Approximately Strategyproof Tournament Rules with Multiple Prizes

EMILY DALE, Princeton University
JESSICA FIELDING, Princeton University
HARI RAMAKRISHNAN, Princeton University
SACHETH SATHYANARAYANAN, Princeton University
S. MATTHEW WEINBERG, Princeton University

We consider the manipulability of tournament rules which take the results of $\binom{n}{2}$ pairwise matches and select a *ranking* over the teams. Prior work designs simple tournament rules such that no pair of teams can manipulate the outcome of their match to improve their probability of *being ranked first* by more than $1/3$, and this is the best possible among any Condorcet-consistent tournament rule (which selects an undefeated team whenever one exists) [15, 16]. We initiate the consideration of teams who may manipulate their match *to improve their ranking* (not necessarily to reach first).

Specifically, teams compete for a monetary prize, and the i^{th} ranked team takes home p_i in prize money ($p_i \geq p_{i+1}$ for all i). In this language, prior work designs tournament rules such that no pair of teams can manipulate the outcome of their match to improve their (collective) expected prize money by more than $1/3$, *when the price vector is* $\langle 1, 0, \ldots, 0\rangle$. We design a simple tournament rule (that we call Nested Randomized King of the Hill) such that: a) no pair of teams can improve their collective expected prize money by more than $1/3$ *for any prize vector in* $[0,1]^n$, and b) no set of *any* teams can gain *any* prize money for the uniform prize vector with $p_i := \frac{n-i}{n-1}$.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**.

Additional Key Words and Phrases: Tournament Rules, Incentive Compatibility, Quick Sort, Voting Theory

ACM Reference Format:
Emily Dale, Jessica Fielding, Hari Ramakrishnan, Sacheth Sathyanarayanan, and S. Matthew Weinberg. 2022. Approximately Strategyproof Tournament Rules with Multiple Prizes. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22), July 11–15, 2022, Boulder, CO, USA.* ACM, New York, NY, USA, 19 pages. https://doi.org/10.1145/3490486.3538242

EC '22, July 11–15, 2022, Boulder, CO, USA

ACM ISBN 978-1-4503-9150-4/22/07.
https://doi.org/10.1145/3490486.3538242

1 INTRODUCTION

A tournament consists of n teams competing to win a championship via pairwise matches, and a tournament rule selects a winner (possibly using randomization) based on the results. The design of tournament rules has received significant attention within Social Choice Theory [3, 7, 8, 11–13, 17] (see [4] for a survey). Tournament rules have also become an object of study within the TCS community over the past decade [1, 2, 6, 9, 10, 15, 16, 18, 19]. In particular, works such as [1, 2, 6, 15, 16] design *fair* tournaments that are as *minimally manipulable as possible* subject to a precise fairness condition. Our work further contributes to this direction.

In particular, all prior works in this direction consider a set of colluding teams who fix the outcomes of their matches in order to *maximize the probability that one of them wins the tournament.* This model is natural for theoretical study, and accurately captures settings where there is truly only one winner and no distinction between the non-winners (e.g. a qualification tournament where only one team can advance, or a tournament rule as a proxy for a voting rule). However, most modern tournaments offer rewards for teams beyond the winner. For example, the top four teams in the English Premier League qualify for the Champions League, and the bottom three are relegated to a less competitive league. While it is of course more prestigious to win the league outright, there is a significant difference between finishing fourth versus fifth, and seventeenth versus eighteenth. For example, one could imagine assigning a monetary value to participation in the Champions League, and to staying in the Premier League, and therefore the tournament rule implicitly assigns a monetary prize based on the final ranking. Some tournaments, such as the League of Legends Championship Series, directly award a monetary prize to teams based on their final ranking. We initiate the study of manipulation of tournament rules where teams care about their final ranking, and not just whether they win the tournament.

To have a running example in mind, consider an eSports tournament where prizes are awarded to top-ranked teams (either implicitly via qualification to a more prestigious tournament, or explicitly via a decreasing sequence of monetary prizes). Specifically, say that the i^{th}-ranked team takes home p_i in prize money. It is not uncommon for one sponsor to back multiple teams in the same tournament, and this sponsor may seek to maximize their own expected prize winnings (independently of exactly which teams the winnings come from). In particular, two teams of the same sponsor may be incentivized to fix the outcome of the match between them in order to maximize their collective winnings.[1] A tournament designer in turn may hope for a tournament rule that minimizes the maximum possible gains that are possible, while maintaining some formal guarantee that the produced ranking is fair.

1.1 Our Results

Prior work seeks winner-selection rules that are both *Condorcet-Consistent* (if there is an undefeated team, that team wins with probability 1)1 and minimally manipulable subject to this [1, 2, 6, 15, 16]. More specifically, they define $\alpha_k(r)$ to be the minimum α such that no set S of k teams can ever manipulate the outcomes of matches within S to improve the probability that r selects a winner in S by more than α. These works design several tournament rules with $\alpha_2(r) = 1/3$, and also prove that $\alpha_2(r) \geq 1/3$ for any Condorcet-Consistent r [6, 15, 16].

We instead consider tournament rules that output a complete ranking of the teams, and the i^{th}-ranked team earns p_i prize money. We now define $\alpha_k^{\vec{p}}(r)$ to be the minimum α such that no set S of k teams can ever manipulate the outcome of matches within S to improve their expected prize

[1] For example, perhaps the top X teams advance to a more prestigious tournament, and the stronger sponsored team feels pretty confident about their chances of advancing. They may choose to take it easy on the weaker sponsored team to improve the probability that both advance.

winnings under r with prize vector $\vec{p}$ by more than α (see Section 2 for a formal definition). In this language, prior works study $\alpha_2^{\langle 1,0,\ldots,0\rangle}(r)$, but not $\alpha_2^{\vec{p}}(r)$ for any other $\vec{p}$. Our main results designs a new rule that we call Nested Randomized King of the Hill (NRKotH), and provide tight bounds on its manipulability over arbitrary prize vectors, and a particular Borda prize vector.

Main Result 1 (see Theorem 3.1). For all nonincreasing prize vectors in $[0,1]^n$, no pair of teams can ever manipulate the match between them to improve their expected prize winnings under Nested Randomized King of the Hill by more than 1/3. This is the best possible guarantee among all Condorcet-Consistent tournament rules, and even just for the prize vector $\langle 1,0,\ldots,0\rangle$.

Main Result 2 (see Theorem 4.1) For a tournament on n participants, we define the Borda prize vector $\vec{p} = (p_1, p_2, \ldots, p_n) \in \mathbb{R}^n$ as the vector having $p_i = \frac{n-i}{n-1}$. Then, we have that $\alpha_k^{\vec{p}}(\text{NRKotH}) = 0$ for all $k \leq n$. That is, no set of teams of any size can gain anything by manipulating NRKotH under the Borda prize vector. Theorem 4.2 further gives an upper bound on the manipulability of NRKotH for $\vec{p}$ as a function of how far $\vec{p}$ is from Borda.

Nested Randomized King of the Hill is quite simple: (a) Pick a uniformly random team u, and fix their ranking to be below all teams who beat u, and above all teams that u beats, then (b) recurse on the set of teams that beat u, and separately on the set of teams that u beats, to determine the relative ranking of teams within these two sets. We also quickly note that, while NRKotH is competitive even with the best Condorcet-Consistent tournament rule, and even with the best guarantee achievable just on $\langle 1,0,\ldots,0\rangle$, it achieves a significantly stronger fairness guarantee. Specifically, NRKotH is *Cover-Consistent*: if team u beats team v *and every team that* v *beats*, then team u is ranked ahead of team v with probability 1 (see Section 2 for a formal definition). We now elaborate on the context for each of our main results.

