

A New Perspective on Implementation by Voting Trees

Felix Fischer*

Ludwig-Maximilians-Universität München
80538 München, Germany
fischerf@tcs.ifi.lmu.de

Ariel D. Procaccia†

Microsoft Israel R&D Center
Herzeliya 46725, Israel
arielpro@cs.huji.ac.il

Alex Samorodnitsky‡

Hebrew University of Jerusalem
Jerusalem 91904, Israel
salex@cs.huji.ac.il

ABSTRACT

Voting trees provide an abstract model of decision-making among a group of individuals in terms of an iterative procedure for selecting a single vertex from a tournament. A family of voting trees is said to implement a given voting rule if for every tournament it chooses according to the rule. While partial results concerning implementable rules and necessary conditions for implementability have been obtained, a complete characterization of voting rules implementable by trees has proven surprisingly hard to find. A prominent rule that *cannot* be implemented by trees is the Copeland rule, which singles out vertices with maximum degree. In this paper, we suggest a new angle of attack and re-examine the implementability of the Copeland solution using paradigms and techniques at the core of theoretical computer science. We study the extent to which voting trees can *approximate* the maximum degree, and give upper and lower bounds on the worst-case ratio between the degree of the vertex chosen by a tree and the maximum degree, both for the deterministic model concerned with a single fixed tree, and for randomizations over arbitrary sets of trees. Our main positive result is a randomization over surjective trees of polynomial size that provides an approximation ratio of at least $1/2$. The proof is based on a connection between a randomization over caterpillar trees and a rapidly mixing Markov chain.

Categories and Subject Descriptors: F.2.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity—*Nonnumerical Algorithms and Problems*; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

*Part of the work was done while the author was visiting the Hebrew University of Jerusalem. This visit was supported by a Minerva Short-Term Research Grant.

†The work was done while the author was at the Hebrew University of Jerusalem, and was supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

‡The author is supported by ISF grant 039-7165.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'09, July 6–10, 2009, Stanford, California, USA.

Copyright 2009 ACM 978-1-60558-458-4/09/07 ...\$5.00.

General Terms: Theory, Economics

Keywords: Computational Social Choice, Voting Trees, Copeland Rule, Approximation

1. INTRODUCTION

Voting is a general method for aggregating the preferences of a set of individuals. In the case of two alternatives, the voting rule that selects the alternative preferred by a majority of the voters, commonly known as the *majority rule*, is uniquely characterized by anonymity, neutrality, and monotonicity [21]. It is thus the only “reasonable” rule for two alternatives that is symmetric with respect to both voters and alternatives. A natural way to generalize the majority rule to more than two alternatives is to base the choice on an asymmetric *dominance relation* over the set of alternatives, in which one alternative dominates the other if it is preferred by a majority of the voters. Assuming an odd number of voters with linear preferences this relation is a *tournament*, *i.e.*, both asymmetric and complete. In graph theoretic terms, a tournament is an orientation of the complete graph on a set of alternatives, with a directed edge from a dominating alternative to a dominated one.

Unfortunately, the situation becomes much less favorable when moving to this more general setting. The dominance relation may contain cycles [7], and in fact any tournament can be obtained from a set of voters of size polynomial in the number of alternatives [22]. As a consequence, an important area in social choice theory is concerned with so-called *tournament solutions* that map each tournament to a nonempty set of “good” alternatives in the absence of maximal, *i.e.*, undominated, ones (see, *e.g.*, [20]).

Voting trees provide an abstract model for decision-making among a group of individuals. They were first studied by Farquharson [9], and have consequently been applied to the analysis of a wide range of decision procedures (see, *e.g.*, [29, 1], and the references therein). A voting tree over a set A of alternatives is a binary tree with leaves labeled by elements of A . Given a tournament T , a labeling for the internal nodes is defined recursively by labeling a node by the label of its child that beats the other child according to T (or by the unique label of its children if both have the same label). The label at the root is then deemed the winner of the voting tree given tournament T . This definition expressly allows an alternative to appear multiple times at the leaves of a tree.

A voting tree over A is said to *implement* a particular solution concept if for every tournament on A it selects an optimal alternative according to said solution concept.

While there have been some partial results concerning particular implementable rules and necessary conditions for implementability, an axiomatic characterization of voting rules implementable by trees remains one of the most intriguing open problems in social choice theory (see, *e.g.*, [26, 31]).

It should be noted that the interest in implementation by voting trees does not stem from their attractiveness as an algorithmic procedure, but rather from their importance as an abstract model of decision-making. To this end, voting trees were first studied in the context of so-called “binary procedures” [9]. Assume that the set of alternatives is split into two, possibly overlapping, subsets. Then, one alternative is chosen from each of the subsets by applying the idea recursively, and finally the voters decide which of these two alternatives is to be chosen. It is easy to see that each instance of this procedure induces a voting tree, with subtrees corresponding to the subsets of alternatives. While the division into two subsets rather than, say, three, seems a bit arbitrary, it nevertheless makes perfect sense given that we know how to choose from two alternatives, namely by majority vote. The class of binary procedures furthermore includes many procedures that are relevant in practice, like “those of Parliament [...] and the usual committee method of voting on amendments to a motion. It also includes some electoral procedures, such as that whereby each candidate is considered in turn for acceptance, and the Scots’ practice of election by paired comparisons” ([9], p. 9).

Binary procedures possess the following remarkable property: if voters have strict preferences over the set of alternatives, and are assumed to vote strategically at each division into subsets in order to achieve the best possible outcome according to their own preference order, then there always exists a unique outcome that survives the iterated elimination of dominated strategies. It turns out that the alternative elected in this outcome is precisely the one chosen by the corresponding voting tree. Implementability by voting trees further implies implementability by backward induction, which is considered one of the most appealing solution concepts in terms of the plausibility of the underlying strategic behavior [15, 8].

A particular solution that *cannot* be implemented by voting trees is the *Copeland solution*, which selects a vertex with maximum (out-)degree [26]. This is rather surprising, because degree is conceivably one of the most basic properties to single out vertices in a tournament, and it is also very easy to find a vertex with maximum degree. In other words, although the Copeland rule is computationally tractable, it cannot be *rationalized* in terms of the iterative comparison-based model of choice provided by voting trees.

In this paper, we suggest a new angle to attack the problem of characterizing implementable rules, and re-examine implementability of the Copeland solution using paradigms and techniques that are at the core of theoretical computer science. Indeed, we study the extent to which voting trees can *approximate* the maximum degree in a tournament, by asking for the largest value of α , such that for any set A of alternatives, there exists a tree Γ , which for every tournament on A selects an alternative with at least α times the maximum degree in the tournament. We address this question both in the *deterministic* model, where Γ is a fixed voting tree, and in the *randomized* model, where voting trees are chosen randomly according to some distribution. Conceptually, randomizations over voting trees provide an abstract

model for decision-making processes that include an element of chance, just like voting trees do for deterministic rules, and we find it quite surprising that seemingly they have not been considered before.¹

This is not the first time that theoretical computer science provides a new perspective on problems in social choice theory. In the end, any model of decision-making is also a computational model, and by studying the power and limitations of the latter in computing certain functions, one directly obtains insights about the former. A prominent example where this observation has been applied successfully is the agenda of using computational complexity as a barrier against manipulation in elections (see, *e.g.*, [2, 5, 34, 12]), thereby circumventing the Gibbard-Satterthwaite impossibility result [13, 28].

