On Approximately Strategy-Proof Tournament Rules for Collusions of Size at Least Three

David Mikšaník[1], Ariel Schvartzman[2](✉), and Jan Soukup[1]

[1] Computer Science Institute, Charles University, Prague, Czechia
soukup@kam.mff.cuni.cz
[2] Google Research, Mountain View, CA, USA
aschvartzman@google.com

Abstract. A tournament organizer must select one of n possible teams as the winner of a competition after observing all $\binom{n}{2}$ matches between them. The organizer would like to find a tournament rule that simultaneously satisfies the following desiderata. It must be *Condorcet-consistent* (henceforth, CC), meaning it selects as the winner the unique team that beats all other teams (if one exists). It must also be *strongly non-manipulable* for groups of size k at probability α (henceforth, k-SNM-α), meaning that no subset of $\leq k$ teams can fix the matches among themselves in order to increase the chances any of it's members being selected by more than α. Our contributions are threefold. First, wee consider a natural generalization of the Randomized Single Elimination Bracket rule from [18] to d-ary trees and provide upper bounds to its manipulability. Then, we propose a novel tournament rule that is CC and 3-SNM-$1/2$, a strict improvement upon the recent work of [7] who proposed a CC and 3-SNM-$31/60$ rule. Finally, we initiate the study of reductions among tournament rules.

Keywords: Tournament design · Strategy-proof rules · Computational Social Choice

1 Introduction

Consider the problem a tournament organizer faces when, after observing all pairwise matches between n teams, they must select one as the winner of the tournament. We model the tournament T as a complete, directed graph on the n teams. A tournament rule r is a (possibly randomized) mapping from the set of tournaments on n teams $\mathcal{T}_n$ to a probability vector in Δ^n, $r : \mathcal{T}_n \rightarrow \Delta^n$. The tournament organizer is thus tasked with designing a tournament rule r and would like the rule to satisfy the following natural properties:

1. If there is a team who beats all other teams, termed a *Condorcet-winner*, they should be picked as the winners of the tournament with probability 1. We call such rules *Condorcet-consistent* (or CC).
2. No team should be incentivized to unilaterally throw their own games in order to obtain a better outcome. We call such rules *monotone*.

R. Freeman and N. Mattei (Eds.): ADT 2024, LNAI 15248, pp. 33–47, 2025.
https://doi.org/10.1007/978-3-031-73903-3_3

3. No subset of $\leq k$ teams should have incentives to fix the matches among themselves in order to improve the chances of any of its members being selected as the winner of the tournament. We call such rules *k-strongly non-manipulable* (or k-SNM).

These properties are motivated by real-world sports competitions. It would be unimaginable to violate Property 1 and not award the top prize to an undefeated team. Violations to Properties 2, 3 have been observed in high-stakes competitions such as the Olympic Games and the FIFA World Cup. An infamous scandal in the Women's Doubles Badminton tournament at the 2012 Olympics saw multiple teams purposefully losing the last games of their group stage matches in order to avoid a difficult match-up in the following single-elimination bracket. This clear violation of monotonicity (and sportsmanship) resulted in the disqualification of 4 teams, including many of the likely medalists. A less investigated but equally egregious scandal occurred during the 1982 FIFA World Cup. West Germany and Austria disputed the last match of their group stage with full knowledge of the outcomes of all other games. It is suspected that they colluded in order to produce an outcome that would see both teams advance to the next stage of the tournament, at the expense of Algeria, who unexpectedly defeated West Germany in their opening match. As a result of this possible violation of strong non-manipulability, the last game of every group in every FIFA World Cup since has been played simultaneously.

Observe that if one wanted to simply satisfy Properties 2, 3, there are numerous simple rules that do so. For example, picking a winner uniformly at random, picking a fixed team as the winner (i.e., a dictatorship) or picking the winner proportional to the number of wins in the tournament all satisfy Properties 2, 3 but not Property 1. Similarly, it is easy to satisfy Properties 1, 2. If there is a Condorcet-winner, pick that team. Otherwise, pick a team uniformly at random. Unfortunately, it is known that Properties 1, 3 are directly at odds with each other: [1] showed that there exists no randomized tournament rule that can satisfy both of these properties at the same time, even for $k = 2$. One way to overcome this impossibility result is to relax Property 3 as follows:

4. No subset of $\leq k$ teams should be able to fix the matches among themselves in order to improve the chances of any of its members being selected as the winner by more than α. We call such rules k-SNM-α.

A growing body of work has asked what is the smallest α for which there exists a rule that satisfies Properties 1, 2 and 4 (for some fixed value of k). First, the work of [18] proves that the Random Single-Elimination Bracket (henceforth RSEB) rule is CC and 2-SNM-$1/3$, and that no other CC rule can do better. Later, [19] show that a rule termed Randomized King-of-the-Hill (henceforth, RKOTH) matches the performance of RSEB and satisfies a condition even stronger than CC. These works completely settle the question of finding CC and minimally manipulable (or *optimal*) strategy-proof rules for collusions of size $k = 2$. On the other hand, very little is known about the case when $k > 2$, even $k = 3$. [18] prove a simple lower bound of $\alpha \geq \frac{k-1}{2k-1}$. [19] proved that there exists an LP-based rule that is CC and k-SNM-$2/3$ *for all* k simultaneously. Unfortunately, this rule is neither explicit nor monotone. More recently, in concurrent work [7] gave the first explicit, CC, monotone, 3-SNM-α rule for $\alpha = 31/60$ (and this value of α is tight for their rule).