For Theorem 3.1, we note that NRKotH is indeed inspired by Randomized King of the Hill, designed in [16].[2] Much of our analysis is inspired by [16] as well. Still, we wish to emphasize that there are many natural extensions from winner-selection rules to full-ranking rules, so the main technical contribution of Theorem 3.1 is nailing down the right one.[3] It is also notable that a simple Cover-Consistent tournament rule has the optimal manipulability guarantee against all prize vectors in $[0,1]^n$, even among Condorcet-Consistent rules, and its analysis is fairly clean.

Theorem 4.1 is the first result of its kind in two ways, and the proof approach is quite distinct from prior work. First, Theorem 4.1 is the only *non*-manipulability result in this model, at all. Indeed, [1, 2] establish that Condorcet-consistent winner-selection rules are all manipulable, even by pairs of teams, [15] clarifies that they are in fact 1/3-manipulable by pairs of teams, and all recent work designs various tournament rules that match this guarantee. Theorem 4.1 shows that by considering an alternative, and still natural, prize vector, non-manipulable rules are now possible. Second, there is scarce prior work considering manipulations from sets of $k > 2$ teams, and what little work exists is *significantly* more technically involved than the $k = 2$ case. For example, all that is currently known for any $k > 2$ is: a) for any Condorcet-consistent r, and all k, $\alpha_k^{\langle 1,0,\ldots,0\rangle}(r) \geq \frac{k-1}{2k-1}$ [15], b) this bound can be slightly improved when $k \geq 939$ [16], and c) a rule r exists with $\alpha_k^{\langle 1,0,\ldots,0\rangle} \leq 2/3$ [16]. Results b) and c) are quite technical, and in particular are significantly more involved than the

[2]Randomized King of the Hill is a winner-selection rule. The winner it selects is the first-ranked team under NRKotH.

[3]For example, perhaps the most natural extension from a winner-selection rule to a full-ranking rule would be to run the winner-selection rule n times, each time removing the most recent winner and ranking them next. It is unclear whether this extension of RKotH achieves the same guarantees as NRKotH. If so, a proof would likely be significantly more involved.

simple arguments used to prove Theorem 4.1. Theorem 4.1 therefore also shows that by considering an alternative prize vector, tractable analysis for manipulating sets of size $k > 2$ is now possible.

Beyond our main results, Section 5 also draws a formal connection between tournament ranking rules and sorting algorithms. For example, NRKotH seems similar to QuickSort, and Section 5 makes this formal. We also consider rules based on other sorting algorithms. These rules don't retain the same nice properties as NRKotH, further clarifying that there is something 'right' about NRKotH, even in comparison to other structured tournament ranking rules.

With this discussion in mind, the main contributions of our work are (a) initiating the study of manipulability of full ranking rules, rather than just winner-selection rules, (b) identifying NRKotH as a simple, natural rule that is optimally-manipulable over worst-case prize vectors, and *non*-manipulable over the Borda prize vector, and (c) establishing that the Borda prize vector, rather than the all-or-nothing prize vector, enables non-manipulable rules and tractable analysis for manipulating sets of size $k > 2$.

1.2 Related Work

The model we consider was first posed in [2]. Early follow-up work of [1] designed tournaments with $\alpha_2(r) = 0$, but not Condorcet-Consistence, with the goal of maximizing the probability that a Condorcet winner is guaranteed to be selected.[4]

Observing that the space of tournaments with $\alpha_2(r) = 0$ is quite restrictive, more recent work of [15] proposed to instead design tournament rules that are Condorcet-Consistent and as minimally-manipulable as possible. Their main result is a simple tournament rule (that they call Randomized Single Elimination Bracket) with $\alpha_2(r) = 1/3$, the best possible among Condorcet-Consistent rules. Follow-up works in this direction design alternate rules, such as Randomized King of the Hill [16] and Randomized Death Match [6] that achieve this same guarantee.

The simplest comparison between Theorem 3.1 and these recent works is as follows: [15] proposes the model we study, and designs the first tournament rule with $\alpha_2(r) = 1/3$. [16] considers colluding sets of more than two teams, and refutes a conjecture of [15] regarding this case. [6] considers a beyond worst-case model, and aims to understand whether improved manipulability guarantees are possible when the outcome of matches are 'close to random'. Each of these works take the study of manipulability of tournament rules in one new direction, and design one new tournament rule . Our work is the first to consider the manipulability of tournament rules that produce a full ranking, and designs the new Nested Randomized King of the Hill.

For $k > 2$, we have previously overviewed prior work. A simple construction in [15] establishes that $\alpha_k^{\langle 1,0,\ldots,0\rangle}(r) \geq \frac{k-1}{2k-1}$ for all Condorcet-consistent r. A significantly more involved construction of [16] improves this to show that $\alpha_k^{\langle 1,0,\ldots,0\rangle}(r) \geq 1/2$ for all $k \geq 939$. [16] also designs a Condorcet-consistent tournament rule with $\alpha_k^{\langle 1,0,\ldots,0\rangle}(r) = 2/3$ for all k (but this rule is not even monotone — losing a match can improve a single team's probability of winning). In comparison to these works, Theorem 4.1 follows from significantly cleaner arguments, and establishes optimal bounds for all k, for the Borda prize vector. As previously mentioned, there is no prior work establishing a non-manipulable Condorcet-consistent tournament rule — Theorem 4.1 is the first of this kind.

Beyond this work, the study of tournament rules within TCS focuses specifically on single-elimination brackets, and a designer manipulating seeding in order to get a certain team to win [9, 10, 18, 19]. Other recent works consider strategic manipulation in different particular tournament rules, including the World Cup qualifying procedure [5, 14]. Aside from the thematic relation, these works are technically disjoint from ours.

[4]Example rules to have in mind from their work is one that selects a uniformly random team as a winner, or one that selects two uniformly random teams to play a match and selects the winner as the tournament winner.

2 PRELIMINARIES

We introduce notation consistent with prior work [1, 6, 15, 16]. We update terminology slightly to reflect that our tournament rules output a ranking, rather than a single winner.

Definition 2.1 (Tournament). *A (round robin) tournament T on n teams is a complete, directed graph on n vertices whose edges denote the outcome of a match between two teams. Team i beats team j if the edge between them points from i to j.*

Definition 2.2 (Tournament Ranking Rule). *A tournament ranking rule r is a function that maps tournaments T to a distribution over rankings σ (where $\sigma(i)$ denotes the ranking of i), where $r_\sigma(T) := \Pr[r(T) = \sigma]$ denotes the probability that ranking σ is output on tournament T under rule r. We use the notation $r_{i,j}(T) := \sum_{\sigma|\sigma(i)=j} r_\sigma(T)$ to denote the probability that team i is ranked j^{th}.*

Definition 2.3 (Prize Vector). *A prize vector is a vector $\vec{p} \in \mathbb{R}^n$ such that $p_j \geq p_{j+1}$ for all j. The semantic meaning is that the team ranked j^{th} receives p_j in prize money. For a prize vector $\vec{p}$, team i, and tournament ranking rule r, we'll use the notation $r_i^{\vec{p}}(T) := \sum_{j=1}^n r_{i,j}(T) \cdot p_j$ to denote the expected prize money earned by team i under rule r on tournament T with prize vector $\vec{p}$. We'll also use the notation $r_S^{\vec{p}}(T) := \sum_{i \in S} r_i^{\vec{p}}(T)$ to denote the collective prize money of teams in S.*

Whenever we use the term 'prize vector' to refer to a vector $\vec{p}$, this implies that $p_j \geq p_{j+1}$ for all j. For example, when we say "all prize vectors in $[0,1]^n$," this refers to "all vectors $\vec{p} \in [0,1]^n$ that satisfy $p_j \geq p_{j+1}$ for all j."

Like prior work, we are interested in tournament rules which satisfy basic notions of fairness. Prior work mostly considers rules which are *Condorcet-Consistent*: whenever a team is undefeated, that team wins (in our language, is ranked first) with probability one.

