential voting system, called the "GT method" because of its use of symmetric two-person zero-sum game theory to determine the winner. Game theory is used not to describe voting as a multi-player game between voters, but rather to define when one voting system is better than another one. The cast ballots determine the payoff matrix, and optimal play corresponds to picking winners optimally.

The GT method is a special case of the "maximal lottery methods" proposed by Fishburn [14], when the preference strength between two candidates is measured by just the margin between them. We suggest that such methods have been somewhat underappreciated and deserve further study.

The GT system, essentially by definition, is *optimal*: no other preferential voting system can produce election outcomes that are preferred more by the voters, on the average, to those of the GT system. We also look at whether the GT system has several standard properties, such as monotonicity, Condorcet consistency, etc. We also briefly discuss a deterministic variant of GT, which we call GTD.

We present empirical data comparing GT and GTD against other voting systems on simulated data.

The GT system is not only theoretically interesting and optimal, but simple to use in practice; it is probably easier to implement than, say, IRV. We feel that it can be recommended for practical use.

1 Introduction

Voting systems have a rich history and are still being vigorously researched. We refer the reader to surveys and texts, such as Börgers [1], Brams [2], Brams and Fishburn [3], Fishburn [13], Kelly [18], and Tideman [33], for overviews.

The purpose of this paper is to describe a preferential voting system, called the "GT method," to study its properties, and to compare it with some other well-known voting systems.

The GT method is a special case of the "maximal lottery methods" discussed by Fishburn [14] (who references Kreweras [19] as the first to mention them). A *lottery* assigns a probability to each candidate; a lottery method outputs such a lottery, and the election winner is chosen randomly according to those probabilities. Maximal lotteries are those that voters prefer at least as well as any single candidate or any other lottery. The preference strength between two lotteries is the expected value of a social evaluation function applied to the vote differential (margin) between candidates. The GT method has the identity function as the social evaluation function (i.e., the strength of the social preference between two candidates is the vote margin between them).

We suggest that such voting systems with probabilistic output have received insufficient attention, both in the literature and in practice, and that they are really the most natural resolution of the "Condorcet cycle" paradox that plagues preferential voting systems.

More generally, at a high level, the approach is based on a "metric" or "quantitative" approach to comparing two voting systems, which is a nice complement to the more usual "axiomatic" or "property-based" approach common in the literature; the metric approach enables a simple comparison of any two voting systems, given a distribution on profiles.

Finally, the GT method is easy to use in practice; we discuss some implementation details.

The contributions of this paper are as follows:

- We define "relative advantage" as a metric to compare two preferential voting systems.
- We define the GT method as the "optimal" preferential voting system with respect to relative advantage. This includes our proposal for resolving ambiguity when the optimal mixed strategy is not unique.
- We compare the GT method and various voting systems experimentally and show a ranking of these systems, relative to GT.
- We propose a deterministic variant of GT, called GTD, which performs nearly as well as GT, and may be more acceptable to those who object to randomized methods.

2 Preliminaries

Candidates and ballots We assume an election where n voters are to select a single winner from m alternatives ("candidates"). We restrict attention to *preferential voting systems*, where each ballot lists candidates in order of preference. We assume that all ballots are *full* (they list all candidates), but it is a simple extension to allow voters to submit truncated ballots, to write in candidates, or to express indifference between candidates (details omitted).

Profiles, preference and margin matrices, and margin graphs A collection C of (cast) ballots is called a *profile*. A profile is a multi-set; two ballots may list candidates in the same order.

A profile has an associated *preference matrix* N—the $m \times m$ matrix whose (x, y) entry is the number of ballots expressing a preference for candidate x over candidate y. Each entry is nonnegative, and $N(x, y) + N(y, x) = n$, since all ballots are assumed to be full.

It is also useful to work with the *margin matrix* M — the $m \times m$ matrix defined by $M(x, y) = N(x, y) - N(y, x)$, so that $M(x, y)$ is the margin of x over y—that is, the number of voters who prefer x over y minus the number of voters who prefer y over x. The matrix M is anti-symmetric with diagonal 0; for all x, y we have: $M(x, y) = -M(y, x)$.

From the margin matrix M we can construct a directed weighted *margin graph* G whose vertices are the candidates and where there is an edge from x to y weighted $M(x, y)$ whenever $M(x, y) > 0$. If $M(x, y) = M(y, x) = 0$ then voters are, on the whole, indifferent between x and y, and there is no edge between x and y.

