On Voting Caterpillars: Approximating Maximum Degree in a Tournament by Binary Trees

Felix Fischer* Ariel D. Procaccia† Alex Samorodnitsky‡

Abstract

Voting trees describe an iterative procedure for selecting a single vertex from a tournament. It has long been known that there is no voting tree that always singles out a vertex with maximum degree. In this paper, we study the power of voting trees in *approximating* the maximum degree. We 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.

Keywords: Computational social choice, Algorithmic mechanism design, Approximation, Markov chains

*Institut für Informatik, Ludwig-Maximilians-Universität München, 80538 München, Germany, email: `fischerf@tcs.ifi.lmu.de`. 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.

†School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, email: `arielpro@cs.huji.ac.il`. The author is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

‡School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, email: `salex@cs.huji.ac.il`. The author is supported by ISF grant 039-7165.

1 Introduction

A problem that pervades the theory of social choice is the selection of "best" alternatives from a *tournament*, *i.e.*, a complete and asymmetric (dominance) relation over a set of alternatives (see, *e.g.*, Laslier, 1997). Such a relation for example arises from pairwise majority voting with an odd number of voters and linear preferences. In graph theoretic terms, a tournament is an orientation of a complete undirected graph, with a directed edge from a dominating alternative to a dominated one. In the presence of cycles the concept of maximality is not well-defined, and so-called tournament solutions have been devised to take over the role of singling out good alternatives. A prominent such solution, known as the Copeland solution, selects the alternatives with *maximum (out-)degree*, *i.e.*, those that beat the largest number of other alternatives in a direct comparison.

An interesting question concerns the implementation of a solution concept using a specific procedure. We shall specifically be interested in the well-known class of procedures given by *voting trees.* 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. It has long been known that there exists no voting tree implementing the Copeland solution, *i.e.*, one that always selects a vertex with maximum degree (Moulin, 1986). In this paper, we ask a natural question from a computer science point of view: "Is there a voting tree that *approximates* the maximum degree?" More precisely, we would like to determine 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 will 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.

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. We 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 further 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. 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 further show 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 (1969), and further explored, for example, by McKelvey and Niemi (1978), Miller (1980), Moulin (1986), Herrero and Srivastava (1992), Dutta and Sen (1993), Srivastava and Trick (1996), and Coughlan and Breton (1999). In particular, Moulin (1986) shows that the Copeland solution is not implementable by voting trees if there are at least 8 alternatives, while Srivastava and Trick (1996) demonstrate that it can be implemented for tournaments with up to 7 alternatives.

Laffond et al. (1994) 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.

In theoretical computer science, the problem studied in this paper is somewhat reminiscent of the problem of determining query complexity of graph properties (see, *e.g.*, Rosenberg, 1973; Rivest and Vuillemin, 1976; Kahn et al., 1984; King, 1988). 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, which lies 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. (2007) characterize the computational complexity of determining different types of winners in voting trees. Procaccia et al. (2007) investigate the learnability of voting trees, as functions from tournaments to alternatives. In a slightly different context, Brandt et al. (2007) 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 questions. Detailed proofs of all results, as well as an analysis of "higher order" caterpillars, are given in the appendix.

2 Preliminaries

Let $A = \{1, \ldots, m\}$ 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 $\ell : V \to A$ is a mapping that assigns an element of A to each leaf of (V, E). Given a tournament T, a unique function $\ell_T : V \to A$ exists such that

$$\ell_T(v) = \begin{cases} \ell(v) & \text{if } v \text{ is a leaf} \\ \ell(u_1) & \text{if } v \text{ has children } u_1 \text{ and } u_2\text{, and } \ell(u_1)T\ell(u_2) \text{ or } \ell(u_1) = \ell(u_2) \end{cases}$$

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.*, Moulin, 1986). 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 derive 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 (1986), 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*[1] 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

[1]Moulin (1986) uses the term "adjacent set".