Definition 2.4 (Condorcet-Consistent). *Team i is a Condorcet winner of a tournament T if i beats every other team (under T). A tournament ranking rule r is Condorcet-consistent if for every tournament T with a Condorcet winner i, $r_{i,1}(T) = 1$ (whenever T has a Condorcet winner, that team is ranked first with probability 1).*

For winner-selection rules, Condorcet-Consistence is a minimal, but reasonable notion of fairness. Like prior work, the quality of our designed tournaments will compete with the best Condorcet-Consistent tournament ranking rule. However, in order to deem these rules desirable, they should satisfy stronger properties.[5] [16] considers the stronger notion of cover-consistence. We extend their definition to ranking rules, and will design rules that satisfy this stronger property.

Definition 2.5 (Cover-Consistent). *Team i* covers *team j in tournament T if i beats j, and i beats every team that j beats. A tournament ranking rule is* Cover-Consistent *if for all T, and all i, j such that i covers j in T, and all σ such that $r_\sigma(T) > 0$, $\sigma(i) < \sigma(j)$. That is, whenever i covers j in T, rule r applied to T should output a ranking where i is ahead of j with probability 1.*

Finally, we are interested in how manipulable tournament ranking rules are.

Definition 2.6 (S-adjacent). *Two tournaments T, T' are S-*adjacent *if they are identical except for matches between two teams in S.*

Definition 2.7 (Manipulating a Tournament). *For a set S of teams, tournament T, tournament rule r, and prize vector $\vec{p}$, we define $\alpha_S^{\vec{p}}(r, T)$ to be the maximum prize money that S can possibly gain in*

[5] For example: Consider a rule that ranks a Condorcet winner first, if it exists, and ranks all remaining teams uniformly at random. Then a Condorcet loser could wind up ranked ahead of a team that beats everyone except the Condorcet winner.

expectation by manipulating T to an S-adjacent T'. That is: $\alpha_S^{\vec{p}}(r,T) := \max_{T':T' \text{ is } S\text{-adjacent to } T}\{r_S^{\vec{p}}(T') - r_S^{\vec{p}}(T)\}$.

We further define $\alpha_k^{\vec{p}}(r) := \max_{T,S:|S|\leq k}\{\alpha_S^{\vec{p}}(r,T)\}$ *to be the maximum prize money under* $\vec{p}$ *that any set of* $\leq k$ *teams can gain in* r *on any underlying tournament. For a class of prize vectors* $\mathcal{P}$ *we define* $\alpha_k^{\mathcal{P}}(r) := \sup_{\vec{p}\in\mathcal{P}}\{\alpha_k^{\vec{p}}(r)\}$ *to be the maximum prize money that any set of* $\leq k$ *teams can gain in* r *under any* $\vec{p} \in \mathcal{P}$.

Finally, we define $\alpha_k^{\mathcal{P}} := \inf_{\text{Condorcet consistent } r}\{\alpha_k^{\mathcal{P}}(r)\}$. *That is,* $\alpha_k^{\mathcal{P}}$ *is the best bound on manipulability achievable by a Condorcet-Consistent tournament ranking rule against collusions of* k *teams that holds for all prize vectors in* $\mathcal{P}$.

In this notation, prior works analyze $\alpha_2^{\mathcal{P}}$ *when* $\mathcal{P}$ *contains the single prize vector* $\langle 1, 0, \ldots, 0\rangle$. The key difference in our work is that we will study $\alpha_2^{\mathcal{P}}$ *when* $\mathcal{P}$ *contains all prize vectors in* $[0,1]^n$, and *when* $\mathcal{P}$ *contains the single prize vector with* $p_i := \frac{n-i}{n-1}$.

3 MAIN RESULT I: WORST-CASE PRIZE VECTORS

In this section we state and prove our first main result, that Nested Randomized King of the Hill is the least manipulable tournament rule for prize vectors in $[0,1]^n$. We begin with a formal definition of NRKotH, and then state and prove the result.

Definition 3.1 (Nested Randomized King of the Hill). *The tournament ranking rule Nested Randomized King of the Hill (NRKotH) proceeds as follows, when given as input a tournament T on n teams:*

1. *If $n = 0$, return an empty ordering. Else, continue.*
2. *Pick a team, u, uniformly at random. Call u the pivot.*
3. *Let B denote the teams that beat u, and L denote the teams that lose to u.*
4. *Run NRKotH on B and L, and call the outputs σ_B and σ_L respectively.*
5. *For all teams $b \in B$, set $\sigma(b) := \sigma_B(b)$.*
6. *Set $\sigma(u) := |B| + 1$.*
7. *For all teams $\ell \in L$, set $\sigma(\ell) := \sigma_L(\ell) + |B| + 1$.*
8. *Output σ.*

That is, NRKotH picks a uniformly random team, u. All teams that beat u are ranked above u, and all teams that lose to u are ranked below u. With each set of teams, NRKotH is called recursively to determine their relative ranking. We state our main theorem below.

Theorem 3.1. *Let $\mathcal{P}$ denote the set of all prize vectors in $[0,1]^n$. Then $\alpha_2^{\mathcal{P}}(\text{NRKotH}) = 1/3 = \alpha_2^{\mathcal{P}}$. That is, for any prize vector in $[0,1]^n$, and any underlying tournament T, no two teams can manipulate their match to gain expected prize money more than $1/3$. Moreover, this is the best possible guarantee of any Condorcet-Consistent tournament ranking rule.*

Before proving Theorem 3.1, we establish that NRKotH is Cover-Consistent,[6] and prove some other basic facts about NRKotH. Recall that Cover-Consistence is a significantly stronger fairness guarantee than Condorcet-Consistence. Therefore, Theorem 3.1 establishes that there is no loss in worst-case manipulability due to this stronger fairness guarantee.

[6]The proof is similar to that of [16, Lemma 6.1], which establishes that no covered team can win in the related tournament rule Randomized King of the Hill.

3.1 Properties of NRKotH

Proposition 3.2. *NRKotH is Cover-Consistent.*

Proof. Consider any team v that covers team w, and an execution of NRKotH containing both v and w. Consider the team u selected as pivot:

- If $u = v$, then because v covers w, $w \in L$, and therefore u finishes ahead of w. Similarly, if $u = w$, $v \in B$, and v finishes ahead of w.
- If v beats u and u beats w, then v finishes ahead of w.
- It is not possible to have w beat u and u beat v, because v covers w.
- If both v, w beat u, or both v, w lose to u, then their relative ranking is determined by a recursive call. However, that recursive call must eventually terminate in one of the first two cases.

Therefore, v finishes ahead of w whenever v covers w. □

Next, we provide an equivalent view of NRKotH that will be helpful in analysis. Essentially, this view just specifies a precise order in which to execute the recursive calls, and pre-determines the selected pivots.

Definition 3.2 (Current Group). *During any execution of NRKotH, every team u is currently in a* group. *This refers to the set of teams that will be present in the next recursive call containing u. We will use the notation $G(u)$ to refer to the group containing u, which updates each time a new pivot is selected.*

For example, initially every team is in the same group. After one round of NRKotH, there are three (possibly empty) groups: the teams that beat the pivot, the pivot, and the teams that lose to the pivot. In general, after one round of NRKotH is run on a group, the pivot is now in a group by themselves, teams that beat the pivot form a group, and teams that lose to the pivot form a group.

It will be helpful to couple outcomes of NRKotH on different tournaments via the pivots selected. Specifically, if $\tau(\cdot)$ denotes a permutation on $[n]$ (where $\tau(i)$ is the i^{th} team in the list), we wish to couple executions so that within every recursive call on a group, *the team that is earliest in τ is selected as the pivot.* The definition below captures this concept formally.

