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

Ronald L. Rivest and Emily Shen

Abstract

We present an *optimal* single-winner preferential voting system, called the "GT method" because of its interesting use of symmetric two-person zero-sum game theory to determine the winner. Game theory is not used 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 quite simple and elegant, and works as follows:

- Voters cast ballots listing their rank-order preferences of the alternatives.
- The margin matrix M is computed so that for each pair (x, y) of alternatives, the pairwise margin $M(x, y)$ is the number of voters preferring x over y, minus the number of voters preferring y over x.
- A symmetric two-person zero-sum game G is defined where the payoff matrix is M. Each player picks an alternative and wins $M(x, y)$ points when he picks x and the opponent picks y.
- An optimal mixed strategy p^* is computed, where $p^*(x)$ is the probability that a player picks x under optimal play. If the optimal mixed strategy is not unique, we use the "most balanced" one. This optimal strategy is the same for both players, since the game is symmetric.
- The election winner is determined by a randomized method, based for example on the use of ten-sided dice, that picks outcome x with probability $p^*(x)$. If there is a Condorcet winner x, then $p^*(x) = 1$, and no dice are needed. We also briefly discuss a deterministic variant of GT, which we call GTD.

The GT system, essentially by definition, is *optimal*: no other voting system can produce election outcomes that are liked better by the voters, on the average, than those of the GT system. We also look at whether the GT system has several standard properties, such as monotonicity, consistency, etc.

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 and Fishburn [2], Fishburn [11], or Tideman [30], for overviews.

The purpose of this paper is to describe a preferential voting system, called the "GT" voting system, to study its properties, and to compare it with some previous voting system proposals.

The GT system appears to be new (which, if true, is a surprise to us, given the great volume of literature already written on voting systems). However, it is nonetheless quite similar to a proposal by Laffond et al. [18] for parties to pick platform issues, a situa-

*Version April 8, 2010

tion attributed by Shubik [28] to Downs [6]. It may well be the case that the idea of using game theory in the manner we propose was privately but widely understood to be a possibility, but wasn't explicitly promulgated because it involved randomization.

Candidates and ballots We assume an election (contest) whose purpose is to select a single winner from a set of m alternatives. We call the alternatives "candidates." Let n denote the number of voters.

We restrict attention to *preferential voting systems*, where each voter's ballot lists the candidates in decreasing order of preference: first choice, second choice, and so on.[1] The reader may assume for now that all ballots are *full* (they list all the candidates); Section 6 explains how variations such as truncated ballots and write-ins are easily handled.

Profiles, preference and margin matrices, and margin graphs A collection C of (cast) ballots is called a *profile*. A profile is a multi-set, since two ballots may have the same rank order listing of the candidates. The size $|C|$ of a profile is the number of ballots it contains.

Given a profile, one can derive the associated *preference matrix* N— the $m \times m$ matrix whose (x, y) entry is the number of ballots in the profile that express a preference for candidate x over candidate y. Each such entry is nonnegative, and

$$N(x, y) + N(y, x) = n \ , \qquad (1)$$

since all ballots are assumed to be full.

It is also useful to work with the *margin matrix*— the $m \times m$ matrix M defined by

$$M(x, y) = N(x, y) - N(y, x) \ , \qquad (2)$$

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:

$$M(x, y) = -M(y, x) \ . \qquad (3)$$

[1] There are many voting systems, such as approval voting or range voting, where ballots do not list candidates in order. Our methods do not apply to such systems.

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 are no edges between x and y. A directed graph is called a *tournament* if there is exactly one edge between each pair of vertices. It is called a *weighted tournament* if each such edge has an associated positive numeric weight. Although when the number of voters is large the margin graph for the profile of cast ballots is almost certainly a weighted tournament, a voting system should be well-defined in all cases, including those cases when voters are indifferent between two candidates and there are no edges between the corresponding vertices.

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 proposals 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.

1.1 Example

Suppose we have an election with four candidates `A`, `B`, `C`, `D` and 100 voters. The cast ballots have the profile:

```
(40) A B C D
(30) B C A D
(20) C A B D
(10) C B A D
```

Here 40 voters list `A` first, `B` second, `C` third, `D` fourth. Nobody likes D.

The preference matrix N for this profile is the following.

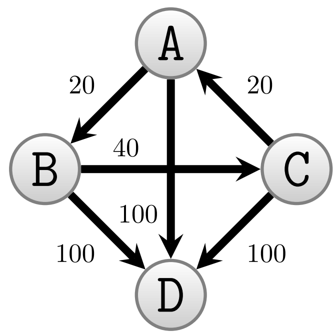

Figure 1: The margin graph for our simple example. Twenty more voters prefer `A` over `B` than prefer `B` over `A`. Candidates `A`, `B`, and `C` are in a cycle and are otherwise unbeaten; so there is no Condorcet winner and this election exhibits a "generalized tie."

	A	B	C	D
A	0	60	40	100
B	40	0	70	100
C	60	30	0	100
D	0	0	0	0

The corresponding margin matrix M is the following.

	A	B	C	D
A	0	20	-20	100
B	-20	0	40	100
C	20	-40	0	100
D	-100	-100	-100	0

Figure 1 gives the associated margin graph.

2 Generalized Ties

A *Condorcet winner* is a candidate x who beats every other candidate in a pairwise comparison; that is, 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 as much as (or more than) x.