As hinted in the previous paragraph, there are two approaches to proving the existence of CC and approximately strategy-proof tournament rules. One approach takes simple rules (such as RSEB, RKotH), in the hopes that they are not too manipulable, and provides tight analysis for them. This does not always work: many simple rules, such as picking the team with the most wins, are extremely manipulable (i.e., have $\alpha = 1 - O(1/n)$, see [18]). The other approach provides rules which are not explicitly implementable. For example, the LP-based rule of [19] arises from fixing the tournament rule for tournaments with a Condorcet-winner, relaxing the manipulability constraints for tournaments that are close to having a Condorcet-winner and proving that the resulting polytope is non-empty for some value of $\alpha < 1$. Our results make use of both of these approaches and introduce another.

Our first contribution, inspired by the positive results of [18], introduces a natural generalization of RSEB. The RSEB rule randomly seeds teams on the leaf nodes of a binary tree and recursively labels inner nodes as the winner of the match between its children. The winner of the bracket is the team whose label appears in the root node of the tree. We define the Random d-ary Single Elimination Bracket (henceforth RdSEB) rule similarly with one key difference: instead of using binary trees, we use d-ary trees. If the sub-tournament induced by an inner node's children has a Condorcet-winner, then the inner node will carry the label of that child. Otherwise, the inner node picks a child uniformly at random to advance[1]. We provide an upper bound $\alpha_{d,k} < 1$ on the manipulability of RdSEB on tournaments with n teams and collusions of size up to $k \leq d$.

Theorem 1. *Let $2 \leq k \leq d$. The* RdSEB *rule is Condorcet-consistent, monotone and* k-SNM-$\alpha_{d,k}$ *for*

$$\alpha_{d,k} \leq 1 - \left(\frac{2 \cdot (d)_k}{d^{k+1}}\right),$$

where $(d)_k = \prod_{i=0}^{k-1}(d-i)$ is the falling factorial of d with k terms.

For $k = d = 3$, we obtain that $\alpha_{3,3} = .8519$. As a consequence of Theorem 1, we get the first explicit family of CC, monotone rules whose manipulability for any (fixed) k is bounded away from 1. This stands in contrast to the LP-based rule of [19] which was neither monotone nor explicit and to several of the rules analyzed in [18] which had $\alpha \to 1$ as $n \to \infty$, even for $k = 2$. In other words, the bound from Theorem 1 is independent of n, the number of competing teams.

Our second contribution, inspired more by the second approach to finding approximately optimal tournament rules, is a new explicit tournament rule which strictly improves upon the results of [7].

Theorem 2. *The* SignificantOnly *rule is Condorcet-consistent, monotone,* 2-SNM-$1/3$ *and* 3-SNM-$1/2$*, and this is tight.*[2]

[1] This decision is inspired by the observation that tournaments on three teams either have a Condorcet-winner or have three teams that beat each other cyclically.

[2] The best lower bound for this problem, due to [18], is $2/5$.

The SignificantOnly rule, while explicitly describable, arises from an approach similar to the LP-based rule of [19]. We identify the tournaments which are close to having a Condorcet-winner as those where teams have more incentives to manipulate outcomes. Given such a close-to-Condorcet tournament T, our rule deems a small number of teams as *significant* and distributes most of the probability mass on these teams according to additional properties of the tournament itself (and the rest uniformly across the remaining teams).

Finally, our last contribution introduces a new way of designing Condorcet-consistent and *asymptotically* optimal tournament rules. If one had substantial computational power, one could compute a top-cycle consistent,[3] optimal rule for fixed values of n, k. How could we use such a rule r_n to construct a rule $r_{n'}$ that works for $n' > n$? First pad the tournament with dummy teams that lose to all real teams until the number of teams $n' := n \cdot M$ is a multiple of n. Partition the teams into n groups of equal size. Within each group, pick a team from the top-cycle uniformly at random as a finalist. The number of finalists will be exactly n. Finally, run r_n on the n finalists and declare its winner as the overall winner. We prove that this simple idea suffices to transform top-cycle consistent, k-SNM-α rules for n teams to CC, k-SNM-α' rules for $n' > n$ teams where α' is close to α.

Theorem 3. *If there exists a top-cycle consistent, and k-SNM-α rule r for n teams, then there exists a top-cycle consistent and k-SNM-α' rule r' for $n' > n$ teams where*

$$\alpha' \leq \alpha\Big(1 - \frac{(k-1)^2}{n}\Big) + \frac{(k-1)^2}{n}.$$

Theorem 3 can be thought of as reducing the problem of finding approximately optimal tournament rules for large n to the same problem for small n. As an example, if we could verify the existence of a top-cycle consistent 3-SNM-$2/5$ rule for $n = 25$, then this would imply top-cycle consistent 3-SNM-α rules for $n \geq 25$ and $\alpha < 1/2$, directly improving on Theorem 2.

1.1 Related Work

Most of the related work has been mentioned already. There exist two other results that are directly related to our problem. Whereas in this paper (and all previously mentioned papers) we evaluate the manipulability of tournament rules based on their worst-case performance, the work of [8] instead studies this question under the lens of average-case analysis. More recently, the work of [6] expanded the model to include prize vectors $v = (v_1, \ldots, v_n)$. Tournament rules with prizes output a complete, linear ranking over the teams where the i-th team earns reward v_i, rather than giving reward 1 to the winner and 0 to everyone else.

Another related line of work involves the Tournament Fixing Problem, where the organizer of the tournament is colluding with a team in order to produce a seeding that

[3] Top-cycle consistency is stronger than Condorcet-consistency. We defer its formal definition but informally, the top-cycle is the smallest non-empty set S of teams such that no team in S loses to a team outside of S. A top-cycle consistent rule would only pick teams from the top-cycle.

selects them as the winner (see, e.g. [2,12,13,20,22,23]). The central questions here are computational (i.e., can the organizer efficiently decide if there is a winning bracket for their favorite team) and structural (i.e., under what conditions does there exist a bracket that selects the organizer's favorite team, see, e.g. [15]). Due to its connections with voting theory and social choice, there is a long history of analyzing properties of particular tournament rules ([4,5,9–11,16,17], to name a few). We refer the reader to the survey by [21] on recent developments in tournaments and computational social choice, or other books on computational social choice [3,14].