Results. Our main negative results are upper bounds of $3/4$ and $5/6$, respectively, on the approximation ratio achievable by deterministic trees and randomizations over trees. It should be noted that this directly implies a strong negative result for exact implementability: *no* rule guaranteed to select an alternative preferred by a majority over at least $3/4$ of the other alternatives can be implemented by voting trees. We also find it quite surprising that randomizations over trees cannot achieve a ratio arbitrarily close to 1.

For most of the paper we concentrate on the randomized model. We study a class of trees we call voting caterpillars, which are characterized by the fact that they have exactly two nodes on each level below the root. We devise a randomization over “small” trees of this type, which crucially satisfies an important property we call *admissibility*: its support only contains trees where every alternative appears in some leaf. Our main positive result is the following.

Theorem 4.1. *Let A be a set of alternatives. Then there exists an admissible randomization over voting trees on A of size polynomial in $|A|$ with an approximation ratio of $1/2 - \mathcal{O}(1/|A|)$.*

We prove this theorem by establishing a connection to a nonreversible, rapidly mixing random walk on the tournament, and analyzing its stationary distribution. The proof of rapid mixing involves reversibilizing the transition matrix, and then bounding its spectral gap via its conductance. To the best of our knowledge, this constitutes the first use of rapid mixing, and in particular of notions like conductance, as a proof technique in Computational Economics. We further show that our analysis is tight, and that voting caterpillars also provide a lower bound of $1/2$ for the second order degree of an alternative, defined as the sum of degrees of those alternatives it dominates.

The paper concludes with negative results about more complex tree structures, which turn out to be rather surprising. In particular, we show that the approximation ratio provided by randomized balanced trees can become arbitrarily bad with growing height. We finally note that “higher-order” caterpillars, with labels chosen by lower-order caterpillars instead of uniformly at random, can also cause the approximation ratio to deteriorate.

Related work. In economics, the problem of implementation by voting trees was introduced by Farquharson [9],

¹Probabilistic choice rules have in fact been studied extensively, for example by Zeckhauser [33], Fishburn [11], Intriligator [16], and Gibbard [14].

and further explored, for example, by McKelvey and Niemi [23], Miller [24], Moulin [26], Herrero and Srivastava [15], Dutta and Sen [8], Srivastava and Trick [31], and Coughlan and Breton [6]. In particular, Moulin [26] shows that the Copeland solution is not implementable by voting trees if there are at least 8 alternatives, while Srivastava and Trick [31] demonstrate that it can be implemented for tournaments with up to 7 alternatives. Laffond et al. [18] compute the *Copeland measure* of several prominent *choice correspondences*—functions mapping each tournament to a set of desirable alternatives. In contrast to the (Copeland) approximation ratio considered in this paper, the Copeland measure is computed with respect to the best alternative selected by the correspondence, so strictly speaking it is not a worst-case measure. More importantly, however, Laffond et al. study properties of given correspondences, whereas we investigate the possibility of *constructing* voting trees with certain desirable properties. In this sense, our work is algorithmic in nature, while theirs is descriptive.

Implementability in the randomized model is related to work that studies the influence of the strength of a player on his winning probability in knockout tournaments (see, e.g., [4], and the references therein).

In theoretical computer science, the problem studied in this paper is somewhat reminiscent of the problem of determining the query complexity of graph properties (see, e.g., [17]). In the general model, one is given an unknown graph over a known set of vertices, and must determine whether the graph satisfies a certain property by querying the edges. The complexity of a property is then defined as the height of the smallest decision tree that checks the property. Voting trees can be interpreted as querying the edges of the tournament in parallel, and in a way that severely limits the ways in which, and the extent up to which, information can be transferred between different queries.

In the area of computational social choice at the boundary of computer science and economics, several authors have looked at the computational properties of voting trees and of various solution concepts. For example, Lang et al. [19] characterize the computational complexity of determining different types of winners in voting trees. Procaccia et al. [27] investigate the learnability of voting trees, as functions from tournaments to alternatives. In a slightly different context, Brandt et al. [3] study the computational complexity of different solution concepts, including the Copeland solution.

Organization. We begin by introducing the necessary concepts and notation. In Section 3 we present upper bounds for the deterministic and the randomized setting. In Section 4, we establish our main positive result using a randomization over voting caterpillars. Section 5 is devoted to balanced trees, and Section 6 concludes with some open problems. Some proofs, as well as an analysis of “higher order” caterpillars, were omitted due to lack of space and can be found in the full version of the paper.

2. PRELIMINARIES

Let $A = \{1, \dots, m\}$ be a set of *alternatives*. A *tournament* T on A is an orientation of the complete graph with vertex set A . We denote by $\mathcal{T}(A)$ the set of all tournaments on A . For a tournament $T \in \mathcal{T}(A)$, we write iTj if the edge between a pair $i, j \in A$ of alternatives is directed from i to j , or i *dominates* j . For an alternative $i \in A$ we denote

by $s_i = s_i(T) = |\{j \in A : iTj\}|$ the *degree* or (Copeland) *score* of i , i.e., the number of outgoing edges from this alternative, omitting T when it is clear from the context.

We then consider computations performed by a specific type of tree on a tournament. In the context of this paper, a (deterministic) *voting tree* on A is a structure $\Gamma = (V, E, \ell)$ where (V, E) is a binary tree with root $r \in V$, and $\ell : V \rightarrow A$ is a mapping that assigns an element of A to each leaf of (V, E) . Given a tournament T , a unique function $\ell_T : V \rightarrow A$ exists such that $\ell_T(v) = \ell(v)$ if v is a leaf, and $\ell_T(v) = \ell(u_1)$ if v has children u_1 and u_2 , and if $\ell(u_1)T\ell(u_2)$ or $\ell(u_1) = \ell(u_2)$. We will be interested in the label of the root r under this labeling, which we call the winner of the tree and denote by $\Gamma(T) = \ell_T(r)$. We call a tree Γ *surjective* if ℓ is surjective. Obviously, surjectivity corresponds to a very basic fairness requirement on the solution implemented by a tree. Other authors therefore view surjectivity as an inherent property of voting trees and define them accordingly (see, e.g., [26]). The sole reason we do not require surjectivity by definition is that our analysis will use trees that are not necessarily surjective.

Finally, a voting tree Γ on A will be said to provide an approximation ratio of α (w.r.t. the maximum degree) if

$$\min_{T \in \mathcal{T}(A)} \frac{s_{\Gamma(T)}}{\max_{i \in A} s_i(T)} \geq \alpha.$$

The above model can be generalized by looking at *randomizations* over voting trees according to some probability distribution. We will call a randomization *admissible* if its support contains only surjective trees. A distribution Δ over voting trees will then be said to provide a (randomized) approximation ratio of α if

$$\min_{T \in \mathcal{T}(A)} \frac{\mathbb{E}_{\Gamma \sim \Delta} [s_{\Gamma(T)}]}{\max_{i \in A} s_i(T)} \geq \alpha.$$

While we are of course interested in the approximation ratio achievable by admissible randomizations, it will prove useful to consider a specific class of randomizations that are not admissible, namely those that choose uniformly from the set of all voting trees with a given structure. Equivalently, such a randomization is obtained by fixing a binary tree and assigning alternatives to the leaves independently and uniformly at random, and will thus be called a *randomized voting tree*.

3. UPPER BOUNDS

In this section we present upper bounds on the approximation ratio achievable by voting trees, both in the deterministic model and in the randomized model. We build on concepts and techniques introduced by Moulin [26], and begin by quickly familiarizing the reader with these.