Voting system – social choice function A voting system provides a *social choice function* that takes as input a profile of cast ballots and produces as output the name of the election winner. (In some systems the output may be a set of winners.) The social choice function may be *deterministic* or *randomized*. While most but not all voting systems in the literature are deterministic, the GT system is randomized. We also describe a deterministic variant, GTD, of the GT system.

Definition 2.1 *Let p be an arbitrary probability distribution over some finite set X. We say that the* support *$supp(p)$ for p is the set $\{x \mid p(x) > 0\}$ of elements in X that p assigns nonzero probability. Similarly, if V is a discrete random variable, we let $supp(V)$ denote the set of values of V that occur with nonzero probability.*

3 Generalized Ties

A *Condorcet winner* is a candidate x who beats every other candidate in a pairwise comparison: for every other candidate y, more voters prefer x to y than prefer y to x. Thus, the margin matrix M has only positive entries in every off-diagonal position of row x. Equivalently, for each other candidate y, the margin graph contains a directed edge from x to y.

If there is no Condorcet winner, we say that there is a "*generalized tie*," since for every candidate x there exists some other candidate y whom voters like at least as much as x.

The interesting question is then: *When there is a generalized tie, how should one do the "tie-breaking" to pick a single winner?*

4 Breaking Ties Using a Randomized Method

We feel strongly that the best way of breaking a generalized tie is to use an appropriate randomized method. Of course, when there is a clear winner (by which we mean a Condorcet winner) then a randomized method is not needed. A randomized method is only appropriate when a tie needs to be broken.

Academic literature on voting systems has often eschewed proposals having a randomized component. For example, Myerson [25, p. 15] says,

> "Randomization confronts democratic theory with the same difficulty as multiple equilibria, however. In both cases, the social choice ultimately depends on factors that are unrelated to the individual voters' preferences (private randomizing factors in one case, public focal factors in the other). As Riker (1982) has emphasized, such dependence on extraneous factors implies that the outcome chosen by a democratic process cannot be characterized as a pure expression of the voters' will."

We would argue that Myerson and Riker have it backwards, since, as we shall see, voting systems can do better at implementing the voters' will if they are randomized.

Arbitrary deterministic tie-breaking rules, such as picking the candidate whose name appears first in alphabetical order, are clearly unfair. And, while much work has gone into devising clever voting systems that break generalized ties in apparently plausible but deterministic manners, the result is nonetheless arguably unfair to some candidates.

The strongest reason for using a randomized tie-breaking method is that for any deterministic voting system there is another voting system whose outcomes are preferred by voters on the average, while there exist randomized voting systems which are not so dominated by another system. This is effectively just a restatement of the minimax theorem, due to von Neumann, that optimal strategies in two-person zero-sum games may need to be randomized.

It is not a new idea to have a voting system that uses randomization, either in theory or in practice. Using a randomized method is in fact a common and sensible way of breaking ties.

Several recent elections have used randomized methods to break ties. In June, 2009, when the city of Cave Creek, Arizona had a tie between two candidates for a city council seat, the two candidates drew cards from a shuffled deck to determine the winner[1]. In November, 2009, the mayor of Wendell, Idaho, was determined by a coin toss, when the challenger and the incumbent were tied. In February, 2010, in Sealy, Texas, dice were used to resolve a tied election for city council membership.

[1] "Election at a Draw, Arizona Town Cuts a Deck," NY Times, June 17, 2009.

Several previous voting system proposals use randomization to determine the outcome. For example, the “random dictator” voting system [15, 31] picks a random ballot, and uses it to name the winner. This method always uses randomization, not just for tie-breaking. Gibbard [15] proves that if a system is strategy-proof (and satisfies certain other conditions), then it must be the random dictator method.

Sewell et al. [31] propose a randomized voting system based on maximum entropy considerations; this is, however, a social welfare function (it produces a complete ordering, not just a single winner), not a social choice function. Potthoff [26] proposes a randomized method for the case of a three-candidate election with a majority cycle. Laffond et al. [20] propose a randomized method based on game theory for parties to pick platform issues, a situation attributed by Shubik [32] to Downs [8].

Other voting systems, such as the Schulze method [30], use randomization as a final tie-breaker.