2 Notation

In this section, we introduce key concepts to contextualize our results. Recall a tournament graph T is a complete, directed graph $G = ([n], E)$ on n labelled vertices. We refer to a tournament graph's vertices as teams (and the number of teams as the size of the tournament), its undirected edges as matches and its directed edges as outcomes (where if $(i, j) \in E$, we say i beats j in T). For a fixed team i, tournament T we let $\delta^+(i, T) = \{j | j \in [n], (i, j) \in E)\}$ be the set of teams that i beats under T, and let $\delta^-(i, T) = \{j | j \in [n], (j, i) \in E)\}$ be the set of teams that i loses to under T. Let $\mathcal{T}_n$ be the set of all tournaments on n teams. Recall a tournament rule is a mapping $r : \mathcal{T}_n \to \Delta^n$. That is, for every tournament T, $r(T)$ denotes the distribution over teams according to which the organizer will select a winner. We use notation $r_i(T)$ to denote team i's probability of being selected by r as the winner under tournament T. We use the shorthand notation $r_S(T) := \sum_{i \in S} r_i(T)$ to denote the probability that a team in S is selected by r in tournament T. The next definitions formalize Properties 1, 2, 3 and 4.

Definition 1. *Team i is the* Condorcet-winner *of tournament T if i beats every other team under T. A rule r is* Condorcet-consistent *if $r_i(T) = 1$ when i is T's Condorcet-winner.*

Definition 2. *A tournament rule r is* monotone *if for all teams i and all tournaments T, T' where all matches not involving team i are identical and $\delta^+(i, T) \supseteq \delta^+(i, T')$, it holds that $r_i(T) \geq r_i(T')$.*

A tournament rule is monotone if it is not in a team's best interest to unilaterally lose matches it would otherwise win. We present the ways in which manipulations are modelled.

Definition 3. *We say tournaments T, T' are S-*adjacent *if the only outcomes where T, T' differ on are those matches that involve two teams in S.*

If T, T' are S-adjacent, outcomes involving at least one team outside of S are identical. Motivated by the results from [1], the following relaxation was introduced by [18].

Definition 4. *A tournament rule is k* strongly non-manipulable at probability *α (k-*SNM*-α) if for all $S \subseteq [n]$ of size at most k, for all tournaments T, T' that are S-adjacent we have $r_S(T') \leq r_S(T) + \alpha$. For $\alpha = 0$, we simply say the rule is k* strongly non-manipulable *(k-*SNM*).*

3 Analysis for the RdSEB Rule

In this section we study the manipulability of the Randomized d-ary Single-Elimination Bracket (RdSEB) rule, a generalization of the Randomized Single-Elimination Bracket (RSEB) rule from [18], against collusions of size $k \leq d$.

Definition 5. *Given a tournament T, a* sub-tournament *on S is the sub-graph induced by T on vertex set S.*

We are now ready to formally define the RdSEB rule.

Definition 6. *The* Randomized d-ary Single-Elimination Bracket *rule operates as follows. Add dummy teams*[4] *until the number of teams $n = d^{\lceil \log_d(n) \rceil}$ is a power of d. Randomly place teams at the leaf nodes of a complete d-ary tree of height $\lceil \log_d(n) \rceil$. Recursively label a parent node with the label of the Condorcet-winner of the sub-tournament induced by the labels of its children, if there is one. Otherwise, choose one of its children uniformly at random and use that label instead. The winner of the tournament is the team whose label appears at the root of the tree.*

In terms of Definition 6, RSEB is the rule that results from setting $d = 2$. The family of RdSEB rules operates in the same way as the RSEB rule except that if there is no Condorcet-winner in the sub-tournament induced at a node, the rule advances a team uniformly at random. This choice is motivated by the following simple observation when $d = 3$. There are only two non-isomorphic sub-tournament graphs on three teams: one where there is a Condorcet-winner and one where the teams beat each other cyclically. In the former case, it is obvious which team to advance. In the latter case, we argue choosing a team uniformly at random is reasonable.

Theorem 1. *Let $2 \leq k \leq d$. The* RdSEB *rule is Condorcet-consistent, monotone and* k-SNM-$\alpha_{d,k}$ *for*

$$\alpha_{d,k} \leq 1 - \left(\frac{2 \cdot (d)_k}{d^{k+1}}\right),$$

where $(d)_k = \prod_{i=0}^{k-1}(d-i)$ is the falling factorial of d with k terms.

The main idea is that if some (at least two) colluding teams meet only in the final round, then they can increase the joint probability that one of them will be winner by at most $(1 - 2/d)$ (which happens when the colluding teams create a Condorcet winner). Obviously, if the colluding teams do not meet in a bracket at all or there exists a Condorcet winner outside the colluding teams, then they cannot increase the chance to win the tournament. In the remaining cases, we simply assume that they can increase the chance to win the tournament by 1.

From Theorem 1, given a fixed k, the RkSEB rule is monotone, CC and its manipulability is bounded away from 1 for all n. This is the first explicit family of rules to exhibit this property (since the LP-based rule of [19] is neither monotone nor explicit). We suspect the bound from Theorem 1 is not tight. A finer argument like the one in [18] might yield a better analysis.

[4] Dummy teams are teams that lose to all non-dummy teams. The outcome of a match between two dummy teams is arbitrary.

Proof. It is more convenient to consider the following equivalent variation of the RdSEB rule. Instead of choosing one of the children of A uniformly at random, chose a number j_A from $[d]$ uniformly at random. If there is no Condorcet-winner in the sub-tournament, label A by the label of j_A-th child of A. In both cases, mark the node A with j_A (even if there is a Condorcet-winner in the sub-tournament).[5] For the sake of the proof, let us denote r^d as this equivalent variation of the RdSEB rule. Observe that the labels of inner nodes of the complete d-ary tree can be deduced from the labels of leaves and marks of inner nodes: run r^d but all random choices are made accordingly to these labels and marks.