Given a tournament T on a set A of alternatives, we say that $C \subseteq A$ is a *component*² of T if for all $i_1, i_2 \in C$ and $j \in A \setminus C$, i_1Tj if and only if i_2Tj . For a component C , denote by \mathcal{T}_C the subset of tournaments that have C as a component. If $T \in \mathcal{T}_C$, we can unambiguously define a tournament T_C on $(A \setminus C) \cup \{C\}$ by replacing the component C by a single alternative. The following lemma states that for two tournaments that differ only inside a particular component, any tree chooses an alternative from that component for one of the tournaments if and only if it does for the other.

²Moulin [26] uses the term “adjacent set”.

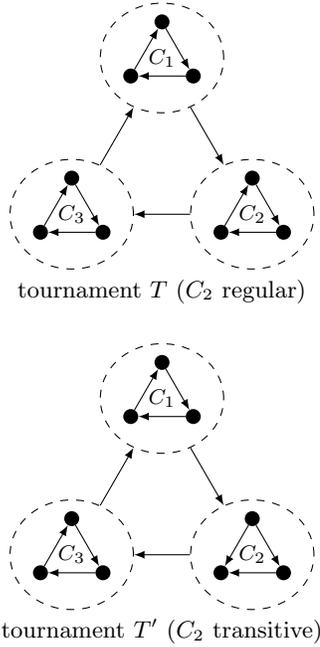


Figure 1: Tournaments used in the proof of Theorem 3.2, illustrated for $k = 3$. A voting tree is assumed to select an alternative from C_1 .

Furthermore, if an alternative outside the component is chosen for one tournament, then the same alternative has to be chosen for the other. Laslier [20] calls a solution concept satisfying these properties *weakly composition-consistent*.

LEMMA 3.1 (Moulin [26]). *Let A be a set of alternatives, Γ a voting tree on A . Then, for all proper subsets $C \subsetneq A$, and for all $T, T' \in \mathcal{T}_C$,*

1. $[T_C = T'_C]$ implies $[\Gamma(T) \in C \text{ if and only if } \Gamma(T') \in C]$, and
2. $[T_C = T'_C \text{ and } \Gamma(T) \in A \setminus C]$ implies $[\Gamma(T) = \Gamma(T')]$.

Our first result asserts that no deterministic tree can always choose an alternative that has a degree significantly larger than $3/4$ of the maximum degree, thereby strengthening the negative result concerning implementability of the Copeland solution [26].

THEOREM 3.2. *Let A be a set of alternatives, $|A| = m$, and let Γ be a deterministic voting tree on A with approximation ratio α . Then, $\alpha \leq 3/4 + \mathcal{O}(1/m)$.*

PROOF. For ease of exposition, we assume $|A| = m = 3k$ for some odd k , but the same result (up to lower order terms) holds for all values of m . Define a tournament T comprised of three components C_1 , C_2 , and C_3 , such that for $r = 1, 2, 3$, (i) $|C_r| = k$ and the subgraph of T induced by C_r is regular, *i.e.*, each $i \in C_r$ dominates exactly $(k-1)/2$ of the alternatives in C_r , and (ii) for all $i \in C_r$ and $j \in C_{(r \bmod 3)+1}$, iTj . An illustration for $k = 3$ is given on the left of Figure 1.

Now consider any deterministic voting tree Γ on A , and assume w.l.o.g. that $\Gamma(T) \in C_1$. Define T' to be a tournament on A such that the subgraphs of T and T' induced by $B \subseteq A$ are identical if $|B \cap C_2| \leq 1$, and the subgraph of T'

induced by C_2 is transitive; in particular, there is $i \in C_2$ such that for any $i \neq j \in C_2$, $iT'j$. An illustration for $k = 3$ is given on the right of Figure 1. By Lemma 3.1, $\Gamma(T') = \Gamma(T)$. Furthermore, T' satisfies

$$s_{\Gamma(T')} = k + \frac{(k-1)}{2} = \frac{3k}{2} - \frac{1}{2}$$

and $\max_{i \in A} s_i = 2k - 1$, and thus

$$\frac{s_{\Gamma(T')}}{\max_{i \in A} s_i(T')} = \frac{3k-1}{4k-2} \leq \frac{3}{4} + \frac{1}{2(k-1)}. \quad \square$$

We now turn to the randomized model. It turns out that one cannot obtain an approximation ratio arbitrarily close to 1 by randomizing over large trees. We derive an upper bound for the approximation ratio by using similar arguments as in the deterministic case above, and combining them with the minimax principle of Yao [32]. A detailed proof is given in the full version of the paper.

THEOREM 3.3. *Let A be a set of alternatives, $|A| = m$, and let Δ be a probability distribution over voting trees on A with an approximation ratio of α . Then, $\alpha \leq 5/6 + \mathcal{O}(1/m)$.*

We point out that the theorem holds in particular for inadmissible randomizations.

4. A RANDOMIZED LOWER BOUND

A weak deterministic lower bound of $\Theta((\log m)/m)$ can be obtained straightforwardly from a balanced tree where every label appears exactly once. While balanced trees will be discussed in more detail in Section 5, they become increasingly unwieldy with growing height, and an improvement of this lower bound or of the deterministic upper bound given in the previous section currently seems to be out of our reach. In the remainder of the paper, we therefore concentrate on the randomized model. In this section we put forward our main result, a lower bound of $1/2$, up to lower order terms, for admissible randomizations over voting trees. Let us state the result formally.

THEOREM 4.1. *Let A be a set of alternatives. Then there exists an admissible randomization over voting trees on A of size polynomial in $|A|$ with an approximation ratio of $1/2 - \mathcal{O}(1/m)$.*

In addition to satisfying the basic admissibility requirement, the randomization also has the desirable property of relying only on trees of polynomial size. The latter is relevant even when assuming a purely descriptive point of view, since a huge voting tree would not provide a very realistic model of real-world decision-making. To prove Theorem 4.1, we make use of a specific binary tree structure known as caterpillar trees.

4.1 Randomized Voting Caterpillars

We begin by inductively defining a family of binary trees that we refer to as *k-caterpillars*. The 1-caterpillar consists of a single leaf. A *k-caterpillar* is a binary tree, where one subtree of the root is a $(k-1)$ -caterpillar, and the other subtree is a leaf. Then, a *voting k-caterpillar* on A is a *k-caterpillar* whose leaves are labeled by elements of A .

It is straightforward to see that an upper and lower bound of $1/2$ holds for the randomized 1-caterpillar, *i.e.*, the uniform distribution over the m possible voting 1-caterpillars.

Indeed, such a tree is equivalent to selecting an alternative uniformly at random. Since we have $\sum_{i \in A} s_i = \binom{m}{2}$, the expected score of a random alternative is $(m-1)/2$, whereas the maximum possible score is $m-1$. This randomization, however, like other randomizations over small trees that conceivably provide a good approximation ratio, is not admissible and actually puts probability one on trees that are not surjective. This leads to absurdities from a social choice point of view; for instance, in a tournament where there are both a *Condorcet winner*, an alternative that beats every other, and a *Condorcet loser*, which loses to every other alternative, the probabilities (under the above inadmissible randomization) of electing the former and the latter are equal, namely $1/m$. In contrast, any admissible randomization would elect a Condorcet winner with probability 1 given a tournament where one exists, and would elect a Condorcet loser with probability 0 given a tournament where one exists.