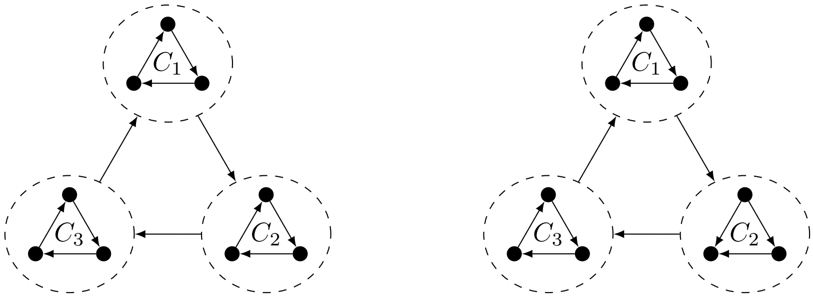

Figure 1: Tournaments used in the proof of Theorem 3.2, illustrated for $k = 3$. A voting tree is assumed to select an alternative from C_1.

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. Furthermore, if an alternative outside the component is chosen for one tournament, then the same alternative has to be chosen for the other. Laslier (1997) calls a solution concept satisfying these properties *weakly composition-consistent.*

Lemma 3.1 (Moulin 1986)**.** *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 $[\Gamma(T) \in C$ if and only if $\Gamma(T') \in C]$, and*
2. *$[T_C = T'_C$ and $\Gamma(T) \in A \setminus C]$ implies $[\Gamma(T) = \Gamma(T')]$.*

We are now ready to strengthen the negative result concerning implementability of the Copeland solution (Moulin, 1986) by showing that no deterministic tree can always choose an alternative that has a degree significantly larger than 3/4 of the maximum degree.

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 $|A| = m = 3k + 1$ 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 restriction of T to 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 restrictions of T and T' to $B \subseteq A$ are identical if $|B \cap C_2| \leq 1$, and the restriction of T' to 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

$$s_{\Gamma(T')} = k + \frac{(k-1)}{2} = \frac{3k}{2} - \frac{1}{2} \qquad \text{and} \qquad \max_{i \in A} s_i = 2k - 1,$$

and thus

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

□

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 (1977).

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, $\alpha \leq 5/6 + \mathcal{O}(1/m)$.*

The proof of this theorem is given in Appendix A. 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. This clearly facilitates its use as a computational procedure. 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.

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

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

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:[2] 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

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

[2]Curiously, this chain bears resemblance to one previously used to define a solution concept called the Markov set (see, *e.g.*, Laslier, 1997). 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.

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.*

The proof of the lemma appears in Appendix B. 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.

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

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

The proof is based on some algebraic manipulations and the Cauchy-Schwarz inequality, and is given in Appendix C.

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.[3]

Lemma 4.6. *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 $\mathfrak{M}(T)$.*

The proof of Lemma 4.6, given in Appendix D, works by reversibilizing the transition matrix of $\mathfrak{M}$ and then bounding the spectral gap of the reversibilized matrix via its conductance.

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 4.6, there exists k polynomial in m and $1/\epsilon$ such that for all $i \in A$, $|p_i^{(k)} - \bar{p}_i^{(k)}| \leq \epsilon/(2\binom{m}{2})$ and $|\bar{p}_i^{(k)} - \pi_i| \leq \epsilon/(2\binom{m}{2})$. 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 follow directly by Lemma 4.5. □

[3] We might be slightly abusing terminology here, 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.

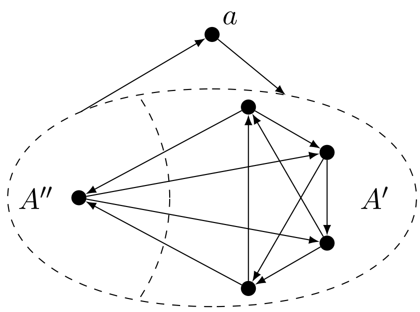

Figure 2: Tournament structure providing an upper bound for the randomized k-caterpillar, example for $m = 6$ and $\epsilon = 1/5$. A' and A'' contain $(1-\epsilon)(m-1)$ and $\epsilon(m-1)$ alternatives, respectively.

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 $\epsilon > 0$. Further consider a tournament $T \in \mathcal{T}(A)$ in which a dominates every alternative in A' and is itself dominated by every alternative in A'', and for which the restriction of T to $A' \cup A''$ is regular. 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,

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

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 4.7. *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*

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

The details of the proof appear in Appendix E.

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.*, Bartholdi et al., 1989). 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 (Bartholdi et al., 1989).

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

Theorem 4.8. *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.

Lemma 4.9. *Let T be a tournament, π the stationary distribution of $\mathfrak{M}(T)$. Then,*

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

It turns out 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. The proof of the lemma, which in turn implies Theorem 4.8, is given in Appendix F.