In the given example, it is clear that `D` should lose, but it is not so clear which of `A`, `B`, and `C` should win. The margin graph contains a cycle: a majority of voters prefer `A` to `B`, a majority of voters prefer `B` to `C`, and a majority of voters prefer `C` to `A`. Such "Condorcet cycles" were first studied by the eighteenth-century philosopher and mathematician Marie Jean Antoine Nicolas de Caritat, marquis de Condorcet.

Even when there is a generalized tie—and thus no clear winners—there may be clear losers, such as `D` in our example. The "Smith set" is defined to be the smallest set of candidates who collectively dominate all of the other candidates, i.e., every candidate in the Smith set beats every candidate outside of the Smith set in a pairwise comparison. The generalized tie is really just between the candidates in the Smith set. See Smith [29] and Börgers [1] for definition and discussion of Smith sets.[2]

The GT method does not need to compute the Smith set; it suffices for our purposes to consider the generalized tie to be a generalized tie between *all* of the candidates. The GT method nonetheless always names as a winner a candidate in the Smith set (see De Donder et al. [5]).

The interesting question is then:

> *When there is a generalized tie, how should one do the "tie-breaking" to pick a single winner?*

3 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 random-

[2] The definition of the Smith set for this purpose needs careful treatment when the margin $M(x, y)$ is zero—it is actually most convenient to assume that each of x and y "beats" the other when this happens. This is equivalent to showing edges both ways in the margin graph, instead of neither way, as we do now.

ized method is not needed. A randomized method is only appropriate when a tie needs to be broken.

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 doing so can yield election outcomes that are liked by more voters on average than the outcomes of any deterministic voting 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. We discuss this relationship in more detail below.

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." Within the last year, several elections used randomized methods to break ties:

- In June, 2009, the city of Cave Creek, Arizona had a tie between two candidates for a city council seat.[3] The two candidates drew cards from a well-shuffled deck to determine the winner.

- In November, 2009, the office of Mayor of the town of Wendell, Idaho, was determined by a coin toss, after the challenger and the incumbent were tied in the election.

- In February, 2010, in the town of Sealy, Texas, dice were used to resolve a tied election for city council membership.

However, academic literature on voting systems has generally eschewed proposals having a randomized component. For example, Myerson [22, 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.

Several previous voting system proposals also have randomized outcomes.

For example, the "Random Dictator" voting system [12, 27] picks the winner by picking a random ballot, and using it to name the winner. This somewhat silly method always uses randomization, not just for tie-breaking. Gibbard [12] studies the strategy-proofness of randomized voting systems, and argues that if a system is strategy-proof (and satisfies certain other conditions), then it must be the random dictator method.

Sewell et al. [27] proposes 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.

Other proposed voting systems use randomization as a final tie-breaker. For example Schulze's "beat-path" method [26] uses randomization in this manner.

3.1 Returning to our example

The example given in Section 1.1 is interesting because even though there is a generalized tie between `A`, `B`, and `C`, it seems not quite right that they should all have an equal chance of winning. But perhaps they should each have *some* chance of winning.

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

But how should the probabilities be determined? The next sections provide an answer, via game theory. For the example, we shall see that the best probability vector p^* turns out to be:

$$\begin{aligned} p^*(\mathtt{A}) &= 0.500000 \\ p^*(\mathtt{B}) &= 0.250000 \\ p^*(\mathtt{C}) &= 0.250000 \\ p^*(\mathtt{D}) &= 0.000000 . \end{aligned}$$

The computation of these values is explained in Section 6. Thus, A should win half the time, B and C should each win 1/4 of the time, and D should never win. This is easily arranged with a pair of coin flips.

The set of candidates chosen with some nonzero probability is of interest to us. It is just those candidates in the support of p^*.

Definition 3.1 *Let p be an arbitrary probability distribution over some finite set X. We say that the* support *for p is the set*

$$supp(p) = \{x \mid p(x) > 0\}$$

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.

The next section motivates the computation of these probabilities.

4 Optimal Voting Systems

How should one compare a voting system P against another voting system Q? Here P and Q are (perhaps 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 or not. The inference is usually that a system with more desirable properties is the better system. But this approach tends to give rather inconclusive and conflicting advice.

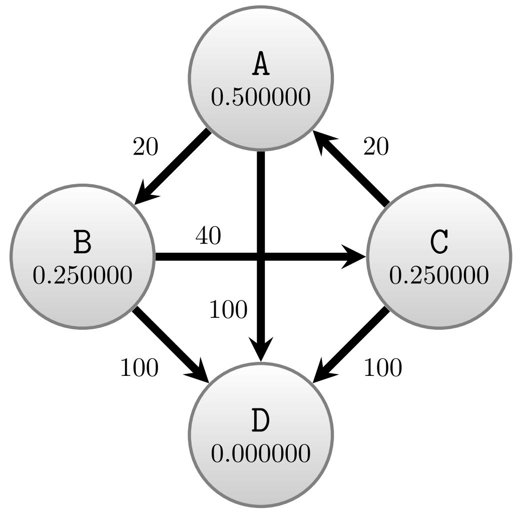

Figure 2: The margin graph for our simple example, with optimal mixed strategy probabilities shown for each candidate.