Definition 3.3 (τ-ordered implementation of NRKotH). *The τ-ordered implementation of NRKotH proceeds as follows. For $i = 1$ to n: Choose the group $G(\tau(i))$ to process and pick $\tau(i)$ as the pivot.*

Observation 3.1. *Drawing τ uniformly at random, and then running the τ-ordered implementation of NRKotH produces a ranking that is identically distributed to NRKotH.*

Proof. This follows as: (a) the order in which groups are processed does not affect the outcome of NRKotH and (b) the pivots chosen for each group are uniformly random among teams in that group, when τ is uniformly random. □

This view of NRKotH will be helpful in proving our main result.

3.2 Proof of Theorem 3.1

Now, we prove Theorem 3.1. We begin with the following lemma, which establishes some cases where u and v cannot profit by manipulating their match.

Lemma 3.3. *If u and v beat exactly the same set of teams in T, then $\alpha^{\vec{p}}_{\{u,v\}}(NRKotH, T) = 0$ for all $\vec{p}$. That is, u and v cannot increase their expected prize money in NRKotH in any tournament where they beat exactly the same set of teams, for any prize vector.*

Lemma 3.3 follows from similar reasoning as [6, Lemma 5.1], as NRKotH is anonymous. We include a proof below for completeness, as notation needs to be updated for ranking rules.

Proof. Observe first that every round of NRKotH will either select u or v as pivot, or it will keep u and v in the same group. In the latter case, this is not impacted by the outcome of the match between u and v (because u and v do not play each other). So we only need to understand what happens in the former case.

Here, we claim that the prize money won is identically distributed, independently of whether u beats v or vice versa. Indeed, let w denote the team that wins the (u, v) matchup, and ℓ denote the other team. Then if w is pivot, ℓ enters L, along with the teams that w beats. Importantly, observe that whether $w = u$ or $w = v$, $L \setminus \{\ell\}$ is exactly the same. Moreover, the final ranking of ℓ depends only on its matches with teams in L, which are also identical whether $\ell = u$ or $\ell = v$. Therefore, when w is selected, the rank of w and rank of ℓ are distributed the same no matter which of $\{u, v\}$ is w and which is ℓ. Identical reasoning holds if ℓ is selected (w enters B, the teams in B are the same, and the matches between w and B are also the same, regardless of whether $w = u$ or $w = v$).

Finally, observe that w and ℓ are equally likely to be pivot. Putting everything together, this means that we can write the distribution of prize money won by u and v together by writing the distribution of prize money won by w and ℓ together, and observe that this is independent of whether $w = u$ or $w = v$. □

We now establish a second case where u and v cannot profitably manipulate their match.

Lemma 3.4. *Let T be a tournament where u beats v, and let T' be the $\{u, v\}$-adjacent tournament that flips the (u, v) match. Then* conditioned on v being selected as the first pivot*, the expected prize money won by u and v is greater in T than T' under NRKotH.*

Proof. Let B_v denote the teams that beat v in T'. Observe that in T, if v is selected as the first pivot, then $\sigma(v) = |B_v| + 2$, and $\sigma(u) \leq |B_v| + 1$ (because u beats v in T). Therefore, the minimum possible prize money u and v could attain is $p_{|B_v|+2} + p_{|B_v|+1}$.

In T', if v is selected as pivot, then $\sigma(v) = |B_v| + 1$, and $\sigma(u) \geq |B_v| + 2$ (because v beats u in T'). Therefore, the maximum possible prize money that u and v could attain in T' is $p_{|B_v|+1} + p_{|B_v|+2}$.

Putting both together, we conclude that u and v have greater expected reward in T than in T'. □

Now, we can complete the proof of Theorem 3.1. The main idea is to show that u and v are unlikely to be able to positively impact their collective prize money (with probability at most 1/3), and that when they do, they gain at most 1 by doing so.

Proof of Theorem 3.1. For the rest of this proof, we will explicitly use the concept of a τ-ordered implementation, and randomly draw τ one step at a time. Specifically:

(1) Initalize τ to be null/empty.
(2) For $i = 0$ to $n - 1$.
   (a) Let X denote the teams that beat u and lose to v, or beat v and lose to u, *and* are in the same group as both u and v after the first i rounds of the τ-ordered imlpementation of NRKotH.[7]
   (b) With probability $(|X| + 2)/(n - i)$, decide that $\tau^{-1}(i + 1) \in X \cup \{u, v\}$.
      (i) Then, draw $\tau^{-1}(i + 1)$ uniformly at random from $X \cup \{u, v\}$.
   (c) Else (i.e., with probability $(n - i - |X| - 2)/(n - i)$), decide that $\tau^{-1}(i + 1) \notin X \cup \{u, v\}$.
      (i) Then, draw $\tau^{-1}(i + 1)$ uniformly at random from the remaining teams not in $X \cup \{u, v\}$.

[7] Observe that to be in the same group as u and v, teams must have not previously been chosen as pivot (as pivots are immediately placed in their own group).

Observe that the procedure above indeed draws a τ uniformly at random: at each step, it picks a uniformly random remaining team.

Now, let us consider the very first i such that step (b) is taken. Our claim breaks down into two cases:

Lemma 3.5. *Consider any run of the process for building τ such that step (b) is first invoked while $|X| = 0$. Then, the outcome of the match between u and v does not affect their collective prize money.*

Proof. Consider the first time that step (b) is invoked, and consider that $|X| = 0$. This means that the group containing u and v contains no teams that beat one but not the other. Therefore, by Lemma 3.3, the outcome of the match between u and v does not affect the distribution of rankings within their group. Also, it is clear that the outcome of the match between u and v has not previously been invoked in any prior iteration of NRKotH, as neither of them were pivot before this round (as step (b) is the only case where u or v could be pivot). □

Lemma 3.6. *Consider any run of the process for building τ such that step (b) is first invoked while $|X| > 0$. Then the probability that the outcome of the match between u and v positively affects their collective prize money is at most $1/3$.*

Proof. Consider two manipulating teams u and v, and say wlog that u beats v in the original tournament. Now consider the round at which step (b) is invoked while building τ.

- With probability $1/(|X| + 2)$, v is selected as pivot. By Lemma 3.4, manipulating the (u, v) match cannot possibly increase their expected prize winnings.
- With probability $|X|/(|X| + 2)$, a team that is currently in the same group as both u and v, and that beats one of them but not the other, is selected as pivot. When this happens, it guarantees that the (u, v) match is never played (because u and v are immediately split into different groups), and therefore its outcome can't possibly affect any prize winnings.
- With probability $1/(|X| + 2)$, u is selected as pivot. This is the only case where u and v can increase their joint prize winnings by manipulating their match. But as $|X| \geq 1$ by hypothesis, this case occurs with probability at most $1/3$. This completes the proof.

□

We've now established that u and v are unlikely to be able to profitably manipulate their match. Our last step is to upper bound their gains by doing so. A trivial upper bound is 2 (because both make at least 0 and at most 1 in all outcomes), but we'll need a slightly stronger bound of 1.

Lemma 3.7. *If the match between u and v is played, its outcome affects the collective prize winnings of u and v by at most $p_1 - p_n$.*

Proof. If the match between u and v is played, this means that one of them (wlog, say it is u) is pivot. Let r denote the resulting $\sigma(u)$ if u beats v.

Then if u beats v, their collective prize money is at least $p_r + p_n$. If v beats u, then $\sigma(u) = r + 1$, and so their collective prize money is at most $p_1 + p_{r+1}$. Because the prize vector is monotone, their difference is at most $p_1 - p_n$. □

Now with Lemmas 3.5 and 3.6, we can wrap up the proof of Theorem 3.1. Observe that the first time step (b) is invoked, we either have $|X| = 0$, or $|X| \geq 1$. In both cases, Lemmas 3.5 and 3.6 establish that the probability of selecting a pivot such that u and v can positively affect their collective prize winnings by manipulating their match is at most $1/3$. By Lemma 3.7, conditioned on this occurring, it impacts the collective winnings of u and v by at most 1. Therefore, the expected prize winnings that u and v can gain by manipulation is at most $1/3$. This proves that $\alpha_2^{\mathcal{P}}(\text{NRKotH}) \leq 1/3$.