We further 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 4.8 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, permutations 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 dominate 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 $K \in \mathbb{N}$ and $\epsilon > 0$, 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, given in Appendix G, constructs a tournament consisting of a 3-cycle of components and shows that the distribution over alternatives chosen by the k-RPT *oscillates* between the different components as k grows.

In Appendix H we 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.

Acknowledgements

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

References

J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.

F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. In *Proceedings of the Eleventh Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 82–91, 2007.

P. J. Coughlan and M. Le Breton. A social choice function implementable via backward induction with values in the ultimate uncovered set. *Review of Economic Design*, 4:153–160, 1999.

B. Dutta and A. Sen. Implementing generalized Condorcet social choice functions via backward induction. *Social Choice and Welfare*, 10:149–160, 1993.

R. Farquharson. *Theory of Voting*. Yale University Press, 1969.

J. A. Fill. Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *The Annals of Applied Probablity*, 1(1):62–87, 1991.

M. Herrero and S. Srivastava. Decentralization by multistage voting procedures. *Journal of Economic Theory*, 56:182–201, 1992.

J. Kahn, M. Saks, and D. Sturtevant. A topological approach to evasiveness. *Combinatorica*, 4: 297–306, 1984.

V. King. Lower bounds on the complexity of graph properties. In *Proceedings of the 20th ACM Symposium on Theory of Computing (STOC)*, pages 468–476, 1988.

G. Laffond, J. F. Laslier, and M. Le Breton. The Copeland measure of Condorcet choice functions. *Discrete Applied Mathematics*, 55:273–279, 1994.

J. Lang, M.-S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Winner determination in sequential majority voting. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1372–1377, 2007.

J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer, 1997.

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.

N. Miller. A new solution set for tournaments and majority voting: Further graph theoretical approaches to the theory of voting. *Americal Journal of Political Science*, 24:68–96, 1980.

J. W. Moon. *Topics on Tournaments*. Holt, Reinhart and Winston, 1968.

H. Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3:271–291, 1986.

A. D. Procaccia, A. Zohar, Y. Peleg, and J. S. Rosenschein. Learning voting trees. In *Proceedings of the 22nd Conference on Artificial Intelligence (AAAI)*, pages 110–115, 2007.

R. Rivest and S. Vuillemin. On recognizing graph properties from adjacency matrices. *Theoretical Computer Science*, 3:371–384, 1976.

A. L. Rosenberg. The time required to recognize properties of graphs: A problem. *SIGACT News*, 5(4):15–16, 1973.

A. Sinclair and M. Jerrum. Approximate counting, uniform generation, and rapidly mixing Markov chains. *Information and Computation*, 82:93–133, 1989.

S. Srivastava and M. A. Trick. Sophisticated voting rules: The case of two tournaments. *Social Choice and Welfare*, 13:275–289, 1996.

A. C. Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proceedings of the 17th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 222–227, 1977.

A Proof of Theorem 3.3

Proof. Reformulating the minimax principle for voting trees, an upper bound on the worst-case performance of the best randomized tree on a set A of alternatives is given by the performance of the best deterministic tree with respect to some probability distribution over tournaments on A.

As in the proof of Theorem 3.2, we assume for ease of exposition that $|A| = m = 3k + 1$ for some odd k, and define a tournament T as a cycle of three regular components C_1, C_2, and C_3, each of size k. Further define three new tournaments T_1, T_2, and T_3 such that for $r = 1, 2, 3$, the restrictions of T and T_r to $B \subseteq A$ are identical if $|B \cap C_r| \leq 1$, and the restriction of T_r to C_r is transitive. Let Γ be any deterministic tree on A. Combining both statements of Lemma 3.1, there exists $i \in \{1, 2, 3\}$ such that for $r = 1, 2, 3$, $\Gamma(T_r) \in C_i$. In particular, Γ selects an alternative with score at most $3k/2 - 1/2$ for two of the three tournaments T_r. Now consider a tournament T drawn uniformly from $\{T_1, T_2, T_3\}$. By the above,

$$\mathbb{E}_{\Gamma\sim\Delta}[s_{\Gamma(T)}] \leq (2(3k/2 - 1/2) + (2k - 1))/3 = 5k/3 - 2/3 \qquad \text{and} \qquad \max_{i\in A} s_i = 2k - 1,$$

and thus

$$\frac{\mathbb{E}_{\Gamma\sim\Delta}[s_{\Gamma(T)}]}{\max_{i\in A} s_i} \leq \frac{5k-2}{6k-3} \leq \frac{5(k-1)+3}{6(k-1)} = \frac{5}{6} + \frac{1}{2(k-1)}.$$

In particular, this ratio tends to 5/6 as k tends to infinity. □

B Proof of Lemma 4.4

Proof (sketch). 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 restriction of T to $tc(T)$ 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 $tc(T)$ 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 $\mathfrak{M}$. □

C Proof of Lemma 4.5

To analyze π, we require the following lemma.