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 $\mathcal{C}$ turn out not to be that important to us, since the GT method 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 $y = Q(C)$.
- We score the systems 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 4.1 *We say that the* advantage *of voting system P over voting system Q, denoted $\mathbf{Adv}_{\mathcal{C}}(P, Q)$, with respect to the distribution $\mathcal{C}$ on profiles, is*

$$\mathbf{Adv}_{\mathcal{C}}(P, Q) = E_{\mathcal{C}}(M(x, y)) \tag{4}$$

where $x = P(C)$ and $y = Q(C)$, and where $E_{\mathcal{C}}$ denotes expectation with respect to choosing profiles C according to the distribution $\mathcal{C}$ and with respect to any randomization within P and Q. Similarly, we define the relative advantage *of voting system P over voting system Q, denoted $\mathbf{Adv}^*_{\mathcal{C}}(P, Q)$, with respect to the distribution $\mathcal{C}$ on profiles,*

$$\mathbf{Adv}^*_{\mathcal{C}}(P, Q) = E_{\mathcal{C}}(M(x, y)/\,|C|) \tag{5}$$

where $x = P(C)$ and $y = Q(C)$, and where $E_{\mathcal{C}}$ denotes expectation with respect to choosing profiles C 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 $\mathbf{Adv}_C(P, Q)$ and $\mathbf{Adv}^*_C(P, Q)$.

Definition 4.2 *We say that voting system P is as good as or better than voting system Q (with respect to probability distribution $\mathcal{C}$ on profiles), if*

$$\mathbf{Adv}_{\mathcal{C}}(P, Q) \geq 0 \ . \tag{6}$$

Definition 4.3 *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*

$$\mathbf{Adv}_C(P, Q) \geq 0 \ , \tag{7}$$

where E_C denotes expectation with respect to any randomization within P and Q.

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. Since an optimal voting system P^* is not beaten by Q even for any fixed profile C the distribution $\mathcal{C}$ doesn't matter.

5 Game Theory

We now describe how to construct an optimal voting systems using standard game theory computations.

In the game $G_C(P, Q)$, the value $M(x, y)$ 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 [23].

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:

$$\sum_x \sum_y p_x q_y M(x, y) \ . \tag{8}$$

An optimal strategy depends on the margin matrix M. When there is a Condorcet winner, then it is easy to see that it is optimal to pick the Condorcet winner as the election winner.

When there is no Condorcet winner, there is a generalized tie, and the optimal strategy is to play according to an *optimal mixed strategy*. Computing this optimal mixed strategy is not hard; see Section 6. Playing this optimal mixed strategy yields an optimal voting system—no other voting system can produce election outcomes that are preferred more.

The set of candidates that have nonzero probability in the optimal mixed strategy for the game associated with profile C is $\mathrm{supp}(GT(C))$. (When there is not a unique optimal mixed strategy, we assume GT uses the most balanced optimal mixed strategy.) Intuitively, $\mathrm{supp}(GT(C))$ is the set of "potential winners" for the election with profile C for the GT voting system. When there is a Condorcet winner x, then $\mathrm{supp}(GT(C)) = \{x\}$. Otherwise, the GT winner will be chosen from $\mathrm{supp}(GT(C))$ using a randomized procedure as described in Section 7.

6 Computing Optimal Mixed Strategies

How does one solve a two-person zero-sum symmetric game with m by m payoff matrix M? Raghavan [23] gives an excellent overview of two-person zero-sum games and their relationship to linear programming and duality. One of the simplest reductions to linear programming is the following (see Raghavan [23, Problem A, page 740]):

- Increase every entry in M by some constant value w, where w is chosen so that every entry in M is now positive. Note that this changes the value of the game defined by M from 0 to w.

- Solve the following linear programming problem for the probability vector p of length m:
  - Minimize $\Sigma_x p_x$
  - Subject to

  $$p_x \geq 0 \quad \text{for all} \quad x \ , \tag{9}$$

  $$Mp \geq \mathbf{e} \tag{10}$$

  where $\mathbf{e}$ is a column vector of length m containing ones.

- Return the solution vector $p^* = w \cdot p$ — that is, p with every entry multiplied by w.

It is easy to see that this linear programming problem has a solution.

It may have more than one solution, but each such solution has the same value w, which, when translated back to the original game, has a value 0 (as every two-person zero-sum symmetric game has value 0). Each such solution provides an optimal mixed strategy for the original game.

Unique optimal mixed strategies When ballots are full and the number of voters is odd, the optimal mixed strategy p^* is uniquely defined. This follows from a result of Laffond et al. [19]. There are other situations for which there is a unique optimal mixed strategy. In practice, with a large number of voters, one would expect that there would almost always be a unique optimal mixed strategy.

Appendix B describes how we propose that GT should handle the situation when there is not a unique optimal mixed strategy—basically, to pick the unique optimal mixed strategy that minimizes the sum of squares $\sum_i p_i^2$, which can be computed easily with standard quadratic programming packages.

7 Selecting the winner

As we have seen, the GT voting system comprises the following steps, given a profile C of cast ballots:

1. [**Margins**] Compute the margin matrix M.

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.)

We now discuss the third step, and a deterministic variant.

For many, the biggest issue with the GT method may be its use of randomized methods for tie-breaking. Yet, as we have argued, randomized tie-breaking methods are both natural and beneficial.

If there is no Condorcet winner, then GT winner selection should proceed by first computing the cumulative probabilities $q_i = \sum_{j \le i} p_j$. Then a public ceremony should be held where a sequence of (say, six) ten-sided dice are rolled in an indisputably random and non-manipulable manner. Six dice rolls yield a six-digit number $x = 0.d_1d_2d_3d_4d_5d_6$ between 0 and 1. Then the winner is declared to be candidate i where i is the least integer such that $x < q_i$. This is easily seen to elect each candidate i with probability p_i (to within the round-off error).