For any non negative number t, let D_t be a reserved set containing t dummy teams. Moreover, let T be a tournament on N' (disjoint from any D_t) of n' teams, $h := \lceil \log_d(n') \rceil$, and $n := d^h$. We assume that the rule r^d initially adds $n - n'$ dummy teams from $D_{n-n'}$ into T.[6] Given $N := N' \cup D_{n-n'}$, a *d-bracket* $G(\pi, m)$ for N is a complete d-ary tree G of height h endowed with a pair (π, m), where

- π is a bijection from leaves of G to N (i.e., labels of leaves),
- m is a mapping from inner nodes of G to $[d]$ (i.e., marks of inner nodes).

Let $\mathcal{B}_N$ be the set of all d-brackets for N. Now we precisely describe how the labels of nodes of G can be deduced from π and m. Given a d-bracket $G(\pi, m)$ for N, the *outcome* of $G(\pi, m)$ under T[7] is a labeling ω_T of nodes of G such that $\omega_T(A) = \pi(A)$, for every leaf A, and if A is a node with children $A_1, \ldots, A_d$, then

$$\omega_T(A) := \begin{cases} x & \text{if } x \text{ is the Condorcet-winner in} \\ & \text{the sub-tournament induced by} \\ & \{\omega_T(A_1), \ldots, \omega_T(A_d)\} \text{ under } T, \\ \omega_T(A_{m(A)}) & \text{otherwise.} \end{cases}$$

Observe that ω_T is a one-to-one correspondence between the set of all outcomes of d-ary brackets on N under T and the set of all runs of r^d on T.

Fix a d-bracket $G(\pi, m)$ for N. The *winner* of $G(\pi, m)$ under T is $\omega_T(R)$, where R is the root of G. Given a team $x \in N$, $G(\pi, m)$ is *winning for* x under T if x is the winner of $G(\pi, m)$ under T. Denote by $\mathcal{B}_{N,T}(x) \subseteq \mathcal{B}_N$ the set of all winning d-brackets for x under T. The motivation behind this reframing of RdSEB is to simply the argument of the proof. Similar to the original argument of [18], we will bound the manipulability of RdSEB by directly counting the number of brackets where colluding teams could gain and compare it to the total number of brackets. We have

$$r^d_x(T) = |\mathcal{B}_{N,T}(x)|/|\mathcal{B}_N|.$$

Observe that $|\mathcal{B}_N| = n! \cdot d^\ell$, where $\ell := \ell(d, h)$ is the number of inner nodes of a complete d-ary tree of height h.

[5] Every inner node has a label and a mark. Note that they can be equal but they have different meaning.

[6] Hence every tournament on n' teams is extended by the same set of dummy teams.

[7] The outcome is well-defined only if the teams in T are the same as the set N.

First, we prove that r^d is monotone. Take an arbitrary team $x \in N'$. It sufficient to show that $r^d_x(T) \geq r^d_x(T')$ for every $\{x, y\}$-adjacent tournaments T' of T such that x beats y under T. Let $G(\pi, m)$ be a d-bracket for N. Observe that if $G(\pi, m)$ is winning for x under T', then $G(\pi, m)$ is also a winning for x under T. Hence $\mathcal{B}_{N,T'}(x) \subseteq \mathcal{B}_{N,T}(x)$, and so $r^d_x(T') \leq r^d_x(T)$ as required.

Second, we prove that r^d is Condorcet-consistent. Suppose that $x \in N'$ is the Condorcet-winner in T. For every d-bracket $G(\pi, \ell)$ for N, consider the unique path P from the leaf in G labeled by x to the root of G. Observe that every node of P is labeled by x by the outcome of $G(\pi, m)$ under T. In particular, the root of G is labeled by x. It follows that every d-bracket $G(\pi, \ell)$ is winning for x under T. Hence $\mathcal{B}_{N,T}(x) = B_N$, and so $r^d_x(T) = 1$ as required.

Lastly, we prove that r^d is k-SNM-$\alpha_{d,k}$ for some $\alpha_{d,k}$ (to be determined). Suppose that $S = \{s_1, s_2, \dots, s_k\} \subseteq N'$ is a subset of colluding teams. For any S-adjacent tournaments T' of T, we show that $r^d_S(T') - r^d_S(T) \leq \alpha$. Recall that $\mathcal{B}_{N,T}(x)$ is the set of all winning d-brackets for x under T. Moreover, define $\mathcal{B}_{N,T}(S) := \bigcup_{s \in S} \mathcal{B}_{N,T}(S)$. In this notation, we can write

$$r_S(T') - r_S(T) = \frac{|\mathcal{B}_{N,T'}(S)|}{|\mathcal{B}_N|} - \frac{|\mathcal{B}_{N,T}(S)|}{|\mathcal{B}_N|}.$$

We upper bound this expression using the idea introduced in [18]. For that, let us denote by $\mathcal{B}^+_N(S)$ the set of all d-brackets $G(\pi, m)$ for N such that the least common ancestor of any leaves A and B with $\pi(A), \pi(B) \in S$ is the root of $B(\pi, m)$. In other words, $\mathcal{B}^+_N(S)$ is the set of all d-brackets for N such that the colluding teams can meet possibly only in the final round. Set $\mathcal{B}^-_N(S) := \mathcal{B}_N \setminus \mathcal{B}^+_N(S)$. Moreover, for a tournament T, let $\mathcal{B}^+_{N,T}(S) := \mathcal{B}^+_N(S) \cap \mathcal{B}_{N,T}(S)$ and $\mathcal{B}^-_{N,T}(S) := \mathcal{B}^-_N(S) \cap \mathcal{B}_{N,T}(S)$. Then