Lemma C.1. *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. □

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

$$\sum_i \pi_i s_i + \sum_i \pi_i w_i = (m-1)\sum_i \pi_i = m-1. \tag{1}$$

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 C.1, $\sum_i (2w_i+1)\pi_i^2 = 1$. Furthermore,

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

and thus,

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

and

$$\sum_i w_i\pi_i \le \frac{m}{2} - \frac{\sum_i \pi_i}{2} = \frac{m-1}{2}. \tag{2}$$

By combining (1) and (2) we obtain

$$\sum_i \pi_i s_i \ge \frac{m-1}{2}. \qquad \square$$

D Proof of Lemma 4.6

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 $i \notin tc(T)$, $|\pi_i^{(k')} - \pi_i| = |\pi_i^{(k')}| \le \epsilon$, 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 $i \in tc(T)$, $\pi_i > 0$. Too 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 (1991) 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 stationary distribution π as well. Denote by $\beta_1(M)$ the second largest eigenvalue of M. Then, by Theorem 2.7 of Fill (1991),

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

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*[4] 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

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

It is known from the work of Sinclair and Jerrum (1989) 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,

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

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. □

E Proof of Theorem 4.7

We shall require two lemmata. The first one is a "geometric" version of the Cauchy-Schwarz inequality. The second one is a well-known result about the sequence of degrees of a tournament, which we state without proof.

Lemma E.1. *Let* $\mathbf{a} = (a_1, \ldots, a_m) \in \mathbb{R}^m$, $\mathbf{b} = (b_1, \ldots, b_m) \in \mathbb{R}^m$. *Then,*

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

[4]The conductance is called *Cheeger constant* by Fill (1991).

Proof.

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

Lemma E.2 (Moon, 1968, Theorem 29)**.** *$s_1 \le s_2 \le \cdots \le s_m$ is the degree sequence of a tournament if and only if for all $k \le m$, $\sum_{i=1}^{k} s_i \ge \binom{k}{2}$.* □

Proof of Theorem 4.7. 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 $\sum_i a_ib_i = (1-2\epsilon)m$. Since $\|a\| = m$ and, by Lemma C.1, $\|b\| = 1$, we have

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

By Lemma E.1,

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

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

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

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 $|B| \le \epsilon'' m$. Assume for contradiction that $|B| > \epsilon'' m$. Then, by Lemma E.2,

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

and

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

We thus have

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

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 $i \notin B$, $\left|s_i - \frac{m}{2}\right| \leq (3\epsilon''/2)m$, *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 $m \geq \frac{1}{(\epsilon'')^2}$,

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

□

F Proof of Lemma 4.9

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:

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

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

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

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,

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

□

G Proof of Theorem 5.1

To prove the theorem, we will show that given a tournament consisting of a 3-cycle of components, the distribution over alternatives chosen by the k-RPT *oscillates* between the different components as k grows. This is made precise in the following lemma.

Lemma G.1. *Let A be a set of alternatives, $T \in \mathcal{T}(A)$ containing three components C_i, $i = 1, 2, 3$, such that for all alternatives $a \in C_i$ and $b \in C_{(i \bmod 3)+1}$, aTb. For $i = 1, 2, 3$ and $k \in \mathbb{N}$, denote by $p_i^{(k)}$ the probability that the k-RPT selects an alternative from C_i. If for some $K \in \mathbb{N}$ and $\epsilon > 0$, $p_1^{(K)} \leq \epsilon \leq 2^{-12}$, then there exists $K' > K$ such that $p_3^{(K')} \leq \epsilon/2$ and $p_2^{(K')} \geq 1 - \sqrt{\epsilon}$.*

Proof. The event that some alternative from C_i is chosen by a perfect tree of height $k+1$ can be decomposed into the following two disjoint events: either an element from C_i appears at the left child of the root, and an element from C_i or $C_{(i \bmod 3)+1}$ at the right child, or an element from C_i appears at the right child and one from $C_{(i \bmod 3)+1}$ at the left. Thus, for all $k > 0$,

$$p_i^{(k+1)} = p_i^{(k)}\left(p_i^{(k)} + p_{(i \bmod 3)+1}^{(k)}\right) + p_i^{(k)} \cdot p_{(i \bmod 3)+1}^{(k)} = p_i^{(k)}\left(p_i^{(k)} + 2p_{(i \bmod 3)+1}^{(k)}\right), \tag{6}$$

It should be noted that (6) is independent of the structure of T inside the different components, but only depends on the relationship between them.