There are of course details that must be taken care of properly with using a randomized method to select a winner when there is a tie; these details are very similar to those that arise when generating suitable random numbers of post-election audits; see Cordero et al. [4].

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

Instead of randomly picking a candidate according to this probability distribution, as the GT method does, GTD just chooses a 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 doesn't 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 can not prove that it is optimal.

8 Properties of the GT voting system

This section reviews some of the properties of the GT voting system.

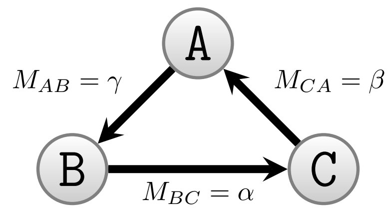

Figure 3: The margin graph for a generic three-cycle.

8.1 Optimality

Optimality is perhaps the most important property of the GT voting system. No other voting system can produce election outcomes that voters prefer better, on the average.

8.2 GT method on three-cycles

The simplest generalized tie has a margin graph that is a three-cycle. See Figure 3

Theorem 8.1 *The optimal mixed strategy probabilities for the three-cycle of Figure 3 are:*

$$p^*_A = \frac{\alpha}{\alpha+\beta+\gamma} \quad (11)$$

$$p^*_B = \frac{\beta}{\alpha+\beta+\gamma} \quad (12)$$

$$p^*_C = \frac{\gamma}{\alpha+\beta+\gamma} \quad (13)$$

Proof See Kaplansky[16, page 479]. ∎

Thus, the probability that the optimal mixed strategy picks a particular vertex in a three-cycle is proportional to the weight of the edge on the opposite side of the cycle: $p^*_{\mathtt{A}}$ is proportional to $M(\mathtt{B},\mathtt{C})$, etc.

Intuitively, as $M(\mathtt{B},\mathtt{C})$ increases, B becomes more attractive to play, and C becomes less attractive, and together these make A substantially more attractive to play, since A beats B and C beats A.

Note that playing any fixed strategy against GT results in zero net value to either player. For example, suppose the opponent picks candidate A and GT picks x. The net point gain for GT is thus

$$\begin{aligned} & p^*_{\mathtt{A}} \cdot 0 + p^*_{\mathtt{B}} \cdot (-\gamma) + p^*_{\mathtt{C}} \cdot \beta \\ = \; & (\alpha/\delta) \cdot 0 + (\beta/\delta) \cdot (-\gamma) + (\gamma/\delta) \cdot \beta \\ = \; & 0 \; . \end{aligned}$$

where $\delta = \alpha + \beta + \gamma$.

Perhaps most surprisingly, the probabilities (11)–(13) remain unchanged if all of the edge directions in Figure 3 are *reversed*! If every voter had submitted a ballot that is the reverse of his submitted ballot, then the GT winning probabilities for each candidate remain unchanged. (This doesn't hold in general, but only for a generalized tie that is a three-cycle.)

8.3 Condorcet winners and losers

Theorem 8.2 *The GT voting system will always pick a Condorcet winner, if one exists. The GT voting system will never pick a Condorcet loser, if one exists.*

Proof (Details omitted.) ∎

As Schulze [26] notes:

> "The Condorcet criterion implies the majority criterion. Unfortunately, compliance with the Condorcet criterion implies violation of other desired criteria like consistency (Young, 1975) [32], participation (Moulin, 1988) [21], later-no-help, and later-no-harm (Woodall, 1997) [31]."

8.4 Monotonicity

The notion of monotonicity needs to be redefined for probabilistic voting systems.

Definition 8.3 *We say that a probabilistic voting system P is* monotonic *if, if a voter raises x on her ballot without changing the order of other candidates, then the probability that P outputs x does not decrease.*

Theorem 8.4 *The GT voting system is not monotonic.*

Proof (Proof sketch.) Note that in the example of Figure 3 and Theorem 8.1, moving `A` in front of `B` on some ballot causes γ to increase, while α and β stay fixed. Thus, $p^*_{\mathtt{A}}$ decreases. ∎

8.5 Reversal Symmetry

The criterion of *reversal symmetry* (see Saari [25]), says that a voting system should not be capable 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 system does not exhibit reversal symmetry: an election where every possible ballot occurs an equal number of times gives each candidate an equal chance of winning. Reversing the ballots doesn't change anything.

GT also gives a three-cycle the same outcome when all ballots are reversed—the probabilities that each candidate wins are unchanged, as noted in Section 8.2.

8.6 Independence of clones

A voting system satisfies the *independence of clones* property if replacing an existing candidate `B` with a set of $k > 1$ *clones* $\mathtt{B}_1, \mathtt{B}_2, \ldots, \mathtt{B}_k$ doesn't change the winning probability for candidates other than `B`. These new candidates are clones in the sense that with respect to the other candidates, voters prefer each $\mathtt{B}_i$ to the same extent that they preferred `B`, and moreover, the voters are indifferent between any two of the clones. (Schulze [26, p. 141] notes that there are some subtleties in the definition of this property, especially when B is already in some sense tied with other candidates.)