$$\begin{aligned} r_S(T') - r_S(T) &= \frac{|\mathcal{B}_{N,T'}(S)|}{|\mathcal{B}_N|} - \frac{|\mathcal{B}_{N,T}(S)|}{|\mathcal{B}_N|} \\ &= \frac{|\mathcal{B}^+_{N,T'}(S)| + |\mathcal{B}^-_{N,T'}(S)|}{|\mathcal{B}_N|} - \frac{|\mathcal{B}^+_{N,T}(S)| + |\mathcal{B}^-_{N,T}(S)|}{|\mathcal{B}_N|} \\ &= \frac{|\mathcal{B}^+_{N,T'}(S)| - |\mathcal{B}^+_{N,T}(S)|}{|\mathcal{B}_N|} + \frac{|\mathcal{B}^-_{N,T'}(S)| - |\mathcal{B}^-_{N,T}(S)|}{|\mathcal{B}_N|} \\ &\leq \frac{|\mathcal{B}^+_{N,T'}(S)| - |\mathcal{B}^+_{N,T}(S)|}{|\mathcal{B}_N|} + \frac{|\mathcal{B}^-_N(S)|}{|\mathcal{B}_N|}. \end{aligned}$$

We upper bound the first term in the last expression. Let $G := G(\pi, m)$ be a d-bracket for N in $\mathcal{B}^+_{N,T'}(S)$. Let $\mathbf{x} := (x_1, \dots, x_d)$ be the d-tuple of finalists in G under T'. More precisely, let R be the root of G with children $A_1, \dots, A_d$. Then d-tuple of finalists of G under T is $(\omega_{T'}(A_1), \dots, \omega_{T'}(A_d))$. The crucial observation is that also $\mathbf{x}$ is d-tuple of finalist of G under T. We say that there is a Condorcet-winner in $\mathbf{x}$ under T' if there is a Condorcet-winner in the sub-tournament induced by x under T'. Notice that:

(i) If $|S \cap \{x_1, x_2, \dots, x_d\}| \leq 1$, then $G \in \mathcal{B}^+_{N,T}(S)$.

(ii) If $|S \cap \{x_1, x_2, \ldots, x_d\}| \geq 2$ and there is no Condorcet-winner in **x** under T', then $G \in \mathcal{B}^+_{N,T}(S)$.
(iii) If $|S \cap \{x_1, x_2, \ldots, x_d\}| \geq 2$ and there is a Condorcet-winner $s_i \in S$ in **x** under T', then $G \notin \mathcal{B}^+_{N,T}(S)$ only if there is no Condorcet-winner $s_j \in S$ in **x** under T and the mark of the root R of G is pointing to a team outside of S (i.e., $\omega_T(A_{m(R)}) \notin S$).

A d-bracket is of type (i) if it satisfies the statement (i). Analogously, for types (ii) and (iii). It follows that, for every d-bracket in $\mathcal{B}^+_{N,T'}(S)$ of type (i) or (ii), there exists at least one d-bracket in $\mathcal{B}^+_{N,T}(S)$ of the same type. Moreover, we claim that, for every d d-bracket in $\mathcal{B}^+_{N,T'}(S)$ of type (iii), there exist at least two d-brackets in $\mathcal{B}^+_{N,T}(S)$ of type (iii). Indeed, take a bracket $G(\pi, m) \in B^+_{N,T'}(S)$ of type (iii). Then $G(\pi, m') \in B^+_{N,T'}(S)$ is of type (iii), where only the mark of the root of G can be changed (there are d ways how to change it). On the other hand, let $i \neq j$ be two indices such that $x_i, x_j \in S$. Then $G(\pi, m_1), (\pi, m_2) \in B^+_{N,T}(S)$ are of type (iii), where m_i and m_j are m but the mark of the root of G is changed to i and j, respectively.

If we denote by p be the number of d-brackets from $\mathcal{B}^+_{N,T'}(S)$ of type (iii), then

$$|\mathcal{B}^+_{N,T'}(S)| - |\mathcal{B}^+_{N,T}(S)| \leq p \cdot \left(1 - \frac{2}{d}\right)$$
$$\leq |\mathcal{B}^+_{N,T'}(S)| \cdot \left(1 - \frac{2}{d}\right) \leq |\mathcal{B}^+_{N}(S)| \cdot \left(1 - \frac{2}{d}\right).$$

Observe that

$$|\mathcal{B}^+_{N}(S)| = (d)_k \cdot \left(\frac{n}{d}\right)^k \cdot (n-k)! \cdot d^\ell$$

and hence

$$|\mathcal{B}^-_{N}(S)| = n! \cdot d^\ell - (d)_k \cdot \left(\frac{n}{d}\right)^k \cdot (n-k)! \cdot d^\ell.$$

Therefore,

$$\begin{aligned} r_S(T') - r_S(T) &\leq \frac{|\mathcal{B}^+_{N,T'}(S)| - |\mathcal{B}^+_{N,T}(S)|}{|\mathcal{B}_N|} + \frac{|\mathcal{B}^-_N(S)|}{|\mathcal{B}_N|} \\ &\leq \frac{\left((d)_k \cdot \left(\frac{n}{d}\right)^k \cdot (n-k)! \cdot d^\ell\right) \cdot \left(1 - \frac{2}{d}\right)}{n! \cdot d^\ell} \\ &+ \frac{n! \cdot d^\ell - (d)_k \cdot \left(\frac{n}{d}\right)^k \cdot (n-k)! \cdot d^\ell}{n! \cdot d^\ell} \\ &= 1 - \frac{1}{n!} \cdot \left(\frac{2}{d} \cdot (d)_k \cdot \left(\frac{n}{d}\right)^k \cdot (n-k)!\right) \\ &= 1 - \frac{n^k \cdot (n-k)}{n!} \cdot \frac{2 \cdot (d)_k}{d^{k+1}} \\ &\leq 1 - \frac{2 \cdot (d)_k}{d^{k+1}}. \end{aligned}$$

□

We conjecture that the manipulability of RdSEB is bounded away from $1/2$ for all $k \leq d$.