Now, consider the largest, possibly empty, set $S = \{K, K+1, K+2, \ldots,\}$ such that for all $k \in S$, $p_1^{(k)} + p_2^{(k)} \leq 1/2$. It then holds for all $k \in S$ that $2p_1^{(k)} + 2p_2^{(k)} \leq 1$, and, by (6), that $p_1^{(k+1)} \leq p_1^{(k)} \leq p_1^{(K)} \leq 2^{-12}$; that is, $p_1^{(k)}$ is weakly decreasing for indices in S, and since we assumed $p_1^{(K)} \leq 2^{-12}$, we have that $p_1^{(k+1)} \leq 2^{-12}$ for all $k \in S$. Since $p_2^{(k)} < 0.5$ and $p_3^{(k)} \geq 0.5$, we have that for all $k \in S$, $p_2^{(k)} + 2p_3^{(k)} > 1.3$. Hence, we conclude by (6) that for all $k \in S$, $p_2^{(k+1)} \geq 1.3 \cdot p_2^{(k)}$.

Choosing K_1 to be the smallest integer such that $K_1 \geq K$ and $K_1 \notin S$, we have that $p_1^{(K_1)} \leq \epsilon$ and $p_3^{(K_1)} \leq 1/2$. Also, by (6), for all $i = 1, 2, 3$ and all $k \in \mathbb{N}$, $p_i^{(k+1)} \leq 2p_i^{(k)}$. Choosing $L \geq 12$ such that $2^{-(L+1)} \leq \epsilon \leq 2^{-L}$, we have for all $k = K_1, \ldots, K_1 + \lfloor L/2 \rfloor - 1$,

$$p_1^{(k)} \leq \epsilon \cdot 2^{\lfloor L/2 \rfloor - 1} \leq \frac{2^{-\lceil L/2 \rceil}}{2} \leq \sqrt{\epsilon}/2. \tag{7}$$

By the assumption that $\epsilon \leq 2^{-12}$, this also implies for all such k that $p_1^{(k)} \leq 2^{-7}$.

We now claim that $K' = K_1 + \lfloor L/2 \rfloor - 1$ is as required in the statement of the lemma. Indeed, by applying (6), we have

$$p_3^{(K_1+1)} = p_3^{(K_1)}(p_3^{(K_1)} + 2p_1^{(K_1)}) \leq \frac{1}{2}(\frac{1}{2} + 2^{-6}) \leq 0.258,$$

and thus

$$p_3^{(K_1+2)} = p_3^{(K_1+1)}(p_3^{(K_1+1)} + 2p_1^{(K_1+1)}) \leq 0.258(0.258 + 2^{-6}) < 0.08.$$

Finally,

$$p_3^{(K_1+3)} = p_3^{(K_1+2)}(p_3^{(K_1+2)} + 2p_1^{(K_1+2)}) \leq 0.08(0.08 + 2^{-6}) < 0.0077.$$

Now, for $k = K_1 + 3, \ldots, K_1 + \lfloor L/2 \rfloor - 2$, $p_3^{(k+1)} \leq p_3^{(k)}(0.0077 + 2^{-6}) < p_3^{(k)}/2^5$, since $p_3^{(k)}$ is strictly decreasing for these values of k.

It also follows directly from the above discussion that

$$p_3^{(K')} \leq p_3^{(K_1+3)} \cdot (2^{-5})^{\lfloor L/2 \rfloor - 4} \leq 2^{-5} \cdot (2^{-5})^{\lfloor L/2 \rfloor - 4} = 2^{-5\lfloor L/2 \rfloor + 15}.$$

For $L \geq 12$, $2^{-5\lfloor L/2 \rfloor + 15} \leq 2^{-(L+2)} \leq \epsilon/2$. We therefore have that $p_3^{(K')} \leq \epsilon/2$, while $p_1^{(K')} \leq \sqrt{\epsilon}/2$ by (7). Furthermore, since $p_2^{(K')} = 1 - (p_1^{(K')} + p_3^{(K')})$, $p_2^{(K')} \geq 1 - \sqrt{\epsilon}$. ☐

We will now prove a stronger version of Theorem 5.1.