The GT voting system satisfies the independence of clones properties in the following sense. If x is an optimal mixed strategy for the game based on the margin matrix for the given election, then when `B` is replaced by $\mathtt{B}_1, \mathtt{B}_2, \ldots, \mathtt{B}_k$ then in the game for the new election it is an optimal mixed strategy to divide `B`'s probability according to x equally among $\mathtt{B}_1, \mathtt{B}_2, \ldots, \mathtt{B}_k$. To see this, note that equations (9) and (10) will continue to hold. However, if the original game does not have a unique optimal mixed strategy, then balancing the probabilities in the derived game may affect probabilities outside of the clones in a different way than the balancing affects those probabilities in the original game—consider what must happen with an empty profile, where all candidates are tied.

8.7 Pareto

The Pareto property says if 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 should never declare y the winner.

While this property is appealing, when x and y are in a cycle or generalized tie it may not be optimal to honor the Pareto property.

It is easy to show that the GT voting system does not have the Pareto property, by considering a three-cycle where *all* voters prefer `A` to `B`, a majority prefer `B` to `C`, and a majority prefer `C` to `A`. The optimal mixed strategy given in Theorem 8.1 will give non-zero probability to each of `A`, `B`, and `C`.

8.8 Strategy-proofness

Our definition allows one to compare two voting systems based on which voting system produces outcomes preferred by the voters, measured by voter preferences as expressed in their ballots. We do not try to take into consideration whether voters might be voting "strategically."

This needs to be studied. We have no reason to believe that GT is more, or less, vulnerable to strategic voting than other preferential voting systems.

8.9 Other properties

There are numerous properties that have been defined for voting systems. Many of them need to be redefined for probabilistic voting systems.

A later version of this paper will include an expanded section here with suitably expanded definitions and results.

9 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 advantage $\mathbf{Adv}_{\mathcal{C}}(P,Q)$ (or relative advantage $\mathcal{A}^*_{\mathcal{C}}(P,Q)$) of one system over the other.

For example, consider the seven voting systems: plurality, IRV, Borda, minimax, Schulze's "beatpath" method [26], GTD, and GT. We used the minimax variant based on margins and the beatpath variant based on "winning votes."

We worked with 10,000 profiles generated randomly for $m = 5$ candidates. Each profile had $n = 100$ full ballots, generated as follows. 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 choice of parameters, about 64.3% of the profiles have a Condorcet winner.

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

Figure 4 gives the cumulative net "point advantage" of each of the seven voting systems against each other in our experiment. For example, the "16380" entry in row "beatpath," column "IRV" means that in an average election, the net number of voters preferring the beatpath outcome to the IRV outcome is about 1.6380 voters (i.e., 1.6380% of the electorate). That is, $\mathbf{Adv}^*_{\mathcal{C}}(\text{beatpath}, \text{IRV}) \approx 0.016380$.

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

In our experiments about 77.13% of the 10,000 simulated elections had a unique optimal mixed strategy.

We note that when comparing another voting sys-

	plurality	IRV	Borda	minimax	beatpath	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
beatpath	32128	16380	4654	58	0	-1402	-76
GTD	32390	15892	5324	1436	1402	0	10
GT	29978	13872	2522	174	76	-10	0

Figure 4: Cumulative margins table for our main experiment. The entry in row X column Y gives the sum, over 10,000 simulated elections with 100 votes each, of the number of voters preferring X's outcome to Y's outcome, minus the number of voters preferring Y's outcome to X's outcome. 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, $\mathbf{Adv}^*_{\mathcal{C}}(GT, IRV) = 1.3872\%$.

	plurality	IRV	Borda	minimax	beatpath	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
beatpath	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 5: 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.

tem with GT, that there is no expected net point gain for GT if the other system picks a candidate that is in supp($GT(C)$), the set of potential winners for GT. Candidates in 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. However, the other system playing other candidates will normally result in an expected positive net point gain for GT, and an expected loss for the other system. Figure 5 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 (only agreeing with GTS 55.15% of the time, as does IRV (72.99%) but beatpath (99.51%) and minimax (99.15%) are nearly perfect.

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 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 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($GT(C)$).

As another example, Schulze gives the following example [26, Section 3.6.2 p. 78]. Figure 6 gives the matrix of pairwise preferences. The beatpath method selects B as the winner. The GT method has the unique optimal mixed strategy:

$$p^*(\mathtt{A}) = 0.333333, \quad p^*(\mathtt{C}) = 0.400000, \quad p^*(\mathtt{D}) = 0.266667,$$
$$p^*(\mathtt{B}) = p^*(\mathtt{E}) = 0.000000$$

so that supp($GT(C)$) $= \{\mathtt{A},\mathtt{C},\mathtt{D}\}$. That is, the beatpath method chooses as a winner a candidate B that is not even a potential winner for the GT method; the GT method gives B no support. For this election, the GT method wins an expected $0.333333*(18-12)+0.400000*(16-14)+0.266667*(13-17) = 1.733333$ points. The GT outcome is preferred by almost 2 more voters than the beatpath outcome, on the average. This is more than 5% of the electorate (30 voters). We also note that the GTD outcome, C, is preferred by exactly two more voters than the beatpath outcome. (It was a bit surprising to us to notice this example in Schulze's paper, given the impressive accuracy with which beatpath generally picks winners from the support of GT.)

	A	B	C	D	E
A	0	18	11	21	21
B	12	0	14	17	19
C	19	16	0	10	10
D	9	13	20	0	30
E	9	11	20	0	0

Figure 6: An example preference matrix provided by Schulze. For this example the beatpath method does not select a candidate that is in the support of the optimal mixed strategy.