5 Optimal Preferential Voting Systems

How should one compare a voting system P against another voting system Q? Here P and Q are (possibly randomized) social choice functions that each take a profile C of cast ballots and produce an election outcome or winner, $P(C)$ or $Q(C)$.

There is a long list of well-studied properties of voting systems, such as monotonicity, consistency, strategy-proofness, etc.; such studies exemplify the “axiomatic” approach to voting systems. One can certainly ask whether a voting system has these desirable properties. The inference is usually that a system with more desirable properties is the better system. But this approach can sometimes give rather conflicting and inconclusive advice.

Here is a more direct approach:

> A voting system P is said to be better than a voting system Q if voters tend to prefer the outcome of P to the outcome of Q.

How can one make this appealing intuition precise?

Let $\mathcal{C}$ be an assumed probability distribution on the profiles of cast ballots. (The details of $\mathcal{C}$ will turn out to be not so important, since GT is optimal on each profile C separately.)

Suppose we play a game $G_{\mathcal{C}}(P, Q)$ between P and Q as follows:

- A profile C of cast ballots for the election is chosen according to the distribution $\mathcal{C}$.
- P and Q compute respective election outcomes $x = P(C)$ and $y = Q(C)$.
- The systems are scored as follows: P wins $N(x, y)$ points, and Q wins $N(y, x)$ points.

Note that the net number of points gained by P, relative to the number of points gained by Q, is just the margin $M(x, y) = N(x, y) - N(y, x)$; more voters prefer P's outcome to Q's outcome than the reverse if $M(x, y) > 0$.

Definition 5.1 *We say that the* relative advantage *of voting system P over voting system Q, denoted $\mathbf{Adv}_{\mathcal{C}}(P, Q)$, with respect to distribution $\mathcal{C}$ on profiles, is*

$$\mathbf{Adv}_{\mathcal{C}}(P, Q) = E_{\mathcal{C}}(M(x, y)/ |C|) \tag{1}$$

where $x = P(C)$ and $y = Q(C)$, where $E_{\mathcal{C}}$ denotes expectation over profiles C chosen according to the distribution $\mathcal{C}$ and with respect to any randomization within P and Q, and where $0/0$ is understood to equal 0 if $|C| = 0$. When $\mathcal{C}$ has all of its support on a single profile C, we write $\mathbf{Adv}_C(P, Q)$.

Definition 5.2 *We say that voting system P is as good as or better than voting system Q (with respect to distribution $\mathcal{C}$ on profiles), if $\mathbf{Adv}_{\mathcal{C}}(P, Q) \geq 0$.*

Definition 5.3 *We say that voting system P is* optimal *if it is as good as or better than every other voting system for any distribution $\mathcal{C}$ on profiles—equivalently, if for every profile C and for every voting system Q we have $\mathbf{Adv}_{C}(P, Q) \geq 0$.*

Intuitively, P will win more points than Q, on the average, according to the extent that voters prefer P's outcomes to Q's outcomes. If P's outcomes tend to be preferred, then P should be considered to be the better voting system. And if P is as good as or better than any other voting system, for any distribution on profiles, then P is optimal.

Note that if P is as good as or better than Q on every distribution $\mathcal{C}$ on profiles, then P must be as good or better than Q on each particular profile C, and vice versa, so the details of distribution $\mathcal{C}$ don't matter.

6 Game Theory

We now describe how to construct an optimal voting system using game theory.

In the game $G_C(P, Q)$, the margin $M(x, y)$ is the "payoff" received by P from Q when P picks x, and Q picks y, as the winner for the election with profile C. The comparison of two voting systems reduces to considering them as players in a distribution on two-person zero-sum games—one such game for each profile C.

The theory of two-person zero-sum games is long-studied and well understood, and optimal play is well-defined. See, for example, the excellent survey article by Raghavan [27].

The expected payoff for P, when P chooses candidate x with probability p_x and when Q independently chooses candidate y with probability q_y is:

$$\sum_x \sum_y p_x q_y M(x, y) \ . \tag{2}$$

An optimal strategy depends on the margin matrix M. When there is a Condorcet winner, the optimal strategy will always pick the Condorcet winner as the election winner. When there is no Condorcet winner, there is a generalized tie, and the optimal strategy is a mixed strategy. Computing the optimal mixed strategy is not hard; see Section 7. Playing this optimal mixed strategy yields an optimal preferential voting system—no other voting system can produce election outcomes that are preferred more by the voters, on average.