To wrap-up, recall that [15, Theorem 3.1] proves that every Condorcet-Consistent tournament rule admits a tournament T such that some pair of teams can manipulate it and improve their joint probability of winning by at least 1/3. In our language, this establishes that even $\alpha_2^{\langle 1,0,\ldots,0\rangle} \geq 1/3$. This implies that $\alpha_2^{\mathcal{P}} \geq 1/3$ as well, and completes the proof. □

4 MAIN RESULT II: BORDA PRIZE VECTOR

In this section, we introduce the notion of *consistency under expectation* and show that NRKotH is consistent under expectation. We use this fact to prove that, under the Borda prize vector, no set of k teams can manipulate NRKotH, for any k. We also provide a bound on the manipulability of NRKotH under a prize vector ℓ_∞-close to the Borda prize vector. We first state our main results:

Definition 4.1 (Borda Prize Vector). *We define a Borda prize vector for a tournament T on n participants as the vector $\vec{p} = (p_1, p_2, \ldots, p_n) \in \mathbb{R}^n$ where $p_i = \frac{n-i}{n-1}$.*

Definition 4.2 (ε−close to Borda Prize Vectors). *We define the class of ε−close to Borda prize vectors $\mathcal{P}^\varepsilon$ to be the set of all prize vectors $\vec{p}$ with $p_i \in \left[\frac{n-i}{n-1} - \varepsilon, \frac{n-i}{n-1} + \varepsilon\right]$ for all i (i.e. every element is within ε of the corresponding element in the Borda prize vector).*

Theorem 4.1. *For the Borda prize vector $\vec{p} \in \mathbb{R}^n$, we have that $\alpha_k^{\vec{p}}(NRKotH) = 0$ for all $k \leq n$.*

Theorem 4.2. *For all k, $\alpha_k^{\mathcal{P}^\varepsilon}(NRKotH) \leq 2k\varepsilon$.*

While Theorem 4.2 implies Theorem 4.1, we separate the two proofs to provide the main ideas first in Theorem 4.1. The main workhorse in both proofs is a property of NRKotH that we term *consistency under expectation*. Intuitively, this property states that the expected rank of a team under NRKotH can be determined solely as a function of the number of teams it beats. The definitions below state this formally.

Definition 4.3. *For any individual team u and tournament T, we define $w_T(u)$ to be the set of teams that u defeats in tournament T.*

Definition 4.4. *For any team u and rule r, we define $\sigma_T^r(u)$ to be the random variable that is the ranking of team u under rule r, applied to tournament T.*

Definition 4.5 (Consistent Under Expectation). *A tournament rule r is* consistent under expectation *if for all n, all tournaments T on n teams, and all u:*

$$\mathbb{E}\left[\sigma_T^r(u)\right] = n - |w_T(u)|$$

We now begin the proofs of Theorems 4.1 and 4.2. The main idea behind both proofs is Theorem 4.3, which we prove first.

Theorem 4.3. *NRKotH is consistent under expectation.*

Proof. We prove this by strong induction on n. The base case of $n = 1$ is is easy to verify: for the single team u, it beats 0 teams and its expected rank is $1 = 1 - 0$.

Consider now some $n > 1$, and assume for inductive hypothesis that the theorem holds for all $n' < n$. Consider now any team u, and let's analyze the distribution of u's rank.

During the first iteration of NRKotH, let v be the pivot. We compute the expectation of $\sigma_T^{\text{NRKotH}}(u)$ by conditioning on whether $v = u$, v defeats u, or v loses to u. Note that since the pivot is selected uniformly at random, the probability that a given team is the pivot is $1/n$. We now observe the following:

- If $v = u$, then the rank of u is immediately set at $n - |w_T(u)|$ (because all $n - |w_T(u)| - 1$ teams that beat u are permanently ranked higher, and all $|w_T(u)|$ teams that lose to u are permanently ranked lower.
- If u defeats v, then u's rank is equal to its rank in the subtournament B_v of teams that beat v.
- If u loses to v, then u's final rank is equal to its rank in the subtournament L_v of teams that lose to v, plus the rank of v, which is $n - |w_T(v)|$.

Therefore, we can conclude the following:

$$\begin{aligned}\mathbb{E}\left[\sigma_T^{\text{NRKotH}}(u)\right] &= \frac{1}{n}\left(n - |w_T(u)| + \sum_{v \in w_T(u)} \mathbb{E}\left[\sigma_{B_v}^{\text{NRKotH}}(u)\right] + \sum_{v \neq u, v \notin w_T(u)} \left(n - |w_T(v)| + \mathbb{E}\left[\sigma_{L(v)}^{\text{NRKotH}}(u)\right]\right)\right) \\ &= \frac{1}{n}\left(n - |w_T(u)| + \sum_{v \in w_T(u)} n - |w_T(v)| - 1 - |w_{B_v}(u)|\right) \\ &\quad + \frac{1}{n}\left(\sum_{v \neq u, v \notin w_T(u)} \left(n - |w_T(v)| + |w_T(v)| - |w_{L_v}(u)|\right)\right)\end{aligned}$$

The first equality follows from the three previous bullets. The second follows from our inductive hypothesis, as both B_v and L_v are tournaments on $< n$ teams, and have $n - |w_T(v)| - 1$ and $|w_T(v)|$ teams in them, respectively. Let us now investigate some terms that appear in the sums.

First, let us analyze the term $|w_{B_v}(u)|$, for a $v \in w_T(u)$. Observe that $w_{B_v}(u)$ contains all teams that lose to u, but beat v. Put another way, it contains all teams that lose to u, removing v and those that lose to v. Therefore:

$$w_{B_v}(u) = w_T(u) \setminus \{w_T(v) \cup \{v\}\}.$$

Similarly, let us analyze the term $|w_{L_v}(u)|$, for a $v \notin w_T(u)$. Observe that $w_{L_v}(u)$ contains exactly the teams that lose to both v and u. Therefore:

$$w_{L_v}(u) = w_T(v) \cap w_T(u).$$

Substituting these back into our prior bounds, we now have:

$$\begin{aligned}
&\frac{1}{n}\left(n - |w_T(u)| + \sum_{v \in w_T(u)} n - |w_T(v)| - 1 - |w_{B_v}(u)|\right) \\
&\qquad + \frac{1}{n}\left(\sum_{v \neq u, v \notin w_T(u)} \left(n - |w_T(v)| + |w_T(v)| - |w_{L_v}(u)|\right)\right) \\
&= \frac{1}{n}\left(n - |w_T(u)| + \sum_{v \in w_T(u)} n - |w_T(v)| - |w_T(u) \setminus w_T(v)|\right) \\
&\qquad + \frac{1}{n}\left(\sum_{v \neq u, v \notin w_T(u)} n - |w_T(u) \cap w_T(v)|\right) \\
&= \frac{1}{n}\left(n - |w_T(u)| + \sum_{v \in w_T(u)} n - |w_T(u)| - |w_T(v) \setminus w_T(u)|\right) \\
&\qquad + \frac{1}{n}\left(\sum_{v \neq u, v \notin w_T(u)} n - |w_T(u)| + |w_T(u) \setminus w_T(v)|\right) \\
&= n - |w_T(u)| + \frac{1}{n}\left(\sum_{v \neq u, v \notin w_T(u)} |w_T(u) \setminus w_T(v)| - \sum_{v \in w_T(u)} |w_T(v) \setminus w_T(u)|\right).
\end{aligned}$$

The first equality follows from our substitutions. The second follows by rewriting $|w_T(u) \cup w_T(v)|$ and $|w_T(u) \cap w_T(v)|$. The third just groups the n terms $n - |w_T(u)|$ together. From here it suffices to prove that the term inside the parentheses is 0.