Lemma G.2. *For $k \in \mathbb{N}$, denote by Δ_k the distribution corresponding to the k-RPT. Then, for every set A of alternatives, $|A| \geq 5$, there exists a tournament $T \in \mathcal{T}(A)$ such that for every $K \in \mathbb{N}$ and $\epsilon > 0$, there exists $K' \geq K$ such that*

$$\frac{\mathbb{E}_{\Gamma \sim \Delta_{K'}}[s_{\Gamma(T)}]}{\max_{i \in A} s_i} \leq \frac{1 + \epsilon}{m - 2}.$$

Proof of Lemma G.2 and Theorem 5.1. Let $m \geq 5$, and define a tournament as in the statement of Lemma G.1 with components $C_1 = \{1\}$, $C_2 = \{2\}$, and $C_3 = \{3, \ldots, m\}$, such that C_3 is transitive.

We first show that there exists K_0 such that, using the notation of Lemma G.1, $p_1^{(K_0)} \leq 2^{-12}$. If $m \geq 2^{12}$, this holds trivially for $K_0 = 0$, since the uniform distribution selects each alternative with probability $1/m \leq 2^{-12}$. For $m < 2^{-12}$, the claim is easily verified using a computer simulation.

Now, by Lemma G.1, there exists K_1 such that $p_3^{(K_1)} \leq 2^{-13}$ and $p_2^{(K_1)} \geq 1 - 2^{-6}$. Renaming the components and applying Lemma G.1 again, there has to exist K_2 such that $p_2^{(K_2)} \leq 2^{-14}$ and $p_1^{(K_2)} \geq 1 - 2^{-13/2}$. Another application yields K_3 satisfying $p_1^{(K_3)} \leq 2^{-15}$ and $p_3^{(K_3)} \geq 1 - 2^{-7}$. Iteratively applying the lemma in this fashion, we get that there exists $K' \geq K$ such that $p_1^{(K')} \geq 1 - \epsilon'$, for $\epsilon' = \epsilon/(m-3)$. In this case, the approximation ratio is at most

$$\frac{(1 - \epsilon') + \epsilon' \cdot (m-2)}{m-2} = \frac{1 + \epsilon}{m-2}.$$

☐

H Composition of Caterpillars

In Section 5 we studied the ability of randomizations over balanced trees to improve the lower bound of Section 4, with somewhat unexpected results. A different approach to improve the randomized lower bound is to take a tree structure that provides a good lower bound, and construct a more complex tree by composing several trees of this type to form a new structure. Since a particular randomized tree chooses alternatives according to some probability distribution, this technique is conceptually closely related to probability amplification as commonly used in the area of randomized algorithms.

In our case, the obvious candidate to be used as the basis for the composition is the RSC, both because it provides the strongest lower bound so far, and because it can conveniently be analyzed using the stationary distribution of a Markov chain. We will thus focus on *higher order caterpillar trees* obtained by replacing each leaf of a caterpillar of sufficiently large height by higher order caterpillars with order reduced by one. To analyze the behavior of these higher order caterpillars on a particular tournament T, we again employ a Markov chain abstraction. Given a tournament T, we inductively define Markov chains $\mathfrak{M}_k = \mathfrak{M}_k(T)$ for $k \in \mathbb{N}$ as follows: for all k, the state space of $\mathfrak{M}_k$ is A. The initial distribution and transition matrix of $\mathfrak{M}_1$ are given by those of $\mathfrak{M}$ as defined in Section 4.1. For $k > 1$, the initial distribution of $\mathfrak{M}_k$ is given by the stationary distribution $\pi^{(k-1)}$ of $\mathfrak{M}_{k-1}$, which can be shown to exist and be unique using similar arguments as in Section 4.1. Its transition matrix $P_k = P_k(T)$ is defined as

$$P_k(i,j) = \begin{cases} \pi_i^{(k-1)} + \sum_{j':iTj'} \pi_{j'}^{(k-1)} & \text{if } i = j \\ \pi_j^{(k-1)} & \text{if } jTi \\ 0 & \text{if } iTj. \end{cases}$$

The class of tournaments used in Section 4.2 to show tightness of our analysis of ordinary caterpillars can also be used to show that the approximation ratio cannot be improved significantly by means of higher order caterpillars of small order. Perhaps more surprisingly, a different class of tournaments can be shown to cause the stationary distribution of $\mathfrak{M}_k$ to oscillate as k increases, leading to a deterioration of the approximation ratio. This phenomenon is similar to the one witnessed by the proof of Theorem 5.1.