The set of candidates with nonzero probability in the optimal mixed strategy for the game associated with profile C is $\text{supp}(GT(C))$. (If there is not a unique optimal mixed strategy, GT uses the most "balanced" optimal mixed strategy, as described in Section 7.) Intuitively, $\text{supp}(GT(C))$ is the set of "potential winners" for the election with profile C for the GT voting system. If there is a Condorcet winner x, then $\text{supp}(GT(C)) = \{x\}$; otherwise, the GT winner is chosen randomly from $\text{supp}(GT(C))$ according to the optimal mixed strategy probabilities.

7 Computing Optimal Mixed Strategies

One can solve a two-person zero-sum symmetric game with $m \times m$ payoff matrix M using a simple reduction to linear programming. Each solution to the linear program provides an optimal mixed strategy for the game. (See Raghavan [27, Problem A, page 740] for details.)

When ballots are full and the number of voters is odd, the optimal mixed strategy p^* is uniquely defined (see Laffond et al. [21]). There are other situations for which there is a unique optimal mixed strategy. With a large number of voters, one would expect the optimal mixed strategy to be unique.

In the case when there is not a unique optimal mixed strategy, we propose that GT picks the unique optimal mixed strategy that minimizes the sum of squares $\sum_i p_i^2$; this strategy can be computed easily with standard quadratic programming packages. This approach then gives a well-defined lottery as output, and treats candidates symmetrically.

8 Selecting the Winner

As we have seen, the GT voting system comprises the following steps:

1. [**Margins**] Compute the margin matrix M from the profile C of cast ballots.
2. [**Optimal mixed strategy**] Determine the optimal mixed strategy p^* for the two-person zero-sum game with payoff matrix M.
3. [**Winner selection**] Select the election winner by a randomized method in accordance with the probability distribution p^*. (If there is a Condorcet winner x, then $p^*(x) = 1$ and this step is trivial.)

There are of course details that must be taken care of properly with using a randomized method to select a winner; these details are very similar to those that arise when generating suitable random numbers of post-election audits; see Cordero et al. [6].

GTD—A Deterministic Variant of GT We now describe a deterministic variant of the GT voting system, which we call GTD. The optimal mixed strategy is computed as in GT, but the winner selection then proceeds in a deterministic manner.

Instead of randomly picking a candidate according to this probability distribution, GTD chooses the candidate with the maximum probability in this optimal mixed strategy. (If there is more than one candidate with the maximum probability in the optimal mixed strategy, then the one with the least name alphabetically is chosen.)

The GTD method does not require any randomness—it is a deterministic social choice function. We expect that in practice it would perform as well as the GT method. However, since GTD is deterministic, one cannot prove that it is optimal.

9 Properties of the GT voting system

Although our focus is on comparing voting systems using "relative advantage," instead of an axiomatic approach, we briefly consider how GT fares with respect to some standard properties.

Optimality. Optimality is perhaps the most important property of the GT voting system. No preferential voting system can produce election outcomes that are preferred more by voters to those of the GT system, on average.

Condorcet winner and loser criteria. Fishburn [14] proves that maximal lotteries satisfy the strong Condorcet property: If the candidates can be partitioned into nonempty subsets A and B such that, for all $a \in A$ and all $b \in B$, more voters prefer a to b than b to a, then the winner will be a candidate in A. This result implies in particular that the GT method will always elect a Condorcet winner, if one exists, and will never elect a Condorcet loser, if one exists. Schulze [30] notes: "The Condorcet criterion implies the majority criterion. Unfortunately, compliance with the Condorcet criterion implies violation of other desired criteria like consistency (Young, 1975) [35], participation (Moulin, 1988) [24], later-no-help, and later-no-harm (Woodall, 1997) [34]."

Pareto optimality. A voting system satisfies *Pareto optimality* if whenever there exist two candidates x and y such that no voter prefers candidate y to x, and at least one voter prefers x to y, then the voting system never elects y. Fishburn [14] proves that maximal lottery methods satisfy Pareto optimality (and thus GT does).

Monotonicity. A voting system satisfies *monotonicity* if, if a voter raises a candidate x on her ballot without changing the order of other candidates, then the probability that the voting system elects x does not decrease. The GT system is not monotonic. This can be seen by analyzing the optimal mixed strategy probabilities of the simplest generalized tie, whose margin graph is a three-cycle. (See Fishburn [14] and Kaplansky [17, p. 479].)