Conjecture 1 (See Table 1). For all $3 \leq k \leq d$, $\alpha_{d,k} \geq 227/420$.

Table 1. Evaluation of $\alpha_{d,k}$ for small values of d and k rounded up to 4 decimals.

d	k				
	3	4	5	6	7
3	0.8519	–	–	–	–
4	0.8125	0.9531	–	–	–
5	0.808	0.9232	0.9846	–	–
6	0.8148	0.9074	0.9691	0.9949	–
7	0.8250	0.9	0.9572	0.9878	0.9983

4 Analysis for the SignificantOnly Rule

In this section we formalize the SignificantOnly rule and outline the proof of Theorem 2. The motivation behind the SignificantOnly rule is captured by the following simple observation. Suppose we are trying to design a rule which is CC and k-SNM-α for some fixed k. Take any tournament graph T. If there is a Condorcet-winner, then the rule is fixed and must declare that team the winner. If T is *far* from having a Condorcet-winner, meaning that the smallest set of teams who could collude and produce a Condorcet-winner among them is larger than k, then pick a team uniformly at random. If T is only a small number of manipulations away from having a Condorcet-winner, then in order to satisfy the SNM constraint (approximately) we must allocate all (resp. most) of the probability mass to the teams who could produce a Condorcet-winner. However, we must be careful in order not to incentivize teams in tournaments that are far from having a Condorcet-winner to manipulate into those which are close to having a Condorcet-winner.

The previous paragraph captures the spirit of the SignificantOnly rule. We first partition the set of all tournament graphs $\mathcal{T}_n$ into three groups: those with a Condorcet-winner (*Condorcet* tournaments), those where (maybe multiple) sets of at most k teams can produce a Condorcet-winner (*near-Condorcet* tournaments) and those where no set of k teams can produce a Condorcet-winner (*far-Condorcet* tournaments). The rule is fixed for the first part of the partition, and will select a winner uniformly at random on the last part of the partition. The middle part of the partition is further partitioned into four categories depending on the exact number and size of the groups that could produce a Condorcet-winner. For each of the parts in this sub-partition, we propose a way to distribute the probability mass. We must balance two different sets of incentives here. On the one side, we must give sufficient probability mass to the groups that can produce a Condorcet-winner. On the other, we can't give them *too* much mass, since otherwise teams in far-Condorcet tournaments might be substantially incentivized into manipulating the tournament into a near-Condorcet one, where some of the colluding members are also in groups that can now produce Condorcet-winners. We formalize the above with the following definitions.

Definition 7. *A tournament T is said to be* near-Condorcet *for k if there is no Condorcet-winner in T but there exists at least one team i with $|\delta^-(i,T)| \leq k-1$. Call the set of teams $MW(i,T) := \{i\} \cup \delta^-(i,T)$ a* minimal winning group *(MW group). We call team i the* leader *of $MW(i,T)$ and any team in $j \in MW(i,T)$* significant. *If $|MW(i,T)| = 2, 3$ we call it an MW pair or an MW triple, respectively.*

We first prove simple structural properties of near-Condorcet tournaments.

Lemma 1. *Every minimal winning group has exactly one leader.*

Proof. Suppose there exists some other leader j of $MW(i,T)$. By definition, j must lose to i and vice versa, a contradiction. □

Notice that even though every MW group has a different unique leader, a leader of one group can be a member in a different group. Moreover, it can happen that one MW group is a subset of another MW group.

Lemma 2. *Let T be a near-Condorcet tournament, and $MW(i,T)$ and $MW(j,T)$ be distinct MW groups in T with leaders i and j, respectively. Then either i is in $MW(j,T)$, or j is in $MW(i,T)$. In particular, $MW(i,T)$ and $MW(j,T)$ must have a non-empty intersection.*

Proof. Suppose that both i and j in $MW(i,T) \cap MW(j,T)$, or they are both outside. Then, by the definitions of $MW(i,T)$, $MW(j,T)$, they should beat each other, a contradiction. Thus, exactly one of them must be in the intersection. □

We now prove the main lemma about the structure of near-Condorcet tournaments for $k = 3$, showing that there are not too many significant teams and leaders.

Lemma 3. *If $k = 3$ and T is a near-Condorcet tournament, then the number of significant teams in T is at most 6, and the number of teams in the union of MW pairs is at most 3. Furthermore, there cannot be more than 3 MW pairs. If there is exactly one MW pair, then the maximal number of significant teams is 5, if there are exactly two MW pairs, then the maximal number of significant teams is 4, and if there are exactly three pairs, then the maximal number of significant teams is 3.*

Proof. Assume there are p MW pairs and t MW triples in T. By Lemmas 1 and 2, we know that each such group has a unique leader, and these leaders are pairwise distinct. Furthermore, each leader of an MW pair loses exactly one match, and each leader of an MW triple loses exactly two matches. Therefore, all leaders together lose exactly $p+2t$ matches. On the other hand, there are exactly $\binom{p+t}{2}$ matches between leaders, and in each of these matches, one leader loses. Hence,

$$\binom{p+t}{2} \leq p + 2t.$$

We can equivalently rewrite is as $p^2 + t^2 + 2pt - 3p - 5t \leq 0$. If $p \geq 4$ we get that $4+3t+t^2 \leq 0$, which does not have a solution for a non-negative integer t. Hence, there are either 3, 2, 1, or no MW pairs. Furthermore, significant teams are either leaders of some groups or they beat some leaders. Hence, there are at most the number of leaders plus the number of matches lost by leaders to some non-leaders many of them. In other words, there at most $(p + t) + p + 2t - \binom{p+t}{2}$ significant teams. If there are three pairs, the maximum of this expression for non-negative integer t is 3. Similarly, if $p = 2$ the maximum is 4, if $p = 1$ the maximum is 5, and if $p = 0$ the maximum is 6. Moreover, there cannot be more than 3 teams in the union of all MW pairs. Otherwise, there would be at least 3 MW pairs (there cannot be only two because they intersect), but we already showed that if we have three or more MW pairs, there can be at most 3 teams in the union of all MW groups. □