To prove Theorem 4.1, we instead use the uniform randomization over surjective k -caterpillars, henceforth denoted k -RSC, which is clearly admissible. Theorem 4.1 can then be restated as a more explicit—and slightly stronger—result about the k -RSC.

LEMMA 4.2. *Let A be a set of alternatives, $T \in \mathcal{T}(A)$. For $k \in \mathbb{N}$, denote by $p_i^{(k)}$ the probability that alternative $i \in A$ is selected from T by the k -RSC. Then, for every $\epsilon > 0$ there exists $k = k(m, \epsilon)$ polynomial in m and $1/\epsilon$ such that*

$$\sum_{i \in A} p_i^{(k)} s_i \geq \frac{m-1}{2} - \epsilon.$$

The lemma directly implies Theorem 4.1 by letting $\epsilon = 1$ and recalling that the maximum score is $m-1$. The remainder of this section is devoted to the proof of this lemma. For the sake of analysis, we will use the randomized k -caterpillar, or k -RC, as a proxy to the k -RSC. We recall that the k -RC is equivalent to a k -caterpillar with labels for the leaves chosen independently and uniformly at random. In other words, it corresponds to the uniform distribution over all possible voting k -caterpillars, rather than just the surjective ones.

Clearly the k -RC corresponds to a randomization that is not admissible. In contrast to very small trees, however, like the one consisting only of a single leaf, it is straightforward to show that the distribution over alternatives selected by the RC is very close to that of the RSC.

LEMMA 4.3. *Let $k \geq m$, and denote by $\bar{p}_i^{(k)}$ and $p_i^{(k)}$, respectively, the probability that alternative $i \in A$ is selected by the k -RC and by the k -RSC for some tournament $T \in \mathcal{T}(A)$. Then, for all $i \in A$,*

$$|\bar{p}_i^{(k)} - p_i^{(k)}| \leq \frac{m}{e^{k/m}}.$$

PROOF. For all $i \in A$, $|\bar{p}_i^{(k)} - p_i^{(k)}|$ is at most the probability that the k -RC does not choose a surjective tree. By the union bound, we can bound this probability by

$$\sum_{i \in A} \Pr[i \text{ not in the } k\text{-RC}] \leq m \cdot \left(1 - \frac{1}{m}\right)^k \leq \frac{m}{e^{k/m}}. \quad \square$$

With Lemma 4.3 at hand, we can temporarily restrict our attention to the k -RC. A direct analysis of the k -RC, and in particular of the competition between the winner

of the $(k-1)$ -RC and a random alternative, shows that for every k , the k -RC provides an approximation ratio of at least $1/3$. It seems, however, that this analysis cannot be extended to obtain an approximation ratio of $1/2$. In order to reach a ratio of $1/2$, we shall therefore proceed by employing a second abstraction. Given a tournament T , we define a Markov chain $\mathfrak{M} = \mathfrak{M}(T)$ as follows.³ The state space Ω of \mathfrak{M} is A , and its initial distribution $\pi^{(0)}$ is the uniform distribution over Ω . The transition matrix $P = P(T)$ is given by

$$P(i, j) = \begin{cases} \frac{s_i+1}{m} & \text{if } i = j \\ \frac{1}{m} & \text{if } jTi \\ 0 & \text{if } iTj. \end{cases}$$

We claim that the distribution $\pi^{(k)}$ of \mathfrak{M} after k steps is exactly the probability distribution $\bar{p}^{(k+1)}$ over alternatives selected by the $(k+1)$ -RC. In order to see this, note that the 1-RC chooses an alternative uniformly at random. Then, the winner of the k -RC is the winner of the $(k-1)$ -RC if the latter dominates, or is identical to, the alternative assigned to the other child of the root. This happens with probability $(s_i+1)/m$ when i is the winner of the k -RC. Otherwise the winner is some other alternative that dominates the winner of the k -RC, and each such alternative is assigned to the other child of the root with probability $1/m$.

We shall be interested in the performance guarantees given by the stationary distribution π of \mathfrak{M} . We first show that \mathfrak{M} is guaranteed to converge to a unique such distribution, despite the fact that it is not necessarily irreducible.

LEMMA 4.4. *Let T be a tournament. Then $\mathfrak{M}(T)$ converges to a unique stationary distribution.*

PROOF. Let A be a set of alternatives. We first observe that any tournament $T \in \mathcal{T}(A)$ has a unique strongly connected component $tc(T) \subseteq A$, the *top cycle* of T , such that there is a directed path in T from every $i \in tc(T)$ to every $j \in A$. Clearly, a is a recurrent state of $\mathfrak{M} = \mathfrak{M}(T)$ if and only if $a \in tc(T)$. It follows that for every $\epsilon > 0$ there exists $k \in \mathbb{N}$ such that $\sum_{i \in tc(T)} \pi_i^{(k)} \geq 1 - \epsilon$. Since the subgraph of T induced by $tc(T)$ is strongly connected, and since there is a positive probability of going from any state of \mathfrak{M} to the same state in one step, the restriction of \mathfrak{M} to state space $tc(T)$ is ergodic and thus has a unique stationary distribution. Moreover, \mathfrak{M} is guaranteed to converge to this distribution as soon as it has reached a state in $tc(T)$, which in turn happens with probability tending to one as the number of steps tends to infinity. Finally, it is easily verified that the distribution which assigns probability zero to every $i \notin tc(T)$ and equals the stationary distribution of the restriction of \mathfrak{M} to $tc(T)$ for every $i \in tc(T)$ is a stationary distribution of \mathfrak{M} . \square

We are now ready to show that an alternative drawn from the stationary distribution will have an expected degree of at least half the maximum possible degree.

³Curiously, this chain bears resemblance to one previously used to define a solution concept called the Markov set (see, e.g., [20]). However, only limited attention has been given to a formal analysis of this chain, concerning properties which are different from the ones we are interested in.

LEMMA 4.5. Let $T \in \mathcal{T}(A)$ be a tournament, π the stationary distribution of $\mathfrak{M}(T)$. Then

$$\sum_{i \in A} \pi_i s_i \geq \frac{m-1}{2}.$$

To analyze π , we first require the following lemma.

LEMMA 4.6. Let T be a tournament, π the stationary distribution of $\mathfrak{M}(T)$. Then

$$\sum_{i=1}^m (2m - 2s_i - 1)\pi_i^2 = 1.$$

PROOF. Let

$$q_i = 2\pi_i \cdot \left(\sum_{j:iTj} \pi_j \right) + \pi_i^2.$$

Then

$$\sum_{i=1}^m q_i = \sum_{i \neq j} \pi_i \pi_j + \sum_{i=1}^m \pi_i^2 = \left(\sum_{i=1}^m \pi_i \right)^2 = 1.$$

On the other hand, since π is a stationary distribution,

$$\pi_i = \frac{s_i + 1}{m} \pi_i + \frac{1}{m} \sum_{j:iTj} \pi_j,$$

and thus

$$\sum_{j:iTj} \pi_j = (m - s_i - 1) \cdot \pi_i.$$

Hence, $q_i = (2m - 2s_i - 1)\pi_i^2$, which completes the proof. \square

We are now ready to prove Lemma 4.5.