Independence of clones. A voting system satisfies the *independence of clones* property if replacing an existing candidate `A` with a set of clones does not change the winning probability for any candidates other than `A`. (Schulze [30, p. 141] notes some of the subtleties in the definition of this property, especially when B is already in some sense tied with other candidates.) The GT voting system satisfies independence of clones, for a careful definition of the property. (See the full version of this paper for details.)

Reversal symmetry. A voting system satisfies *reversal symmetry* (see Saari [29]) if it is incapable of naming the same candidate as both the best candidate and the worst candidate (e.g. if the election were run over with every ballot reversed in order). The GT voting system satisfies reversal symmetry in cases where the GT support consists of a unique candidate, which may be the only cases when it makes sense to consider reversal symmetry.

Strategy-proofness. Our definition of relative advantage allows one to compare two voting systems based on which voting system's outcomes are preferred more by the voters, according to voter preferences as expressed in their ballots. We do not take into consideration whether voters might be voting strategically.

The resistance of GT to strategic voting needs study, although we have no reason to believe that GT is more, or less, vulnerable to strategic voting than other preferential voting systems.

10 Empirical Comparison with Other Voting Systems

The approach we are recommending allows one to compare any two voting systems P, Q on a given distribution $\mathcal{C}$ of profiles, by computing the relative advantage $\mathbf{Adv}_{\mathcal{C}}(P, Q)$ of one system over the other.

We compared seven voting systems: plurality, IRV, Borda, minimax, the Schulze method [30], GTD, and GT. We used the margins variant of minimax and the "winning votes" variant of the Schulze method.

We randomly generated 10,000 profiles for $m = 5$ candidates, as follows. Each profile had $n = 100$ full ballots. Each candidate and each voter was randomly assigned a point on the unit sphere—think of these points as modeling candidates' and voters' locations on Earth. A voter then lists candidates in order of increasing distance from her location. With this "planetary" distribution, about 64.3% of the profiles had a Condorcet winner, and about 77.1% of the 10,000 simulated elections had a unique optimal mixed strategy.

	plurality	IRV	Borda	minimax	Schulze	GTD	GT
plurality	0	-23740	-31058	-32030	-32128	-32390	-29978
IRV	23740	0	-14148	-16296	-16380	-15892	-13872
Borda	31058	14148	0	-4546	-4654	-5324	-2522
minimax	32030	16296	4546	0	-58	-1436	-174
Schulze	32128	16380	4654	58	0	-1402	-76
GTD	32390	15892	5324	1436	1402	0	10
GT	29978	13872	2522	174	76	-10	0

Figure 1: Cumulative "point advantages" for our main experiment. Row X column Y shows the sum, over 10,000 simulated elections with 100 votes each, of the margin of X's winner over Y's winner. For example, the entry `13872` in row `GT`, column `IRV` means that on average for a random election from our distribution $\mathcal{C}$ on profiles, 1.3872% more of the electorate prefers the GT outcome to the IRV outcome than the reverse; that is, $\mathbf{Adv}_{\mathcal{C}}(GT, IRV) = 1.3872\%$.

We also tried our experiments under the "impartial anonymous culture" distribution (i.e., the uniform distribution). However, under this distribution there were Condorcet winners almost all (about 93%) of the time, so we chose another distribution.

The code we used, and detailed output data, is available at `http://people.csail.mit.edu/rivest/gt` .

Figure 1 gives the cumulative "point advantage" of each of the seven voting systems against each other in our experiment. For example, the "16380" entry in row "Schulze," column "IRV" means that in an average election, the net number of voters preferring the Schulze outcome to the IRV outcome is about 1.6380 voters (i.e., 1.6380% of the electorate). That is, $\mathbf{Adv}_{\mathcal{C}}(\text{Schulze}, \text{IRV}) \approx 0.016380$.

With this distribution on profiles, there is a clear improvement in quality of output (as measured by relative advantage compared to GT) as one goes from plurality to IRV to Borda to minimax to Schulze. GT and GTD are perfect by definition in this metric, but Schulze is amazingly close. Although GTD and GT are by definition in a dead heat against each other, GTD appears to be a better competitor against the other systems than GT.