This lemma implies that in every near-Condorcet tournament at most 6 teams can be directly part of some manipulating group creating a Condorcet-winner. Hence, we can directly design a rule that is 3-SNM-$\frac{2}{3}$. It is sufficient to assign probability 1 to Condorcet-winners, probability $\frac{1}{6}$ to significant teams in near-Condorcet tournaments, and distribute the remaining probabilities in all tournaments uniformly between the remaining teams. We will not prove this formally since we design a better rule, but we believe this intuition is useful in understanding our construction. We now formalize the partition of tournament graphs $\mathcal{T}_n$.

– Let $\mathcal{CC}_n \subset \mathcal{T}_n$ be the set of tournaments with a Condorcet-winner.
– Let $\mathcal{FC}_n \subset \mathcal{T}_n$ be the set of far-Condorcet tournaments.
– Let $i{-}\mathcal{NCP}_n \subset \mathcal{T}_n$ be the set of near-Condorcet tournaments with exactly i MW pairs for $i = 0, 1, 2, 3$.

Now we are ready to define our main rule.

Definition 8 (SIGNIFICANTONLY Tournament Rule). *For $n \geq 6$ teams, the* SIGNIFICANTONLY tournament rule *does the following.*

1. *If $T \in \mathcal{CC}_n$ the Condorcet-winner gets 1, and the remaining teams get zero.*
2. *If $T \in \mathcal{FC}_n$, pick a winner uniformly at random.*
3. *If $T \in 3{-}\mathcal{NCP}_n \cup 2{-}\mathcal{NCP}_n$, teams in MW pairs get $\frac{1}{3}$, and the remaining teams get zero.*
4. *If $T \in 1{-}\mathcal{NCP}_n$, teams in the only MW pair get $\frac{1}{3}$, teams in MW triples that are not in any MW pair get $\frac{1}{9}$, and the remaining teams get the remaining probability mass evenly distributed between them.*
5. *If $T \in 0{-}\mathcal{NCP}_n$, teams in MW triples get $\frac{1}{6}$, and the reaming teams get the remaining probability mass evenly distributed between them.*

We restate the main result of this section.

Theorem 2. *The* SIGNIFICANTONLY *rule is Condorcet-consistent, monotone,* 2-SNM-$1/3$ *and* 3-SNM-$1/2$*, and this is tight.* (See Footnote 2)

The fact that SIGNIFICANTONLY is Condorcet-consistent follows directly from the definition. The proof of Theorem 2 and other missing proofs can be found in the full version of this paper.

5 Almost Optimal Rules for Large n from Optimal Rules for Small n

For fixed values of n, k, [18] present an LP that can compute the minimally manipulable, Condorcet-consistent and monotone rule. Unfortunately, solving this LP is highly intractable as the number of variables and constraints grows exponentially in the number of teams. For small values of n, k, and with sufficient computational power, one could compute such solutions, in the hope that one could use that rule in order to construct approximately optimal ones for larger n. Consider the following simple procedure to scale a k-SNM-α rule r_n for n teams to a k-SNM-α' rule $r_{n'}$ for $n' > n$. First, increase n' to the nearest multiple of n, $n' := nM$ of n. Partition the teams into n groups of equal size. Within each group, compute the top-cycle on the induced sub-tournament and select a team from it uniformly at random as a finalist. This will reduce the number of teams to exactly n finalists. Run r_n on the n finalists and declare that rule's winner as the overall winner. The result of this section bounds the value of α' for $r_{n'}$.

Theorem 3. *If there exists a top-cycle consistent, and k-SNM-α rule r for n teams, then there exists a top-cycle consistent and k-SNM-α' rule r' for $n' > n$ teams where*

$$\alpha' \leq \alpha\Big(1 - \frac{(k-1)^2}{n}\Big) + \frac{(k-1)^2}{n}.$$

We now outline the proof of Theorem 3 and defer its proof to the full version of this paper. It is easy to show that r' is top-cycle consistent. The derivation of the upper bound on α' is more complicated but it is based on a fairly simple idea. Let S be a set of colluding teams in a tournament T on a set of n' teams. If a permutation π on a set of n' teams is chosen uniformly at random, then either every group in the partition contains at most one team from S or there exists a group containing at least two teams from S. In the former case, we use the fact that fixing matches inside S does not change a team's chances of surviving the group. Then we use the assumption that r is k-SNM-α to conclude that S can increase the probability to win the tournament T by at most α by fixing matches inside S. In the latter case, we simply assume that they can increase the probability by 1. The latter case happens with probability $(k-1)^2/n$, which finishes the proof. An explicit consequence of this result is the following. If there exists a top-cycle consistent 3-SNM-2/5 rule for $n = 25$ teams (which would be the best possible as per [18]), then there exists a 3-SNM-α rule for all $n \geq 25$ for some $\alpha < 1/2$.

6 Conclusion and Future Directions

This paper extends our knowledge of non-manipulable tournament rules in several ways. First, we generalize RSEB from [18] into RdSEB. This rule, at every node, picks a Condorcet-winner if one exists among its children and otherwise chooses a team uniformly at random. We show that for $k \leq d$ this rule is k-SNM-$\alpha_{d,k}$ for some $\alpha_{d,k}$ bounded away from 1, providing the first explicit family of rules that are bounded away from 1 for any n. We suspect that a more careful analysis for small values of d, k might yield better rules but conjecture, however, that $\alpha_{d,k} > 1/2$ for all $k \leq d$.