10 Practical considerations

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

Note that since the GT voting system only depends on the pairwise preference matrix N, and since the preference matrix for the combination of two profiles is just the sum of the preference matrices for the two profiles, 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 (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 isn't so clear how to efficiently "audit" a GT election. (The last property is common to many preferential voting systems).

11 Variations and Extensions

Working with Truncated Ballots and Write-ins So far, we have assumed that each ballot lists all candidates. The GT method works equally well when ballots may be truncated, or where there are write-ins. A *truncated* ballot lists only some of the candidates, where we assume that the unlisted candidates are assumed to be preferred less than the listed candidates, but equal to each other. More generally, voters may be allowed to specify that two candidates rank equally. For our purposes, the important fact is that for each pair of distinct candidates, a ballot expresses a preference one way or the other, or expresses no preference. Write-in candidates are also handled smoothly within this framework.

Nothing really changes; the margin matrix M is still the payoff matrix for a two-person zero-sum game. The computation proceeds as before. Other properties of the GT method are unaffected.

12 Discussion

Note that the GT voting system is directed towards preferential voting systems; it is not applicable to scoring systems (like approval voting or range voting) where cycles can't occur. (You can have only simple ties with equality of total scores, which are easily handled.)

When a voting system is randomized, the notion of "margin of victory" needs to be redefined, and the methods of post-election audits need to be correspondingly adjusted. Further research is needed to clarify this situation.

13 Related Work

Fishburn [10] gives an excellent overview of voting systems that satisfy the Condorcet principle.

The idea of using a two-person 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. [17] introduce the notion of a "bipartisan set", which is the support of a two-person "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 such 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 study (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. [18] explicitly propose the use of two-party game theory to provide solutions to elections, including the use of randomized methods. They call a weighted tournament game a *plurality game*. 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 [3, 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 [7] study the "minimal covering set" of a tournament (as proposed by Dutta [8]), which is the same as Shapley's notion of a "weak saddle" for the corresponding tournament game.

De Donder et al. [5] consider tournament games and plurality games and the relationship of the support of the optimal mixed strategy for a plurality game (the "weighted bipartisan set") to various other set-theoretic notions.

Michael and Quint [20] provide further results on optimal strategies in tournament games and plurality games.

Dutta et al. [9] study *comparison functions*, which correspond to general skew-symmetric matrices, and study axiomatic properties associated with such functions.

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 more voters.

We feel optimality is an important criterion for voting systems. It would seem hard to argue that some other property X was sufficiently important that in return for obtaining property X one should settle for reducing the average number of voters preferring the election outcome.

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

Since the GT system does share 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 Schulze's "beatpath" method, or even 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 balloting is to be used.

Acknowledgments

We thank the Stanford University Computer Science Department and Microsoft Research (Redmond) 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. Thanks also to Rob Schapire and Josh Benaloh for helpful suggestions and discussions.

References

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

[2] Steven J. Brams and Peter C. Fishburn. Voting procedures. In Kenneth Arrow, Amartya Sen, and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, pages 173–236. Elsevier, 2002.

[3] Michel 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.

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

[5] Philippe De Donder, Michel Le Breton, and Michel Truchon. Choosing from a weighted tournament. *Mathematical Social Sciences*, 40(1):85–109, Jul 2000.

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

[7] John Duggan and Michel Le Breton. Dutta's minimal covering sets and Shapley's saddles. *J. Economic Theory*, 70:257–265, 1996.

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

[9] Bhaskar Dutta and Jean-Francois Laslier. Comparison functions and choice correspondences. *Soc. Choice Welfare*, 16:513–532, 1999.

[10] Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, Nov 1977.

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

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

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

[14] P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic Press, 1981.

[15] E. T. Jaynes. Information theory and statistical mechanics I. *Physical Review*, 106:620–630, 1957.

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

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

[18] Gilbert Laffond, Jean François Laslier, and Michel Le Breton. Social choice mediators. *The American Economic Review*, 84(2):448–453, May 1994. Proc. 106th Annual Mtg. of the American Economic Association.

[19] Gilbert Laffond, Jean-Francois Laslier, and Michel Le Breton. A theorem on symmetric two-person zero-sum games. *Journal of Economic Theory*, 72:426–431, 1997.

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

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

[22] Roger B. Myerson. Fundamentals of social choice theory, Sep 1996. Discussion paper 1162. Center for Mathematical Studies in Economics and Management Science. Northwestern Univ., Evanston, IL.

[23] TES 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.

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

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

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

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

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

[29] John H. Smith. Aggregation of preferences with a variable electorate. *Econometrica*, 41(6):1027–1041, Nov 1973.

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

[31] Douglas Woodall. Monotonicity of single-seat preferential election rules. *Discrete Applied Mathematics*, 77:81–98, 1997.

[32] H. P. Young. Social choice scoring functions. *SIAM Journal on Applied Mathematics*, 28(4):824–838, Jun 1975.

Appendix A. When there is not a unique optimal mixed strategy

It may happen that there is not a unique optimal mixed strategy. for the two-person zero-sum symmetric game defined by M.

In this case, one might say that "it doesn't matter" which optimal strategy is used to break a generalized tie, since any such optimal mixed strategy maximizes the number of voters preferring the voting system outcome.

As a trivial example, if the matrix M of margins is all zeros, then *any* probability vector p is an optimal mixed strategy.

While some might feel that this lack of a unique optimal mixed strategy is a defect, it is in fact an opportunity. Given that we have done as well as we can to follow voter preferences, we can then use the additional degrees of freedom to treat the candidates equitably. That is, we can use any additional degrees of freedom to ensure that ties are broken in a "balanced" manner.