Note that when comparing GT with another voting system, there is no expected net point gain for GT if the other system picks a candidate that is in $\text{supp}(GT(C))$. Candidates in $\text{supp}(GT(C))$ have the property that playing any one of them has an expected payoff equal to zero (the value of the game) against GT. If the other system plays a candidate outside of $\text{supp}(GT(C))$, GT will have an expected net point gain and the other system will have an expected loss.

Figure 2 illustrates the number of times each pair of voting systems produced results that "agree with" each other. The column "GTS" refers to the support of GT; a method "agrees with" GTS if it produces an output that is in the support of GT. In our view, level of agreement with the support of GT is an interesting measure of the quality of the results produced by each voting system. Plurality does quite poorly (agreeing with GTS only 55.15% of the time), as does IRV (72.99%), but minimax (99.15%) and the Schulze method (99.51%) have nearly perfect agreement with the support of GT.

Thus, one can perhaps view the evolution of voting system proposals as a continuing effort to identify candidates that are in the support for the optimal mixed strategy for the associated two-person game, without quite realizing that this is the natural goal. That is, voting systems should be (at the minimum) returning winners that are in $\text{supp}(GT(C))$, the set of potential winners for the GT voting system. To do otherwise does not serve the voters as well as can be done. However, since determining the support for the optimal mixed

	plurality	IRV	Borda	minimax	Schulze	GTD	GT	GTS
plurality	10000	5557	4107	4356	4366	4335	4262	5515
IRV	5557	10000	5584	6047	6048	5999	5802	7299
Borda	4107	5584	10000	7854	7874	7813	7193	8913
minimax	4356	6047	7854	10000	9953	8869	8232	9915
Schulze	4366	6048	7874	9953	10000	8895	8246	9951
GTD	4335	5999	7813	8869	8895	10000	8377	10000
GT	4262	5802	7193	8232	8246	8377	10000	10000
GTS	5515	7299	8913	9915	9951	10000	10000	10000

Figure 2: Agreement between pairs of voting systems. Row X column Y gives the number of times that method X produced an outcome that agreed with the outcome of method Y, in our 10,000 trials. Here the "GTS method" refers to the support of GT, and a method "agrees with" GTS if it produces an outcome that is in the support of GT. In our view, frequency of agreement with GTS (producing outcomes in the support of GT) is an important measure of the quality of a preferential voting system.

strategy intrinsically involves linear programming, this computation is non-trivial, so we see a variety of quite complex voting system proposals in the literature, which are, in this view, just approximate computations for (a member of) supp($GT(C)$).

11 Practical Considerations

We believe that the GT voting system is suitable for practical use.

Since the GT voting system depends only on the pairwise preference matrix N, ballot information can be easily aggregated at the precinct level and the results compactly transmitted to central election headquarters for final tabulation; the number of data items that need to be transmitted is only $O(m^2)$, which is much better than for, say, IRV.

Perhaps the only negative aspects with respect to using GT in practice are that (1) its game-theoretic rationale may be confusing to some voters and election officials, (2) it is a randomized method, and may require dice-rolling or other randomized devices in the case of generalized ties, and (3) it is not so clear how to efficiently audit a GT election. (The last aspect is common to many preferential voting systems).

12 Other Related Work

Fishburn [12] gives an excellent overview of voting systems with the Condorcet property.

The idea of using a two-player zero-sum game based on a payoff matrix derived from a profile of ballots is not new; there are several papers that study this and related situations.

Laffond et al. [22] introduce the notion of a "bipartisan set," which is the support of the optimal mixed strategy of a two-player "tournament game." A tournament game is based on an unweighted complete directed graph (a tournament) where each player picks a vertex, and the player picking x wins one point from the player picking y if there is an edge from x to y. They show that any tournament game has a unique optimal mixed strategy, and study the properties of its support.

The weighted version of such a tournament game corresponds to the voting situation we consider (assuming no edge weights are zero); the weight of an edge from x to y corresponds to the margin $M(x, y)$. For the margin graph to be a tournament, no margin may be zero.

Laffond et al. [20] explicitly propose the use of two-player game theory to provide solutions to elections, including the use of randomized methods. However, their focus is on the way political parties choose platform issues, whereas our focus is on "competition" between voting systems rather than between political parties. Our work should nonetheless be viewed as further explorations along the directions they propose.

Le Breton [4, p. 190] proves a general version of Laffond et al.'s earlier result, showing that if all edges satisfy certain congruence conditions, then the weighted tournament game has a unique optimal mixed strategy.