PROOF OF LEMMA 4.5. For any $i \in A$, define $w_i = m - s_i - 1$. It then holds that

$$\sum_i \pi_i s_i + \sum_i \pi_i w_i = (m-1) \sum_i \pi_i = m-1. \quad (1)$$

By the Cauchy-Schwarz inequality,

$$\sum_i (2w_i + 1)\pi_i \leq \sqrt{\sum_i (2w_i + 1)} \cdot \sqrt{\sum_i (2w_i + 1)\pi_i^2}.$$

Using Lemma 4.6, $\sum_i (2w_i + 1)\pi_i^2 = 1$. Furthermore,

$$\sum_i (2w_i + 1) = 2m^2 - 2 \binom{m}{2} - m = m^2,$$

and thus,

$$\sum_i (2w_i + 1)\pi_i \leq \sqrt{m^2} \cdot \sqrt{1} = m$$

and

$$\sum_i w_i \pi_i \leq \frac{m}{2} - \frac{\sum_i \pi_i}{2} = \frac{m-1}{2}. \quad (2)$$

By combining (1) and (2) we obtain

$$\sum_i \pi_i s_i \geq \frac{m-1}{2}. \quad \square$$

The last ingredient in the proof of Lemma 4.2 and Theorem 4.1 is to show that for some k polynomial in m , the distribution over alternatives selected by the k -RC, which we recall to be equal to the distribution of \mathfrak{M} after $k-1$ steps, is close to the stationary distribution of \mathfrak{M} . In other words, we want to show that for every tournament T , $\mathfrak{M}(T)$ is rapidly mixing.⁴ The proof works by reversibilizing the transition matrix of \mathfrak{M} and then bounding the spectral gap of the reversibilized matrix via its conductance.

LEMMA 4.7. Let T be a tournament. Then, for every $\epsilon > 0$ there exists $k = k(m, \epsilon)$ polynomial in m and $1/\epsilon$, such that for all $k' > k$ and all $i \in A$, $|\pi_i^{(k')} - \pi_i| \leq \epsilon$, where $\pi^{(k)}$ is the distribution of $\mathfrak{M}(T)$ after k steps and π is the stationary distribution of $\mathfrak{M}(T)$.

PROOF. We make use of the fact that for every tournament $T \in \mathcal{T}(A)$ and every alternative $i \in A$ with maximum degree, there exists a path of length at most two from i to any other alternative. To see this, assume for contradiction that $i \in A$ has maximum degree, and that $j \in A$ is not reachable from i in two steps. Then jTi , and for all $j' \in A$, iTj' implies jTj' . Thus, $s_j > s_i$, a contradiction. This observation implies that at any given time, \mathfrak{M} either is in a state corresponding to an alternative with maximum degree, or it will reach such a state within two steps with probability at least $1/m^2$. It further implies that any alternative with maximum degree is in $tc(A)$, defined as in the proof of Lemma 4.4. We recall that once \mathfrak{M} reaches the top cycle, it stays there indefinitely. Hence, for every $\epsilon > 0$ there exists k polynomial in m and $1/\epsilon$, such that for all $k' > k$ and all $i \notin tc(T)$, $|\pi_i^{(k')} - \pi_i| = |\pi_i^{(k')}| \leq \epsilon$, where the equality follows from the fact that the support of π is contained in $tc(T)$ (see the proof of Lemma 4.4).

We further observe that π is positive on $tc(T)$, i.e., for all $i \in tc(T)$, $\pi_i > 0$. To see this, consider the largest subset of $tc(T)$ that is assigned probability zero by π , and assume that this set is nonempty. Then, for π to be a stationary distribution, no alternative in this subset can dominate an alternative in $tc(T)$ but outside the subset, contradicting the fact that $tc(T)$ is strongly connected. By all the above, we can thus focus on the restriction of \mathfrak{M} to $tc(T)$. For notational convenience, we henceforth assume w.l.o.g. that \mathfrak{M} , rather than its restriction, is irreducible and has a stationary distribution that is positive everywhere.

Conveniently, the state space Ω of \mathfrak{M} has size m , and all entries of its transition matrix P are either 0 or polynomial in m . However, there exist tournaments T such that the stationary distribution of $\mathfrak{M}(T)$ has entries that are positive but exponentially small. Furthermore, things are complicated by the fact that \mathfrak{M} is usually not reversible. We follow Fill [10] in defining the *time reversal* of P as

$$\tilde{P}(i, j) = \frac{\pi_j P(j, i)}{\pi_i},$$

and the *multiplicative reversibilization* of P as $M = M(P) = P\tilde{P}$. Then, both P and \tilde{P} are ergodic with stationary distribution π , and M is a reversible transition matrix that has

⁴We might be abusing terminology here slightly, since the theory of rapidly mixing Markov chains usually considers chains with an exponential state space, which converge in time poly-logarithmic in the size of the state space. In our case the size of the state space is only m , and the mixing rate is polynomial in m .

stationary distribution π as well. Denote by $\beta_1(M)$ the second largest eigenvalue of M . Then, by Theorem 2.7 of Fill [10],

$$4\|\pi^{(k)} - \pi\|^2 \leq (\beta_1(M))^k |\Omega|, \quad (3)$$

where $\|\sigma - \pi\| = \frac{1}{2} \sum_i |\sigma_i - \pi_i|$ is the *variation distance* between a given probability mass function σ and π . Since $|\Omega| = m$, it is sufficient to show that $\beta_1(M)$ is polynomially bounded away from 1.

To this end, we will look at the *conductance*⁵ of M , which measures the ability of M to leave any subset of the state space that has small weight under π . For a nonempty subset $S \subseteq A$, denote $\bar{S} = A \setminus S$ and $\pi_S = \sum_{i \in S} \pi_i$, and define $Q(i, j) = \pi_i M(i, j)$ and $Q(S, \bar{S}) = \sum_{i \in S, j \in \bar{S}} Q(i, j)$. The conductance of M is then given by

$$\Phi = \min_{S \subseteq A: \pi(S) \leq 1/2} \frac{Q(S, \bar{S})}{\pi_S}.$$

It is known from the work of Sinclair and Jerrum [30] that for a Markov chain reversible with respect to a stationary distribution that is positive everywhere,

$$1 - 2\Phi \leq \beta_1(A) \leq 1 - \frac{\Phi^2}{2}.$$

It thus suffices to bound Φ polynomially away from 0. For any S with $\pi_S \leq 1/2$ it holds that

$$\frac{Q(S, \bar{S})}{\pi_S} \geq \frac{Q(S, \bar{S})}{2\pi_S\pi_{\bar{S}}} = \frac{\sum_{i \in S, j \in \bar{S}} Q(i, j)}{2 \sum_{i \in S, j \in \bar{S}} \pi_i \pi_j} \geq \min_{i \in S, j \in \bar{S}} \frac{Q(i, j)}{2\pi_i \pi_j}.$$

In our case,

$$\begin{aligned} Q(i, j) &= \pi_i \left[\sum_{r \in A} P(i, r) \tilde{P}(r, j) \right] \\ &\geq \pi_i [P(i, i) \tilde{P}(i, j) + P(i, j) \tilde{P}(j, j)] \\ &\geq \frac{1}{m} [\pi_i P(i, j) + \pi_j P(j, i)]. \end{aligned} \quad (4)$$

A crucial observation is that for every $i \neq j$, either $P(i, j) = 1/m$ or $P(j, i) = 1/m$, since either iTj or jTi . Now, let $i_0 \in S$ and $j_0 \in \bar{S}$ be the two alternatives for which the minimum above is attained. If $P(i_0, j_0) = 1/m$, then by (4),

$$\frac{Q(i_0, j_0)}{2\pi_{i_0}\pi_{j_0}} \geq \frac{\frac{\pi_{i_0}}{m^2}}{2\pi_{i_0}\pi_{j_0}} = \frac{1}{2m^2\pi_{j_0}},$$

whereas if $P(j_0, i_0) = 1/m$, then

$$\frac{Q(i_0, j_0)}{2\pi_{i_0}\pi_{j_0}} \geq \frac{1}{2m^2\pi_{i_0}}.$$

In both cases, $\Phi \geq 1/(2m^2)$, which completes the proof. \square

We now have all the necessary ingredients in place.