We present a new rule, the SignificantOnlyRule, which is monotone, Condorcet -consistent, 3-SNM-$1/2$ and 2-SNM-$1/3$ (which is best possible). The rule identifies a small set of teams as significant, awards them substantial probability mass and distributes it uniformly among non-significant teams. The motivation for the rule is that tournaments with a Condorcet-winner and tournaments which are far from having a Condorcet-winner are easy to resolve. We find a way to resolve the intermediate tournaments in a way to avoid substantial gains from manipulation.

Finally we propose a way of reducing the problem of finding good rules for large values of n to the problem of finding good rules for small values of n. Our result implies that if there exists a 3-SNM-$2/5$ just for $n = 25$, then there exist 3-SNM-α rules for $n \geq 25$ and $\alpha < 1/2$, which would directly improve on the state of the art results.

Acknowledgement. This research is part of a project that has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 823748, and while D. M. and J. S. were participants in the DIMACS REU program at Rutgers University, supported by NSF grant CNS-2150186.

References

1. Altman, A., Kleinberg, R.: Nonmanipulable randomized tournament selections. In: Proceedings of the National Conference on Artificial Intelligence, vol. 2, pp. 686–690 (2010)
2. Bartholdi, J.J., Tovey, C.A., Trick, M.A.: How hard is it to control an election? Math. Comput. Model. **16**(8), 27–40 (1992). https://doi.org/10.1016/0895-7177(92)90085-Y. http://www.sciencedirect.com/science/article/pii/089571779290085Y
3. Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D.: Handbook of Computational Social Choice. Cambridge University Press, Cambridge (2016)
4. Copeland, A.: A 'reasonable' social welfare function. Seminar on Mathematics in Social Sciences (1951)
5. Csato, L.: 2018 FIFA World Cup qualification can be manipulated (2017)
6. Dale, E., Fielding, J., Ramakrishnan, H., Sathyanarayanan, S., Weinberg, S.M.: Approximately strategyproof tournament rules with multiple prizes. In: Proceedings of the 23rd ACM Conference on Economics and Computation, EC 2022, pp. 1082–1100. Association for Computing Machinery, New York (2022). https://doi.org/10.1145/3490486.3538242
7. Dinev, A., Weinberg, S.M.: Tight bounds on 3-team manipulations in randomized death match. In: Hansen, K.A., Liu, T.X., Malekian, A. (eds.) WINE 2022. LNCS, vol. 13778, pp. 273–291. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-22832-2_16
8. Ding, K., Weinberg, S.M.: Approximately strategyproof tournament rules in the probabilistic setting. In: 12th Innovations in Theoretical Computer Science Conference, LIPIcs. Leibniz International Proceedings in Informatics, vol. 185, p. 20. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2021). Art. No. 14
9. Dutta, B.: Covering sets and a new Condorcet choice correspondence. J. Econ. Theory **44**(1), 63–80 (1988). https://doi.org/10.1016/0022-0531(88)90096-8. http://www.sciencedirect.com/science/article/pii/0022053188900968
10. Fishburn, P.C.: Condorcet social choice functions. SIAM J. Appl. Math. **33**(3), 469–489 (1977). https://doi.org/10.1137/0133030
11. Gibbard, A.: Manipulation of voting schemes: a general result. Econometrica **41**(4), 587–601 (1973)

12. Kim, M.P., Suksompong, W., Vassilevska Williams, V.: Who can win a single-elimination tournament? In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, Phoenix, Arizona, USA, 12–17 February 2016, pp. 516–522 (2016). http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12194
13. Kim, M.P., Vassilevska Williams, V.: Fixing tournaments for kings, chokers, and more. In: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, 25–31 July 2015, pp. 561–567 (2015). http://ijcai.org/Abstract/15/085
14. Laslier, J.F.: Tournament Solutions and Majority Voting, vol. 7. Springer, Heidelberg (1997)
15. Maurer, S.B.: The king chicken theorems. Math. Mag. **53**(2), 67–80 (1980). http://www.jstor.org/stable/2689952
16. Miller, N.R.: A new solution set for tournaments and majority voting: further graph-theoretical approaches to the theory of voting. Am. J. Polit. Sci. **24**(1), 68–96 (1980). http://www.jstor.org/stable/2110925
17. Moulin, H.: Choosing from a tournament. Soc. Choice Welf. **3**(4), 271–291 (1986). http://www.jstor.org/stable/41105842
18. Schneider, J., Schvartzman, A., Weinberg, S.M.: Condorcet-consistent and approximately strategyproof tournament rules. In: 8th Innovations in Theoretical Computer Science Conference, LIPIcs. Leibniz International Proceedings in Informatics, vol. 67, p. 20. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2017). Art. No. 35
19. Schvartzman, A., Weinberg, S.M., Zlatin, E., Zuo, A.: Approximately strategyproof tournament rules: on large manipulating sets and cover-consistence. In: 11th Innovations in Theoretical Computer Science Conference, LIPIcs. Leibniz International Proceedings in Informatics, vol. 151, p. 25. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2020). Art. No. 3
20. Stanton, I., Vassilevska Williams, V.: Rigging tournament brackets for weaker players. In: IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, 16–22 July 2011, pp. 357–364 (2011). https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-069
21. Suksompong, W.: Tournaments in computational social choice: recent developments. In: Zhou, Z.H. (ed.) Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-2021, pp. 4611–4618. International Joint Conferences on Artificial Intelligence Organization (2021). https://doi.org/10.24963/ijcai.2021/626. Survey Track
22. Vassilevska Williams, V.: Fixing a tournament. In: Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, 11–15 July 2010 (2010). http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1726
23. Vu, T., Altman, A., Shoham, Y.: On the complexity of schedule control problems for knockout tournaments. In: 8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009), Budapest, Hungary, 10–15 May 2009, vol. 1, pp. 225–232 (2009). https://doi.org/10.1145/1558013.1558044