Duggan and Le Breton [9] study the "minimal covering set" of a tournament (proposed by Dutta [10] as a choice function on tournaments), and show that it is the same as Shapley's notion of a "weak saddle" for the corresponding tournament game.

De Donder et al. [7] consider various solution concepts for tournament and weighted tournament games and make set-theoretic comparisons between the corresponding social choice functions.

Michael and Quint [23] provide further results characterizing when there exists a unique optimal strategy in tournament and weighted tournament games.

Dutta et al. [11] introduce "comparison functions," which correspond to general skew-symmetric matrices, as a framework for generalizing choice functions on tournaments.

13 Open Problems

There are many aspects of the GT method, and of probabilistic voting systems in general, that deserve further study. Here are a few such open questions:

- For which pairs of voting systems P and Q, and for which distributions $\mathcal{C}$ on profiles, can $\mathbf{Adv}_{\mathcal{C}}(P, Q)$ be analytically determined? Can one show analytically that GTD performs better than GT against some well-known voting system?
- Can one lower bound (for some assumed distribution $\mathcal{C}$ on profiles) the penalty paid for being deterministic, consistent, or monotonic (i.e., in terms of the advantage of GT over systems with the given property)?
- How sensitive are the output probabilities of GT to the input votes? More generally, how resistant is GT to manipulation, for various notions of manipulation of probabilistic voting systems (e.g., that of [5])?
- Is it possible to modify the Schulze method in a straightforward manner so that it always chooses a winner in the support of GT, while retaining its deterministic character and its other desirable properties?
- To what extent would changing the social evaluation function (see Fishburn [14]) change the perceived relative quality of various voting systems (e.g., via simulation results)?

14 Conclusions

We have described the GT voting system for the classic problem of determining the winner of a single-winner election based on voters' preferences expressed as (full or partial) rank-order listings of candidates.

The GT scheme is arguably optimal among preferential voting systems, in the sense that no other voting system P can produce election outcomes that on the average are preferred by voters to those of GT. We feel that optimality is an important criterion for voting systems.

We believe that the GT voting system is suitable for practical use, when preferential voting is desired. When there is a clear (Condorcet) winner, GT elects that winner. When there is no Condorcet winner, GT produces a "best" set of probabilities that can be used in a tie-breaking ceremony. If one is to use preferential ballots, the GT system can be recommended.

Since the GT system shares some potentially confusing properties, such as non-monotonicity, with many other preferential voting systems, election authorities might reasonably consider alternatives to the GT system, such as a non-optimal but monotonic preferential voting system like the Schulze method, or non-preferential voting systems such as approval voting or range voting.

However, we feel that the optimality property of GT makes it worthy of serious consideration when preferential ballots are to be used.

Acknowledgments

We thank the Stanford CS Department and Microsoft Research for hosting the first author on portions of his sabbatical when some of this research was performed. The second author gratefully acknowledges support from a Bell Labs Graduate Research Fellowship. We also thank Josh Benaloh, Steven Brams, Gilbert Laffond, Jean-François Laslier, Richard Potthoff, Ariel Procaccia, Peter Ryan, Rob Schapire, Warren Smith, and Vanessa Teague for helpful feedback.

References

[1] C. Börgers. *Mathematics of Social Choice: Voting, Compensation, and Division.* SIAM, 2010.

[2] S. J. Brams. *Mathematics and Democracy.* Princeton, 2008.

[3] S. J. Brams and P. C. Fishburn. Voting procedures. In K. Arrow, A. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, pages 173–236. Elsevier, 2002.

[4] M. Le Breton. On the uniqueness of equilibrium in symmetric two-person zero-sum games with integer payoffs. *Économie Publique*, 2005/2(17):187–195, 2005.

[5] V. Conitzer and T. Sandholm. Universal voting protocal tweaks to make manipulation hard. In *Proc. 18th International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 781–788, 2003.

[6] A. Cordero, D. Wagner, and D. Dill. The role of dice in election audits. In *Proc. IAVOSS Workshop on Trustworthy Elections (WOTE 2006)*, 2006.

[7] P. De Donder, M. Le Breton, and M. Truchon. Choosing from a weighted tournament. *Math. Social Sciences*, 40(1):85–109, Jul 2000.

[8] Anthony Downs. *An Economic Theory of Democracy.* Harper and Row, 1957.