PROOF OF LEMMA 4.2 AND THEOREM 4.1. Let $\epsilon > 0$. By Lemma 4.3 and Lemma 4.7, there exists k polynomial in m and $1/\epsilon$ such that for all $i \in A$, $|p_i^{(k)} - \tilde{p}_i^{(k)}| \leq \epsilon/(2\binom{m}{2})$ and $|\tilde{p}_i^{(k)} - \pi_i| \leq \epsilon/(2\binom{m}{2})$. By the triangle inequality, $|p_i^{(k)} - \pi_i| \leq \epsilon/\binom{m}{2}$. Now,

$$\sum_i \pi_i s_i - \sum_i p_i^{(k)} s_i \leq \sum_i |\pi_i - p_i^{(k)}| s_i \leq \frac{\epsilon}{\binom{m}{2}} \sum_i s_i = \epsilon.$$

Lemma 4.2 and thus Theorem 4.1 now follow directly by Lemma 4.5. \square

⁵The conductance is called *Cheeger constant* by Fill [10].

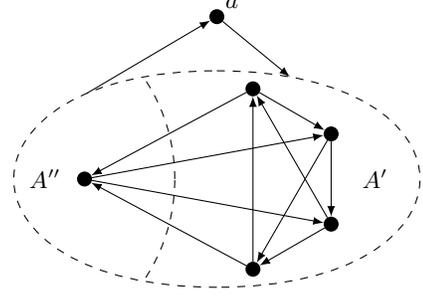


Figure 2: Tournament structure providing an upper bound for the randomized k -caterpillar, example for $m = 6$ and $\epsilon = 1/5$. A' and A'' contain $(1 - \epsilon)(m - 1)$ and $\epsilon(m - 1)$ alternatives, respectively.

4.2 Tightness and Stability of the Caterpillar

It turns out that the analysis in the proof of Theorem 4.1 is tight. Indeed, since we have seen that the stationary distribution π of \mathfrak{M} is very close to the distribution of alternatives chosen by the k -RSC, it is sufficient to see that π cannot guarantee an approximation ratio better than $1/2$ in expectation. Consider a set A of alternatives, and a partition of A into three sets A' , A'' , and $\{a\}$ such that $|A'| = (1 - \epsilon)(m - 1)$ and $|A''| = \epsilon(m - 1)$ for some $\epsilon > 0$. Consider further a tournament $T \in \mathcal{T}(A)$ where a dominates every alternative in A' and is itself dominated by every alternative in A'' , and such that the subgraph induced by $A' \cup A''$ is regular, *i.e.*, has each alternative dominate exactly $(|A'| + |A''| - 1)/2$ other alternatives. The structure of T is illustrated in Figure 2. It is easily verified that the stationary distribution π of $\mathfrak{M}(T)$ satisfies

$$\pi_a = \frac{\sum_{j: aTj} \pi_j}{m - s_a - 1} \leq \frac{1}{m - s_a - 1} \leq \frac{1}{\epsilon(m - 1)},$$

and therefore,

$$\begin{aligned} \sum_i \pi_i s_i &\leq \frac{1}{\epsilon(m - 1)} (m - 1) + \frac{\epsilon(m - 1) - 1}{\epsilon(m - 1)} \cdot \left(\frac{m - 1}{2} + 1 \right) \\ &\leq \frac{m - 1}{2} + \frac{1}{\epsilon} + 1. \end{aligned}$$

Furthermore, a has degree $(1 - \epsilon)(m - 1)$. If we choose, say, $\epsilon = 1/\sqrt{m}$, then the approximation ratio tends to $1/2$ as m tends to infinity.

We proceed to demonstrate that the above tournament is a generic bad example. Indeed, Lemma 4.5 will be shown to possess the following stability property: in every tournament where π achieves an approximation ratio only slightly better than $1/2$, almost all alternatives have degree close to $m/2$, as it is the case for the example above. In particular, this implies that \mathfrak{M} either provides an expected approximation ratio better than $1/2$, or selects an alternative with score around $m/2$ with very high probability.

THEOREM 4.8. *Let $\epsilon > 0$, $m \geq 1/(2\sqrt{\epsilon})$. Let T be a tournament over a set of m alternatives, π the stationary distribution of $\mathfrak{M}(T)$. If $\sum_i \pi_i s_i = (m - 1)/2 + \epsilon m$, then*

$$\left\{ i \in A : \left| s_i - \frac{m}{2} \right| > \frac{3\sqrt[4]{4\epsilon}}{2} m \right\} \leq \sqrt[4]{4\epsilon} \cdot m.$$

We shall require two lemmata: a “geometric” version of the Cauchy-Schwarz inequality and a well-known result about the sequence of degrees of a tournament. The proof of the former is given in the full version of the paper, the latter we state without proof.

LEMMA 4.9. Consider two vectors $\mathbf{a} = (a_1, \dots, a_m) \in \mathbb{R}^m$ and $\mathbf{b} = (b_1, \dots, b_m) \in \mathbb{R}^m$. Then,

$$\sum_{i=1}^m \left(\frac{a_i}{\|a\|} - \frac{b_i}{\|b\|} \right)^2 = \epsilon \iff \sum_{i=1}^m a_i b_i = (1 - \frac{\epsilon}{2}) \|a\| \cdot \|b\|.$$

LEMMA 4.10 (Moon [25], Theorem 29). A sequence $s_1 \leq \dots \leq s_m$ is the degree sequence of a tournament if and only if for all $k \leq m$, $\sum_{i=1}^k s_i \geq \binom{k}{2}$. \square

PROOF OF THEOREM 4.8. Define $w_i = m - s_i - 1$, $a_i = \sqrt{2w_i + 1}$, and $b_i = \sqrt{2w_i + 1} \pi_i$. By the assumption that $\sum_i \pi_i s_i = \frac{m-1}{2} + \epsilon m$ and by (1) in the proof of Lemma 4.5, we have that $\sum_i a_i b_i = (1 - 2\epsilon)m$. Since $\|a\| = m$ and, by Lemma 4.6, $\|b\| = 1$, we have

$$\sum_i a_i b_i = (1 - 2\epsilon) \|a\| \cdot \|b\|.$$

By Lemma 4.9,

$$\sum_i \left(\frac{a_i}{\|a\|} - \frac{b_i}{\|b\|} \right)^2 = 4\epsilon.$$

Denoting $\epsilon' = 4\epsilon$,

$$\sum_i \left(\frac{\sqrt{2w_i + 1}}{m} - \sqrt{2w_i + 1} \cdot \pi_i \right)^2 = \epsilon'.$$

By simplifying and rearranging, we get

$$\sum_i (2w_i + 1) \left(\pi_i - \frac{1}{m} \right)^2 = \epsilon'. \quad (5)$$

Now let $\epsilon'' = \sqrt[4]{\epsilon'}$, and

$$B = \left\{ i \in A : \left| \pi_i - \frac{1}{m} \right| > \frac{\epsilon''}{m} \right\}.$$