To this end, it is helpful to consider any edge $e = (x, y)$ and compare its contribution to each sum. Indeed:

$$\begin{aligned}
\sum_{v \neq u, v \notin w_T(u)} |w_T(u) \setminus w_T(v)| &= \sum_{x \neq u, y \notin \{x,u\}} \mathbb{I}[x \notin w_T(u) \wedge y \in w_T(u) \wedge y \notin w_T(x)] \\
\sum_{v \in w_T(u)} |w_T(v) \setminus w_T(u)| &= \sum_{y \neq u, x \notin \{y,u\}} \mathbb{I}[y \in w_T(u) \wedge x \in w_T(y) \wedge x \notin w_T(u)] \\
&= \sum_{y \neq u, x \notin \{y,u\}} \mathbb{I}[y \in w_T(u) \wedge y \notin w_T(x) \wedge x \notin w_T(u)]
\end{aligned}$$

The first two equality follows by considering the case when $x = v$, and which $y \in w_T(u) \setminus w_T(v)$. The second equality follows by considering the case where $y = v$, and which $x \in w_T(v) \setminus w_T(u)$. The third equality just observes that $x \in w_T(y) \leftrightarrow y \notin w_T(x)$. We now conclude that the two sums are equal, which completes the entire proof as well. □

4.1 Proof of Theorem 4.1

Theorem 4.3 is the main workhorse towards Theorem 4.1, as it draws a connection between the expected rank of a team and the number of matches it wins. The remaining steps are to show: (a) no set of k teams can manipulate the total number of matches they win, and (b) under the Borda prize vector, expected rank directly determines the expected prize winnings.

Lemma 4.4. *For any tournament T, any set S of teams, and any tournament T' that is S-adjacent to T: $\sum_{u\in S} w_{T'}(u) = \sum_{u\in S} w_T(u)$.*

Proof. Observe that we can write $w_T(u)$ as the number of matches u wins against teams in S, plus the matches they win against teams $\notin S$. Therefore:

$$\begin{aligned}\sum_{u\in S} |w_{T'}(u)| &= \sum_{u\in S} |w_{T'}(u) \cap S| + \sum_{u\in S} |w_{T'}(u) \setminus S| \\ &= \binom{|S|}{2} + \sum_{u\in S} |w_{T'}(u) \setminus S| \\ &= \sum_{u\in S} |w_T(u) \cap S| + \sum_{u\in S} |w_T(u) \setminus S| \\ &= \sum_{u\in S} |w_T(u)|\end{aligned}$$

The first and last equalities follow by breaking the teams that u beats into those in S and those not in S. The second inequality follows by observing that the total matches won by teams in S against teams in S must be exactly $\binom{|S|}{2}$ in both T and T'. The third equality follows as T, T' are S-adjacent.

□

Lemma 4.4 doesn't reference a tournament rule, nor a prize vector. Lemma 4.5 below connects the expected rank of a team to its expected prize money under the Borda prize vector.

Lemma 4.5. *Let $\vec{p}$ be the Borda prize vector. Then for all tournament rules r, all tournaments T, and all sets S of teams:*

$$r_S^{\vec{p}}(T) = \frac{|S| \cdot n - \mathbb{E}[\sum_{u\in S} \sigma_T^r(u)]}{n-1}.$$

Proof. The proof follows from the following calculations:

$$\begin{aligned}r_S^{\vec{p}}(T) &= \sum_{u\in S} \sum_{i=1}^{n} \Pr[\sigma_T^r(u) = i] \cdot p_i \\ &= \sum_{u\in S} \sum_{i=1}^{n} \Pr[\sigma_T^r(u) = i] \cdot \frac{n-i}{n-1} \\ &= \frac{|S| \cdot n - \mathbb{E}[\sum_{u\in S} \sigma_T^r(u)]}{n-1}.\end{aligned}$$

□

Now, we can wrap up the proof of Theorem 4.1. Essentially, there are three ingredients to the proof: Theorem 4.3 is specific to NRKotH, Lemma 4.4 is just a fact about tournaments, and Lemma 4.5 is specific to the Borda prize vector.

Proof of Theorem 4.1. Let T be any tournament, S be any set of teams, T' be any S-adjacent tournament to T, and let $\vec{p}$ denote the Borda prize vector. Then we have the following equalities:

$$\begin{aligned} r_S^{\vec{p}}(T) &= \frac{|S| \cdot n - \mathbb{E}[\sum_{u \in S} \sigma_T^r(u)]}{n-1} \\ &= \frac{|S| \cdot n - \sum_{u \in S}(n - |w_T(u)|)]}{n-1} \\ &= \frac{\sum_{u \in S} |w_T(u)|}{n-1} \\ &= \frac{\sum_{u \in S} |w_{T'}(u)|}{n-1} \\ &= \frac{|S| \cdot n - \sum_{u \in S}(n - |w_{T'}(u)|)]}{n-1} \\ &= \frac{|S| \cdot n - \mathbb{E}[\sum_{u \in S} \sigma_{T'}^r(u)]}{n-1} \\ &= r_S^{\vec{p}}(T') \end{aligned}$$

The first and final equalities follow from Lemma 4.5. The second and penultimate equalities follow from Theorem 4.3. The third and fifth equalities are basic algebra, and the fourth equality follows by Lemma 4.4. This concludes the proof. □

4.2 Proof of Theorem 4.2

The proof of Theorem 4.2 follows from Theorem 4.1 with one additional lemma. Essentially, for two prize vectors that are close in ℓ_∞ distance, the expected prize money won by a player in the same tournament under the different prize vectors cannot be far apart.

Lemma 4.6. *Let $|\vec{p} - \vec{q}|_\infty \leq \varepsilon$. Then for any tournament T, any tournament rule r, and any set S of teams:*

$$|r_S^{\vec{p}}(T) - r_S^{\vec{q}}(T)| \leq |S| \cdot \varepsilon.$$

Proof. The proof follows by coupling executions of r on T with $\vec{p}$ and $\vec{q}$ so that the same ranking is selected. Every team $u \in S$ is in the same position in both rankings (by definition). Because $|\vec{p} - \vec{q}|_\infty \leq \varepsilon$, the prize money u wins is within ε in both executions. Therefore, the sum of prize money won by S is within $|S| \cdot \varepsilon$ in the two executions. □

Proof of Theorem 4.2. The proof now follows immediately from Theorem 4.1, and two applications of Lemma 4.6. Let $\vec{q}$ denote the Borda prize vector, and r denote NRKotH. Then for any $\vec{p} \in \mathcal{P}^\varepsilon$, any tournament T, any set of teams S, and any S-adjacent T', we have:

$$r_S^{\vec{p}}(T') \leq r_S^{\vec{q}}(T') + |S| \cdot \varepsilon = r_S^{\vec{q}}(T) + |S| \cdot \varepsilon \leq r_S^{\vec{p}}(T) + 2 \cdot |S| \cdot \varepsilon.$$

□

4.3 An Additional Implication of Consistence Under Expectation

Before wrapping up this section, we briefly observe a stronger implication than our main result, that is implied by consistency under expectation. Specifically, our work (and all prior work in this model) studies the maximum gains from manipulation only by fixing matches between manipulating teams. The rules designed in our work (and all prior works) are also monotone: no single team can improve their own ranking by unilaterally throwing a match. However, prior works do not consider manipulations that both fix within-coalition matches *and* throw matches to non-colluding teams (that is, perhaps by team u throwing a match to team w, it increases the joint expected

prize-winnings of $\{u, v\}$). In general, it is not clear that techniques developed to study match-fixing can also apply to match-throwing. But, our notion of Consistence Under Expectation immediately enables extensions to this case. In fact, we can reuse most of our prior lemmas, and just need to update Lemma 4.4.