By using a least-squares approach, we can find the optimal mixed strategy that is also as "balanced" as possible, where "balanced" now means "minimizing $\sum_x (p^*_x)^2$." Such a minimization will tend to make the probabilities p^*_x as equal as possible, given the constraint that p^* be an optimal mixed strategy.

A least-squares approach is very convenient, since there are excellent software packages available for quadratic programming with constraints, such as the MATLAB routine `lsqlin`)[4] or the `cvxopt` package for Python.

More precisely, the new optimization method works as follows, given as input a skew-symmetric real matrix M.

[4]See the MATLAB documentation for `lsqlin` and Gill et al. [14]. A handful of lines of MATLAB code suffice to produce the desired optimal mixed strategy. Code is available at `http://people.csail.mit.edu/rivest/gt/`.

- Increase every entry in M by some constant value w, where w is chosen so that every entry in M is now positive. Note that this changes the value of the game defined by M from 0 to w.
- Solve the following quadratic programming problem for the probability vector p of length m:
  - Minimize $\Sigma_x p_x^2$.
  - Subject to

$$p_x \geq 0 \text{ for all } x \;, \tag{14}$$

$$\sum_x p_x = 1/w \tag{15}$$

$$Mp \geq \mathbf{e} \tag{16}$$

  where $\mathbf{e}$ is a column vector of length m containing ones.
- Return the solution vector $p^* = w \cdot p$ — that is, p with every entry multiplied by w.

Since the feasible region is convex and nonempty, and the function being optimized is convex, there will be a unique solution to this quadratic programming problem. The relationship with the previous optimization problem should be evident, given that the optimum value for $\sum_x p_x$ in the previous problem was known to be $1/w$.

The above procedure, together with a randomized method for picking the winner according to p^*, constitutes the voting system GT.

Maximum entropy variant There are plausible approaches other than least squares for finding an optimal mixed strategy that is "balanced." For example, one could aim for the *maximum entropy* solution that maximizes $\sum_x p_x \ln(1/p_x)$, instead of minimizing $\sum_x p_x^2$. (See Jaynes [15] for an introduction to maximum entropy methods). We suspect that using maximum-entropy methods would not produce noticeably different results in practice, and there are practical advantages to working with least-squares, given the availability of software packages for quadratic programming.

Appendix B. More examples

This appendix gives a number of simple examples. In each diagram, each vertex gives the candidate's name above the candidate's winning probability as computed by the GT method. The figure name `m_k` gives the number of candidates m and the number of edges k.