We claim that $|B| \leq \epsilon'' m$. Assume for contradiction that $|B| > \epsilon'' m$. Then, by Lemma 4.10,

$$\sum_{i \in B} s_i = \binom{m}{2} - \sum_{i \notin B} s_i \leq \binom{m}{2} - \binom{m - |B|}{2},$$

and

$$\sum_{i \in B} w_i \geq |B|(m - 1) - \binom{m}{2} + \binom{m - |B|}{2} = \binom{|B|}{2}.$$

We thus have

$$\begin{aligned} \sum_{i \in B} (2w_i + 1) \left(\pi_i - \frac{1}{m} \right)^2 &> \frac{\sqrt{\epsilon'}}{m^2} \sum_{i \in B} (2w_i + 1) \\ &\geq \frac{\sqrt{\epsilon'}}{m^2} \left(2 \frac{|B|(|B| - 1)}{2} + |B| \right) \\ &> \frac{\sqrt{\epsilon'}}{m^2} \cdot \sqrt{\epsilon'} m^2 = \epsilon', \end{aligned}$$

contradicting (5). The first inequality holds because $|\pi_i - 1/m| > \epsilon''/m$ for all $i \in B$, the last one follows from the assumption that $|B| > \epsilon'' m$.

It now suffices to show that for all $i \notin B$, $|s_i - \frac{m}{2}| \leq (3\epsilon''/2)m$, i.e., that B contains all alternatives with degree significantly different from $m/2$. Let $i \in A \setminus B$. Since π is a stationary distribution,

$$(m - s_i - 1)\pi_i = \sum_{j:iTj} \pi_j.$$

At most $\epsilon'' m$ of the alternatives dominated by i can be in B , and thus

$$m - s_i - 1 \geq \frac{(s_i - \epsilon'' m) \left(\frac{1}{m} - \frac{\epsilon''}{m} \right)}{\frac{1}{m} + \frac{\epsilon''}{m}}.$$

It should be noted that this holds even if $s_i - \epsilon'' m < 0$. By rearranging and simplifying,

$$(m - s_i - 1)(1 + \epsilon'') \geq (1 - \epsilon'')s_i - m\epsilon''(1 - \epsilon''),$$

and thus

$$s_i \leq \frac{m}{2} + \epsilon'' m.$$

On the other hand,

$$\sum_{j \notin B} \pi_j \geq (1 - \epsilon'')m \cdot \frac{1 - \epsilon''}{m},$$

and therefore

$$(m - s_i - 1) \leq \frac{s_i \frac{1 + \epsilon''}{m} + \left(1 - (1 - \epsilon'')m \frac{1 - \epsilon''}{m} \right)}{\frac{1 - \epsilon''}{m}}.$$

The last implication is true because i dominates at most s_i alternatives outside B , and the overall probability assigned to alternatives in B is at most $1 - (1 - \epsilon'')m \frac{1 - \epsilon''}{m}$. Now,

$$(m - s_i - 1)(1 - \epsilon'') \leq s_i(1 + \epsilon'') + m(2\epsilon'' - (\epsilon'')^2).$$

Thus, for $m \geq \frac{1}{(\epsilon'')^2}$,

$$s_i \geq \frac{m}{2} - \frac{3}{2}\epsilon'' m. \quad \square$$

4.3 Second Order Degrees

So far we have been concerned with the Copeland solution, which selects an alternative with maximum degree. Recently, a related solution concept, sometimes referred to as *second order Copeland*, has received attention in the social choice literature (see, e.g., [2]). Given a tournament T , this solution breaks ties with respect to the maximum degree toward alternatives i with maximum *second order degree* $\sum_{j:iTj} s_j$. Second order Copeland is the first rule, and one of only two natural voting rules, known to be computationally easy to compute but difficult to manipulate [2].

Interestingly, the same randomization studied in Section 4.1 also achieves a 1/2-approximation for the second order degree.

THEOREM 4.11. Let A be a set of alternatives, $T \in \mathcal{T}(A)$. For $k \in \mathbb{N}$, let $p_i^{(k)}$ denote the probability that alternative $i \in A$ is selected by the k -RSC for T . Then, there exists $k = k(m)$ polynomial in m such that

$$\frac{\sum p_i^{(k)} \sum_{j:iTj} s_j}{\max_{i \in A} \sum_{j:iTj} s_j} \geq \frac{1}{2} + \Omega(1/m).$$

Clearly, the sum of degrees of alternatives dominated by an alternative i is at most $\binom{m-1}{2}$. The lower bound is then obtained from an explicit result about the second order degree of alternatives chosen by the k -RSC. Along similar lines as in the proof of Theorem 4.1, it suffices to prove that the stationary distribution of $\mathfrak{M}(T)$ provides an approximation. The following lemma is the second order analog of Lemma 4.5. It turns out, however, that the technique used in the proof of Lemma 4.5, namely directly manipulating the stationary distribution equations and applying Cauchy-Schwarz, does not work for the second order degree. We instead formulate a suitable LP and bound the primal by a feasible solution to the dual.

LEMMA 4.12. *Let T be a tournament, π the stationary distribution of $\mathfrak{M}(T)$. Then,*

$$\sum_{i \in A} \left(\pi_i \sum_{j: iTj} s_j \right) \geq \frac{m^2}{4} - \frac{m}{2}.$$

PROOF. Fix some tournament $T \in \mathcal{T}(A)$, and consider the degrees s_i in T . The minimum expected second order degree of an alternative drawn according to the stationary distribution of $\mathfrak{M}(T)$ is given by the following linear program with variables π_i :

$$\begin{aligned} \mathbf{min} \quad & \sum_{i \in A} \pi_i \left(\sum_{j: iTj} s_j \right) \\ \mathbf{s.t.} \quad & \forall i, (m - s_i - 1)\pi_i - \sum_{j: iTj} \pi_j = 0, \\ & \sum_{i \in A} \pi_i = 1, \\ & \forall i, \pi_i \geq 0. \end{aligned}$$

The dual is the following program with variables x_i and y :

$$\begin{aligned} \mathbf{max} \quad & y \\ \mathbf{s.t.} \quad & \forall i, (m - s_i - 1)x_i - \sum_{j: jTi} x_j + \sum_{j: iTj} s_j \geq y. \end{aligned}$$

By weak duality, any feasible solution to the dual provides a lower bound on the optimal assignment to the primal. Consider the assignment $x_i = -s_i$ to the dual. The maximum feasible value of y given this assignment is the minimum over the left hand side of the constraints. We claim that for any i , the value of the left hand side is at least $m^2/4 - m/2$. Indeed, for all i ,

$$\begin{aligned} & (m - s_i - 1)(-s_i) - \sum_{j: jTi} (-s_j) + \sum_{j: iTj} s_j \\ &= (m - s_i - 1)(-s_i) + \sum_{j \neq i} s_j \\ &= (m - s_i - 1)(-s_i) + \left(\binom{m}{2} - s_i \right) \\ &= m^2/2 - m/2 - s_i(m - s_i) \geq m^2/4 - m/2. \quad \square \end{aligned}$$

We point out that the analysis is tight. Indeed, the second order degree of any alternative in a regular tournament, *i.e.*, one where each alternative dominates exactly $(m-1)/2$ other alternatives, is $(m-1)/2 \cdot (m-1)/2 = m^2/4 - m/2 + 1/4$. Theorem 4.11 itself is also tight, by the example given in Section 4.2.