Lemma 4.7. *For any tournament T, and any set of teams S, let T' be S-adjacent to T. Further let T'' be such that if u beats v in T', but v beats u in T'', then $u \in S$ and $v \notin S$. Then: $\sum_{u \in S} |w_{T''}(u)| \leq \sum_{u \in S} |w_T(u)|$.*

Proof. We already know from Lemma 4.7 that $\sum_{u \in S} |w_{T'}(u)| = \sum_{u \in S} |w_T(u)|$. So we just need to show that $|w_{T''}(u)| \leq \sum_{u \in S} |w_{T'}(u)|$. This is fairly quick to see as no $u \in S$ can win additional matches in T'' compared to T', while some $u \in S$ can lose matches. □

Corollary 4.7.1. *For the Borda prize vector, no set S of teams can manipulate matches within S and/or throw matches to teams outside of S and improve their expected prize winnings. If $\vec{p}$ is ε-close to Borda, then S can gain at most $2|S|\varepsilon$ by manipulating matches within S and throwing matches to teams outside of S.*

Proof. The proof follows from Theorem 4.3, Lemma 4.7, and Lemma 4.5 (plus Lemma 4.6). □

5 TOURNAMENT RULES AND SORTING ALGORITHMS

In this section, we draw parallels between certain tournament rules and sorting algorithms. In particular, we notice that NRKotH is, according to some notion of equivalence, "equivalent to" the quicksort algorithm. Further, we provide examples of tournament rules equivalent to other sorting algorithms that are not consistent under expectation.

We can think of a sorting algorithm as a series of batches of comparisons over several rounds. The comparisons in a given round are determined by the results of the comparisons made in all previous rounds. For example, during selection sort, in the first round, we perform a batch of comparisons to find the smallest element. In the second round, we perform a batch of comparisons to find the second smallest element. In general, in the k^{th} round, we perform a batch of comparisons to find the k^{th} smallest element. Note that this representation of a sorting algorithm is not necessarily unique. Thus, we define the canonical batching of a sorting algorithm as follows:

Definition 5.1 (Canonical Batching). *The canonical batching of a (deterministic) sorting algorithm is defined as follows:*

1. *The first batch consists of all comparisons that do not depend upon the results of any other comparison. That is, these comparisons will get made by the algorithm on every input.*
2. *The second batch consists of all remaining comparisons that do not depend upon the results of any other remaining comparisons. That is, these comparisons will get made by the algorithm on every input, conditioned on the results of the first batch.*
3. *In general the k^{th} batch consists of all comparisons that are not in the first $k-1$ batches, and do not depend on the results of any other comparisons outside the first $k-1$ batches.*

For example, the canonical batching of QuickSort is as follows: the first batch consists of all comparisons made while running the QuickSelect algorithm on the first pivot. The second batch consists of all comparisons made while running the QuickSelect algorithm on a pivot from each half. In general, the k^{th} batch consists of all comparisons made while running the QuickSelect algorithm on a pivot from each of the 2^{k-1} blocks.

We notice that this definition of a canonical batching can easily extend to certain tournament rules. For instance, we notice that NRKotH has the same canonical batching as quick sort. We formalize this notion as follows:

Definition 5.2 (Equivalence of a Tournament Rule and a Sorting Algorithm). *A deterministic sorting algorithm s and a deterministic tournament rule r are said to be equivalent if they share the same canonical batching. Note that while computing the canonical batching of a sorting algorithm, we compare the values of numbers; in contrast, while computing the canonical batching of a sorting algorithm, we compare the results of individual matches. Consider running the rule r on a complete DAG, and running the sorting algorithm s on the nodes in the same DAG (where the comparator is the direction of the edge between the nodes). Then, we say s and r share the same canonical batching if the comparisons made in each batch are identical.*

If a randomized sorting algorithm s and a randomized tournament rule r can be coupled so that they are distributions over deterministic sorting algorithms and deterministic tournament rules that are equivalent, we say s and r are equivalent as well.

Observation 5.1. *NRKotH is equivalent to QuickSort.*

We now provide examples of tournament rules that are equivalent to the mergesort and bubblesort algorithms:

Definition 5.3 (MergeTR). *The tournament ranking rule MergeTR proceeds as follows, when given as input a tournament T on n teams:*

1. *If $n = 0$, return an empty ordering. If $n = 1$, return the ordering where this team has rank 1. Else, continue.*
2. *For $n > 1$, randomly divide the teams into two sets, A of size $\lfloor n/2 \rfloor$ and B of size $\lceil n/2 \rceil$. Run MergeTR on A and B and call the outputs σ_A and σ_B respectively. Let the order of the teams in A under σ_A be $a_1, \ldots, a_{|A|}$ and let the order of the teams in B under σ_B be $b_1, \ldots, b_{|B|}$.*
3. *Now set $ind = 1$, $i = 1$ and $j = 1$. While $i \leq |A|$ and $j \leq |B|$, compare a_i and b_i. If a_i beat b_i, set $\sigma(a_i) = ind$, increment i and ind by 1 and continue. Otherwise, set $\sigma(b_i) = ind$, increment j and ind by 1 and continue.*
4. *While $i \leq |A|$, set $\sigma(a_i) = ind$ and increment ind and i by 1.*
5. *While $j \leq |B|$, set $\sigma(b_j) = ind$ and increment ind and j by 1.*
6. *Output σ.*

Definition 5.4 (BubbleTR). *The tournament ranking rule BubbleTR proceeds as follows, when given as input a tournament T on n teams:*

1. *First arrange the n teams in some random order.*
2. *Do the following operation n times: starting from $i = 1$ and ending at $i = n - 1$, serially swap the positions of the i^{th} and $(i+1)^{st}$ teams if and only if the i^{th} team defeated the $(i+1)^{st}$ team.*

Lemma 5.1. *The MergeTR tournament ranking rule is not consistent under expectation, nor 2-SNM on the Borda prize vector.*

Proof. Consider a tournament with 4 teams (A, B, C, and D) with the following results: A beats to everyone except B, B loses to everyone except A, and D beats C.

There are 2 teams that A beats. We now analyze the expected rank for A outputted by MergeTR. With probability $1/3$, A plays B and C plays D in the first round, which would result in A getting rank 2. With probability $1/3$, A plays C and B plays D in the first round, which would result in A getting rank 1. Finally, with probability $1/3$, A plays D and B plays C in the first round, which would result in A getting rank 1. Therefore, the expected rank of A is $4/3$, which is not equal to $4 - 2$, so this algorithm is not consistent under expectation.

To see that MergeTR is not 2-SNM on the Borda prize vector, consider the tournament where A beats all three teams, and B beats C, C beats D, and D beats B. Consider simultaneously the tournament where A and B fix their match (so B beats A). Then because MergeTR is anonymous

and cover-consistent, A is always ranked first, and B is equally likely to be anywhere in $\{2, 3, 4\}$. So A's expected rank is 1, and B's expected rank is 3.

If instead A and B collude, then the execution of MergeTR proceeds as follows:

- Perhaps the teams are split so that A and B are together. Then B will be ahead of A in the list, while C will be ahead of D in the other list. B will get compared to C and win. Then A will get compared to C and win. Their ranks are 1 and 2, and this happens with probability $1/3$.
- Perhaps the teams are split so that A and D are together. Then A will be ahead of D, and B ahead of C. A will get compared to B, and B will rank first. A will then get compared to C and rank second. Their ranks are 1 and 2, and this happens with probability $1/3$.
- Perhaps the teams are split so that A and C are together. Then A will be ahead of C, and D beats B. A will get compared to D and rank first. C will get compared to D and lose, so D will get ranked second. Then C will get compared to B and lose, so B will rank third. So their ranks are 1 and 3, and this happens with probability $1/3$.