Theorem H.1. *Let A be a set of alternatives, $|A| \geq 6$, and let $K \in \mathbb{N}$. Then there exists a tournament $T \in \mathcal{T}(A)$ and $k \in \mathbb{N}$ such that $K \leq k \leq K+5$ and the stationary distribution $\pi^{(k)}$ of $\mathfrak{M}_k(T)$ satisfies*

$$\sum_i \pi_i^{(k)} s_i \leq \frac{3}{m-2}.$$

Proof. Consider a tournament T with three components C_i, $1 \leq i \leq 3$ such that $C_i T C_j$ if $j = (i \ mod\ 3) + 1$ (as in the proof of Theorem 5.1).

For $i = 1, 2, 3$ and $k \in \mathbb{N}$, denote by $p_i^{(k)}$ the probability that an alternative from C_i is chosen from the stationary distribution of $\mathfrak{M}_k$. In particular, define $p_i^0 = |C_i|/m$. Since $p_i^{(0)} > 0$ for all i, and since T is strongly connected, $p_i^{(k)} > 0$ for all i and all $k \in \mathbb{N}$.

Then, for all $k \in \mathbb{N}$ and $i = 1, 2, 3$, and taking the subsequent index modulo three,

$$p_i^{(k+1)} = (1 - p_{i+2}^{(k)}) p_i^{(k+1)} + p_i^{(k)} p_{i+1}^{(k+1)},$$

and thus

$$p_i^{(k+1)} = \frac{p_i^{(k)}}{p_{i+2}^{(k)}} p_{i+1}^{(k+1)}.$$

Taking two steps, replacing $p_{i+1}^{(k+1)}$, and simplifying, we get

$$p_i^{(k+2)} = \frac{p_i^{(k+1)}}{p_{i+2}^{(k+1)}} p_{i+1}^{(k+2)} = \frac{p_i^{(k+1)}}{p_{i+2}^{(k+1)}} \cdot \frac{p_{i+1}^{(k+1)}}{p_i^{(k+1)}} \cdot p_{i+2}^{(k+2)} = \frac{p_i^{(k+1)} p_{i+1}^{(k)} p_{i+2}^{(k+1)} p_{i+2}^{(k+2)}}{p_{i+2}^{(k+1)} p_i^{(k)} p_i^{(k+1)}} = \frac{p_{i+1}^{(k)} p_{i+2}^{(k+2)}}{p_i^{(k)}},$$

and thus

$$\frac{p_{i+2}^{(k+2)}}{p_i^{(k+2)}} = \frac{p_i^{(k)}}{p_{i+1}^{(k)}}. \tag{8}$$

Analogously,

$$\frac{p_{i+1}^{(k+2)}}{p_i^{(k+2)}} = \frac{p_{i+2}^{(k)}}{p_{i+1}^{(k)}}. \tag{9}$$

Summing (8) and (9) and adding one,

$$\frac{p_i^{(k+2)} + p_{i+1}^{(k+2)} + p_{i+2}^{(k+2)}}{p_i^{k+2}} = \frac{p_i^{(k)} + p_{i+1}^{(k)} + p_{i+2}^{(k)}}{p_{i+1}^{(k)}}$$

and thus

$$p_i^{(k+2)} = p_{i+1}^{(k)}.$$

Choosing T such that $|C_1| = |C_2| = 1$ and $|C_3| = m - 2$, it holds for all k that

$$p_1^{(6k+4)} = p_3^{(0)} = \frac{m-2}{m}$$

and, since the sole vertex in C_1 has degree 1,

$$\sum_{i=1}^{m} \pi_i^{(6k+4)} s_i \leq \frac{m-2}{m} + \frac{2}{m} \cdot m \leq 3.$$

Observing that the sole vertex in C_2 has degree $m - 2$ completes the proof. ☐