[9] J. Duggan and M. Le Breton. Dutta's minimal covering sets and Shapley's saddles. *J. Econ. Theory*, 70:257–265, 1996.

[10] B. Dutta. Covering sets and a new condorcet choice correspondence. *J. Econ. Theory*, 44:63 – 80, 1988.

[11] B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Soc. Choice Welfare*, 16:513–532, 1999.

[12] P. C. Fishburn. Condorcet social choice functions. *SIAM J. Appl. Math.*, 33(3):469–489, Nov 1977.

[13] P. C. Fishburn. Dimensions of election procedures: Analyses and comparisons. *Theory and Decision*, 15(4):371–397, Dec 1983.

[14] P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *Rev. Economic Studies*, 51(4):683–692, Oct 1984.

[15] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, Apr 1977.

[16] A. Gibbard. Straightforwardness of game forms with lotteries as outcomes. *Econometrica*, 46(3):595–614, May 1978.

[17] Irving Kaplansky. A contribution to Von Neumann's theory of games. *Annals. Math. (Second Series)*, 46(3):474–479, Jul 1945.

[18] J. S. Kelly. Social choice bibliography. *Soc. Choice Welfare*, 8:97–169, 1991.

[19] G. Kreweras. Aggregation of preference orderings. In S. Sternberg, editor, *Math. and Soc. Sciences I*, pages 73–79. Moulton, 1965.

[20] G. Laffond, J.-F Laslier, and M. Le Breton. Social choice mediators. *American Economic Review*, 84(2):448–453, May 1994. Proc. 106th Ann. Mtg. Amer. Economic Assn.

[21] G. Laffond, J.-F. Laslier, and M. Le Breton. A theorem on symmetric two-person zero-sum games. *J. Economic Theory*, 72:426–431, 1997.

[22] G. Laffond, J. F. Laslier, and M. LeBreton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5:182–201, 1993.

[23] T. S. Michael and T. Quint. Optimal strategies for node selection games on oriented graphs: Skew matrices and symmetric games. *Linear Algebra and its Applications*, 412:77–92, 2006.

[24] H. Moulin. Condorcet's principle implies the no show paradox. *J. of Economic Theory*, 45:53–64, 1988.

[25] R. B. Myerson. Fundamentals of social choice theory, Sep 1996. Discussion paper 1162. Center for Math. Studies in Economics and Mgt. Science. Northwestern Univ., Evanston, IL.

[26] R. F. Potthoff. The problem of the three-way election. *Essays in Probability and Statistics*, pages 603–620, 1969.

[27] T.E.S. Raghavan. Zero-sum two-person games. In R.J. Aumann and S. Hart, editors, *Handbook of Game Theory*, volume 2, pages 735–759. Elsevier, 1994.

[28] W. H. Riker. *Liberalism Against Populism*. Freeman, 1982.

[29] D. G. Saari. *Geometry of Voting*. Springer-Verlag, 1994.

[30] M. Schulze. A new monotonic, clone-independent, reversal-symmetric, and condorcet-consistent single-winner election method, Aug 2009. `http://m-schulze.webhop.net/schulze1.pdf`.

[31] R. Sewell, D. MacKay, and I. McLean. Probabilistic electoral methods, representative probability, and maximum entropy. *Voting Matters*, 26:16–38, Jan 2009. `http://www.votingmatters.org.uk/ISSUE26/I26P3.pdf`.

[32] M. Shubik. Voting, or a price system in a competitive market structure. *The American Political Science Review*, 64(1):179–181, Mar 1970.

[33] N. Tideman. *Collective Decisions and Voting: the potential for public choice*. Ashgate, 2006.

[34] D. Woodall. Monotonicity of single-seat preferential election rules. *Discrete Applied Math.*, 77:81–98, 1997.

[35] H. P. Young. Social choice scoring functions. *SIAM J. Appl. Math.*, 28(4):824–838, Jun 1975.

Ronald L. Rivest (Email: `rivest@mit.edu`)
MIT Computer Science and Artificial Intelligence Laboratory
Room 32G-692, 32 Vassar Street
Cambridge, MA 02139 USA

Emily Shen (Email: `eshen@csail.mit.edu`)
MIT Computer Science and Artificial Intelligence Laboratory
Room 32G-626, 32 Vassar Street
Cambridge, MA 02139 USA