So in total, the sum of expected ranks is $1 + (2/3) \cdot 2 + (1/3) \cdot 3 = 10/3 < 4$. So A and B strictly gain from this manipulation (because total payoff under the Borda prize vector is proportional to expected rank). Intuitively, this manipulation is profitable because of the first two cases: if A beats B, then B is forced to compete with D, and will fall lower in the ranks. But if B beats A, then B is ranked highly, and A certainly wins its next match and is ranked immediately afterwards. □

Lemma 5.2. *The BubbleTR tournament ranking rule is not consistent under expectation, nor 2-SNM on the Borda prize vector.*

Proof. Consider a tournament with 4 teams (A, B, C, and D) with the following results: A loses to B and C but wins against D. B wins against C but loses to D and C loses to D. Over the randomness of the initial shuffling, the expected rank of B is $2.75 \neq 4 - 2 = 2$.

To see that BubbleTR is not 2-SNM on the Borda prize vector, consider a tournament where A loses to all three teams, and B beats D, D beats C, and C beats B. Then because BubbleTR is anonymous, the rank of A is always 4, and the expected rank of B is 2.

If instead A and B flip the outcome of their match, we have the following:

- If in the initial ordering A and B first and second (in either order) and D is third, then A is ranked 1 and B is ranked 2. This happens with probability $1/12$.
- If in the initial ordering D is last, then A will finish ranked 2, and B ranked 3. This happens with probability $1/4$.
- If in the initial ordering C and D are first and second (in either order), then A will finish ranked 3 and B will finish ranked 4. This happens with probability $1/6$.
- If in the initial ordering, A and D are first and second (in either order), then A will finish ranked 3 and B will finish ranked 4. This happens with probability $1/6$.
- If in the initial ordering, B and D are first and second (in either order), then B will finish ranked 1 and A will finish ranked 4. This happens with probability $1/6$.
- If in the initial ordering, A and C are first and second (in either order), and D is third, then A will finish 2, and B will finish 3. This happens with probability $1/12$.
- If in the initial ordering, B and C are first and second (in either order) and D is third, then A will finish 4 and B will finish 2. This happens with probability $1/12$.

So in total, the sum of expected ranks is $3 \cdot (1/12) + 5 \cdot (1/4) + 7 \cdot (1/6) + 7 \cdot (1/6) + 5 \cdot (1/6) + 5 \cdot (1/12) + 6 \cdot (1/12) = 67/12 < 6$. Therefore, A and B strictly profit from this manipulation. □

6 CONCLUSION AND OPEN PROBLEMS

We initiate the study of manipulation of tournament rules where teams wish to improve their rankings (rather than just their probability of winning). Specifically, we consider tournaments with a prize vector $\vec{p}$, where manipulators wish to improve their collective expected prize winnings. We design a cover-consistent tournament rule, Nested Randomized King of the Hill, and prove that no pair of teams can manipulate their match to gain more than $1/3$ in expected prize winnings when all rewards lie in $[0, 1]$ (and this is optimal, even among all Condorcet-Consistent rules). Indeed, better guarantees are not possible even when restricting attention to the prize vector $\langle 1, 0, \ldots, 0\rangle$. Furthermore, we prove that no set of teams can manipulate their matches to gain any reward under the uniform prize vector. This shows that the uniform prize vector enables both (a) the first *non*-manipulability results in this model, at all, and (b) a significantly simpler analysis of manipulability by sets of $k > 2$ teams when compared to any prior work. We further extend this result to near-uniform prize vectors with some approximation loss.

A nice direction for future work would be to consider other classes of prize vectors besides the worst-case and the near-uniform case. In particular, it would be interesting to see whether there are other classes of prize vectors for which NRKotH is optimal.

REFERENCES

[1] Alon Altman and Robert Kleinberg. 2010. Nonmanipulable Randomized Tournament Selections. (2010). https://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1703

[2] Alon Altman, Ariel D. Procaccia, and Moshe Tennenholtz. 2009. Nonmanipulable Selections from a Tournament. In *Proceedings of the 21st International Joint Conference on Artifical Intelligence* (Pasadena, California, USA) *(IJCAI'09)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 27–32. http://dl.acm.org/citation.cfm?id=1661445.1661451

[3] J. S. Banks. 1985. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare* 1, 4 (01 Dec 1985), 295–306. https://doi.org/10.1007/BF00649265

[4] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.

[5] Laszlo Csato. 2017. European qualifiers 2018 FIFA World Cup qualification can be manipulated. (09 2017). https://mpra.ub.uni-muenchen.de/id/eprint/83437

[6] Kimberly Ding and S. Matthew Weinberg. 2021. Approximately Strategyproof Tournament Rules in the Probabilistic Setting. In *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference (LIPIcs, Vol. 185)*, James R. Lee (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 14:1–14:20. https://doi.org/10.4230/LIPIcs.ITCS.2021.14

[7] Peter C. Fishburn. 1977. Condorcet Social Choice Functions. *SIAM J. Appl. Math.* 33, 3 (1977), 469–489. https://doi.org/10.1137/0133030 arXiv:http://dx.doi.org/10.1137/0133030

[8] David C. Fisher and Jennifer Ryan. 1992. Optimal Strategies for a Generalized "Scissors, Paper, and Stone" Game. *The American Mathematical Monthly* 99, 10 (1992), 935–942. https://doi.org/10.1080/00029890.1992.11995957 arXiv:https://doi.org/10.1080/00029890.1992.11995957

[9] Michael P. Kim, Warut Suksompong, and Virginia Vassilevska Williams. 2016. Who Can Win a Single-Elimination Tournament?. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*. 516–522. http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12194

[10] Michael P. Kim and Virginia Vassilevska Williams. 2015. Fixing Tournaments for Kings, Chokers, and More. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*. 561–567. http://ijcai.org/Abstract/15/085

[11] G. Laffond, J.F. Laslier, and M. Le Breton. 1993. The Bipartisan Set of a Tournament Game. *Games and Economic Behavior* 5, 1 (1993), 182 – 201. https://doi.org/10.1006/game.1993.1010

[12] J.F. Laslier. 1997. *Tournament Solutions and Majority Voting*. Springer. https://books.google.com/books?id=vYbGAAAAIAAJ

[13] Nicholas R. Miller. 1980. A New Solution Set for Tournaments and Majority Voting: Further Graph-Theoretical Approaches to the Theory of Voting. *American Journal of Political Science* 24, 1 (1980), 68–96. http://www.jstor.org/stable/2110925

[14] Marc Pauly. 2014. Can strategizing in round-robin subtournaments be avoided? *Social Choice and Welfare* 43, 1 (2014), 29–46. http://EconPapers.repec.org/RePEc:spr:sochwe:v:43:y:2014:i:1:p:29-46

[15] Jon Schneider, Ariel Schvartzman, and S. Matthew Weinberg. 2017. Condorcet-Consistent and Approximately Strategyproof Tournament Rules. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*. 35:1–35:20. https://doi.org/10.4230/LIPIcs.ITCS.2017.35

[16] Ariel Schvartzman, S. Matthew Weinberg, Eitan Zlatin, and Albert Zuo. 2020. Approximately Strategyproof Tournament Rules: On Large Manipulating Sets and Cover-Consistence. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA (LIPIcs, Vol. 151)*, Thomas Vidick (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 3:1–3:25. https://doi.org/10.4230/LIPIcs.ITCS.2020.3

[17] Kenneth A. Shepsle and Barry R. Weingast. 1984. Uncovered Sets and Sophisticated Voting Outcomes with Implications for Agenda Institutions. *American Journal of Political Science* 28, 1 (1984), 49–74. http://www.jstor.org/stable/2110787

[18] Isabelle Stanton and Virginia Vassilevska Williams. 2011. Rigging Tournament Brackets for Weaker Players. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*. 357–364. https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-069

[19] Virginia Vassilevska Williams. 2010. Fixing a Tournament. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*. http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1726