5. BALANCED TREES

In the previous section we presented our main positive results, all of which were obtained using randomizations over caterpillars. Since caterpillars are maximally unbalanced, one would hope to do much better by looking at *balanced trees*, *i.e.*, trees where the depth of any two leaves differs by at most one. We briefly explore this intuition. Consider a balanced binary tree where each alternative in a set A appears exactly once at a leaf. We will call such a tree a *permutation tree* on A . As we have already mentioned in the previous section, permutation trees provide a very weak deterministic lower bound. Indeed, the winning alternative must dominate the $\Theta(\log m)$ alternatives it meets on the path to the root, all of which are distinct. Since there always exists an alternative with score at least $(m-1)/2$, we obtain an approximation ratio of $\Theta((\log m)/m)$. On the other hand, no voting tree in which every two leaves have distinct labels can guarantee to choose an alternative with degree larger than the height of the tree, so the above bound is tight. More interestingly, it can be shown that no composition of permutation trees, *i.e.*, no tree obtained by replacing every leaf of an arbitrary binary tree by a permutation tree, can provide a lower bound better than $1/2$. Unfortunately, larger balanced trees not built from permutation trees have so far remained elusive.

Can we obtain a better bound by randomizing? Intuitively, a randomization over large balanced trees should work well, because one would expect that the winning alternative dominates a large number of randomly chosen alternatives on the way to the root. Surprisingly, the complete opposite is the case. In the following, we call *randomized perfect voting tree* of height k , or k -RPT, a voting tree where every leaf is at depth k and labels are assigned uniformly at random. This tree obviously corresponds to a randomization that is not admissible, but a similar result for admissible randomizations can easily be obtained by using the same arguments as before.

THEOREM 5.1. *Let A be a set of alternatives, $|A| \geq 5$. For every $K \in \mathbb{N}$ there exists $K' \geq K$ such that the K' -RPT provides an approximation ratio of at most $\mathcal{O}(1/m)$.*

The proof of this theorem, which is given in the full version of the paper, works by constructing a tournament and showing that as k grows, the distribution over alternatives chosen by the k -RPT *oscillates* between different parts of this tournament in which alternatives have large and small degree, respectively. In the full version we further analyze higher order voting caterpillars obtained by replacing each leaf of a caterpillar of sufficiently large height by higher order caterpillars of smaller order (in particular, of order reduced by one). As in the case of the k -RPT, this construction does not provide better bounds but instead causes the approximation ratio to deteriorate.

6. OPEN PROBLEMS

Many interesting questions arise from our work. Perhaps the most enigmatic open problem in the context of this paper concerns tighter bounds for deterministic trees. Some results for restricted classes of trees have been discussed in Section 5, but in general there remains a large gap between the upper bound of $3/4$ derived in Section 3 and the straightforward lower bound of $\Theta((\log m)/m)$. In the randomized

model our situation is somewhat better. Nevertheless, an intriguing gap remains between our upper bound of $5/6$, which holds even for inadmissible randomizations over arbitrarily large trees, and the lower bound of $1/2$ obtained from an admissible randomization over trees of polynomial size. It might be the case that the height of a k -RPT could be chosen carefully to obtain some kind of approximation guarantee. For example, one could investigate the uniform distribution over permutation trees. The analysis of this type of randomization is closely related to the theory of dynamical systems, and we expect it to be rather involved.

More generally, the problem of characterizing voting rules implementable by trees remains open. While we have seen that techniques from theoretical computer science can contribute to the solution of this problem, they are perhaps even more promising for attacking questions concerning the implementability of probabilistic rules by randomized trees or upper and lower bounds on the size of trees implementing certain rules.

Acknowledgements

We thank Julia Böttcher, Felix Brandt, Shahar Dobzinski, Dvir Falik, Paul Harrenstein, Jeff Rosenschein, and Michael Zuckerman for many helpful discussions.

References

- [1] J. S. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 1:295–306, 1985.
- [2] J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.
- [3] F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. In *Proc. of 11th TARK Conference*, pages 82–91, 2007.
- [4] R. Chen and F. K. Hwang. Stronger players win more balanced knockout tournaments. *Graphs and Combinatorics*, 4(1):95–99, 1988.
- [5] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.
- [6] P. J. Coughlan and M. L. Breton. A social choice function implementable via backward induction with values in the ultimate uncovered set. *Review of Economic Design*, 4:153–160, 1999.
- [7] M. de Condorcet. *Essai sur l'application de l'analyse à la probabilité de décisions rendues à la pluralité de voix*. Imprimerie Royal, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.
- [8] B. Dutta and A. Sen. Implementing generalized Condorcet social choice functions via backward induction. *Social Choice and Welfare*, 10:149–160, 1993.
- [9] R. Farquharson. *Theory of Voting*. Yale University Press, 1969.
- [10] J. A. Fill. Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *The Annals of Applied Probability*, 1(1):62–87, 1991.
- [11] P. C. Fishburn. Even-chance lotteries in social choice theory. *Theory and Decision*, 3(1):18–40, 1972.
- [12] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proc. of 49th FOCS Symposium*, pages 243–249, 2008.
- [13] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
- [14] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- [15] M. Herrero and S. Srivastava. Implementation via backward induction. *Journal of Economic Theory*, 56:70–88, 1992.
- [16] M. D. Intriligator. A probabilistic model of social choice. *Review of Economic Studies*, 40(4):553–560, 1973.
- [17] J. Kahn, M. Saks, and D. Sturtevant. A topological approach to evasiveness. *Combinatorica*, 4:297–306, 1984.
- [18] G. Laffond, J. F. Laslier, and M. L. Breton. The Copeland measure of Condorcet choice functions. *Discrete Applied Mathematics*, 55:273–279, 1994.
- [19] J. Lang, M.-S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Winner determination in sequential majority voting. In *Proc. of 20th IJCAI*, pages 1372–1377, 2007.
- [20] J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer, 1997.
- [21] K. May. A set of independent, necessary and sufficient conditions for simple majority decisions. *Econometrica*, 20:680–684, 1952.
- [22] D. C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.
- [23] R. D. McKelvey and R. G. Niemi. A multistage game representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, 18:1–22, 1978.
- [24] N. Miller. A new solution set for tournaments and majority voting: Further graph theoretical approaches to the theory of voting. *American Journal of Political Science*, 24:68–96, 1980.
- [25] J. W. Moon. *Topics on Tournaments*. Holt, Reinhart and Winston, 1968.
- [26] H. Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3:271–291, 1986.
- [27] A. D. Procaccia, A. Zohar, Y. Peleg, and J. S. Rosenschein. Learning voting trees. In *Proc. of 22nd AAAI Conference*, pages 110–115, 2007.
- [28] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [29] K. A. Shepsle and B. R. Weingast. Uncovered sets and sophisticated outcomes with implications for agenda institutions. *American Journal of Political Science*, 28(1):49–74, 1984.
- [30] A. Sinclair and M. Jerrum. Approximate counting, uniform generation, and rapidly mixing Markov chains. *Information and Computation*, 82:93–133, 1989.
- [31] S. Srivastava and M. A. Trick. Sophisticated voting rules: The case of two tournaments. *Social Choice and Welfare*, 13:275–289, 1996.
- [32] A. C. Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proc. of 17th FOCS Symposium*, pages 222–227, 1977.
- [33] R. Zeckhauser. Majority rule with lotteries on alternatives. *Quarterly Journal of Economics*, 83(4):696–703, 1969.
- [34] M. Zuckerman, A. D. Procaccia, and J. S. Rosenschein. Algorithms for the coalitional manipulation problem. In *Proc. of 19th SODA*, pages 277–286, 2008.