`2_0`

`2_1`

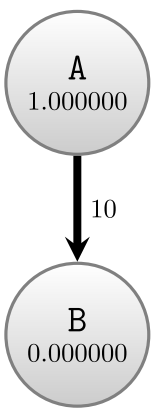

`3_0`

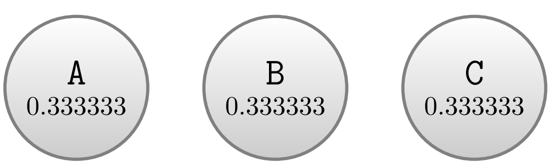

`3_1`

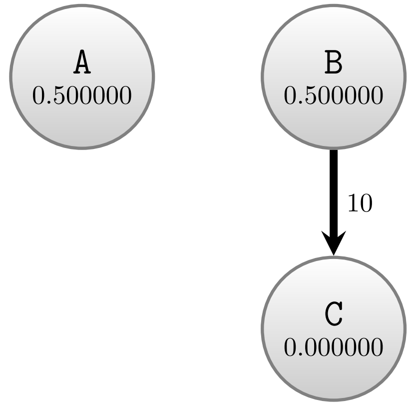

`3_2a`

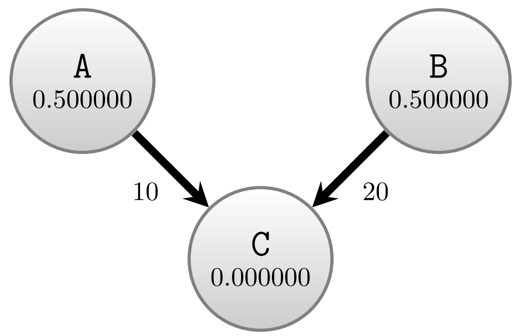

`3_2b`

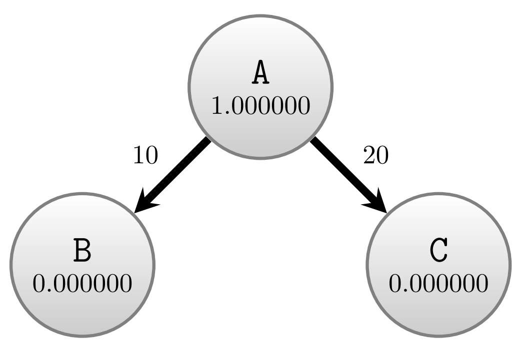

`3_3a`

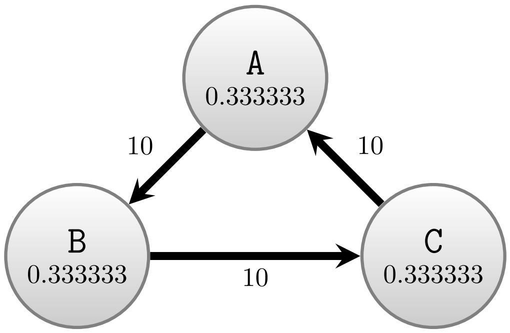

`3_3b`

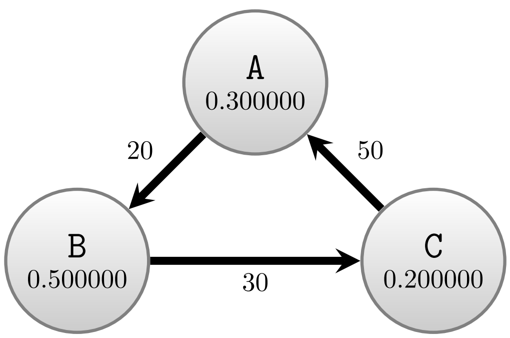

`4_2a`

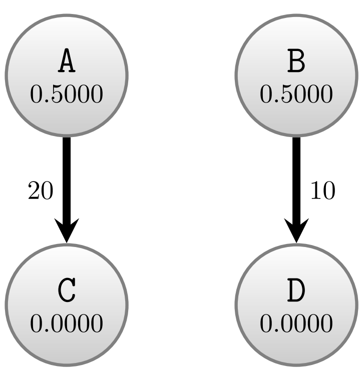

`4_3a`

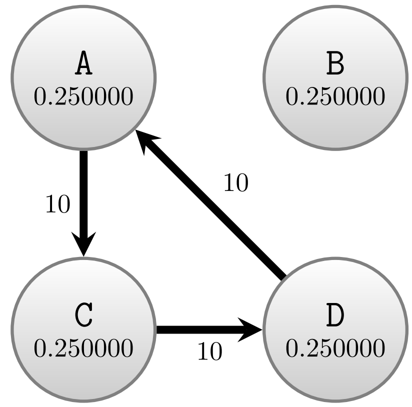

`4_3b`

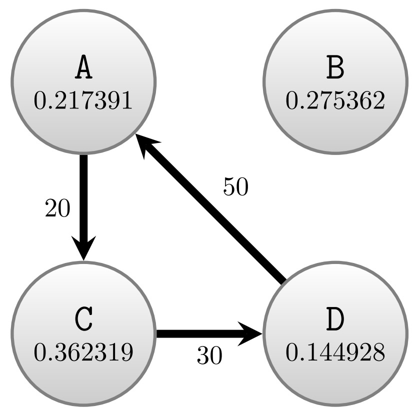

`4_4a`

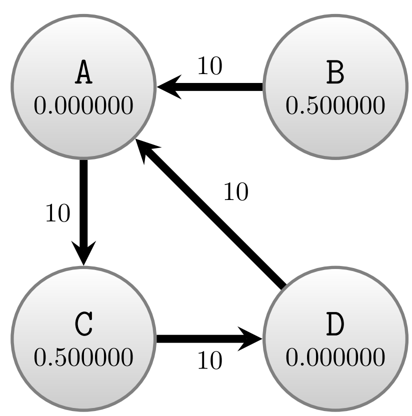

`4_4b`

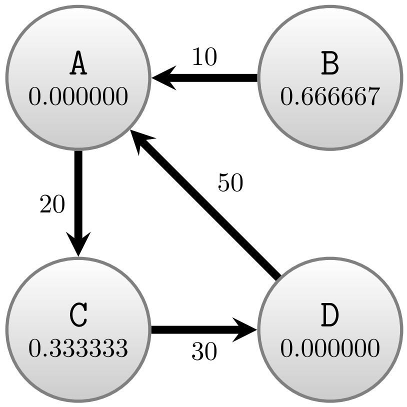

`4_5b`

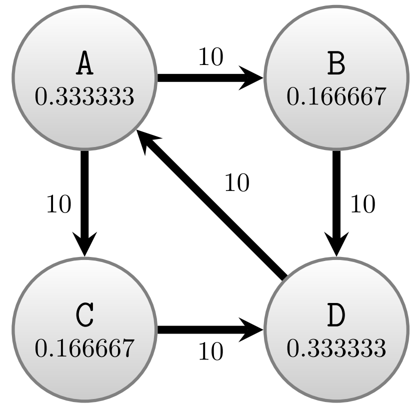

`4_4c`

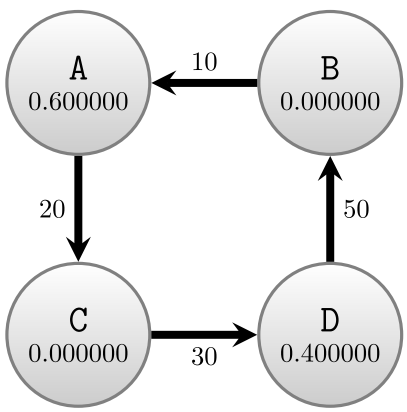

`4_6a`

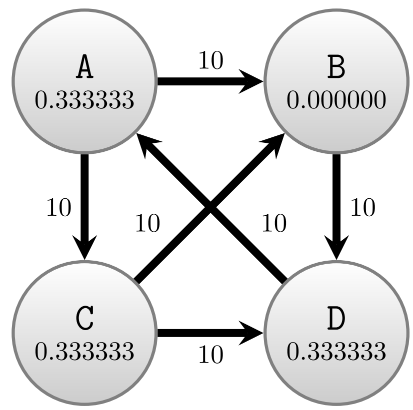

`4_5a`