

Approximately Strategyproof Tournament Rules in the Probabilistic Setting

Kimberly Ding

Computer Science, Princeton University, NJ, USA

kding@princeton.edu

S. Matthew Weinberg

Computer Science, Princeton University, NJ, USA

smweinberg@princeton.edu

Abstract

We consider the manipulability of tournament rules which map the results of $\binom{n}{2}$ pairwise matches and select a winner. Prior work designs simple tournament rules such that no pair of teams can manipulate the outcome of their match to improve their probability of winning 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) [14, 15]. These lower bounds require the manipulators to know precisely the outcome of all future matches.

We take a beyond worst-case view and instead consider tournaments which are “close to uniform”: the outcome of all matches are independent, and no team is believed to win any match with probability exceeding $1/2 + \varepsilon$. We show that Randomized Single Elimination Bracket [14] and a new tournament rule we term Randomized Death Match have the property that no pair of teams can manipulate the outcome of their match to improve their probability of winning by more than $\varepsilon/3 + 2\varepsilon^2/3$, for all ε , and this is the best possible among any Condorcet-consistent tournament rule.

Our main technical contribution is a recursive framework to analyze the manipulability of certain forms of tournament rules. In addition to our main results, this view helps streamline previous analysis of Randomized Single Elimination Bracket, and may be of independent interest.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory

Keywords and phrases Tournaments, Incentive Compatibility, Recursive Analysis, Social Choice Theory

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.14

Funding *S. Matthew Weinberg*: Supported by NSF CAREER Award CCF-1942497.

1 Introduction

A tournament consists of n teams competing to win a championship via pairwise matches, and a tournament rule (possibly randomly) selects a single winner as a result of these matches. Tournament rules have been studied within Social Choice Theory for decades [6, 12, 16, 3, 7, 10, 11], see also [4] for a survey, and have gained attention from a few angles within TCS more recently [2, 1, 18, 17, 9, 8, 14, 15]. Our work follows the model studied in [2, 1, 14, 15] and seeks to design tournaments which are both fair (in that they select a reasonable winner, based on the match outcomes), and “as strategyproof as possible” subject to this.

More specifically, these works acknowledge that an undefeated team, if one exists, should surely win any reasonable tournament format. Formally, this property is termed *Condorcet-consistent* (Definition 4). These works also consider the possibility of two teams strategically manipulating the match between them to improve the probability that one of them wins. The situation to have in mind is that perhaps two teams sponsored by the same company enter an eSports tournament, and wish to maximize the probability that either of them take home the prize money. A rule is 2-Strongly Non-Manipulable (2-SNM, Definition 7) if no pair of teams can manipulate it to improve the probability that one of them wins.



© Kimberly Ding and S. Matthew Weinberg;

licensed under Creative Commons License CC-BY

12th Innovations in Theoretical Computer Science Conference (ITCS 2021).

Editor: James R. Lee; Article No. 14; pp. 14:1–14:20



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Initial works quickly established that no tournament rule exists which is both Condorcet-consistent and 2-SNM [2, 1]. More recent works study the extent to which tournament rules can be Condorcet-consistent and *approximately* 2-SNM. Specifically, a rule is 2-SNM- α if no pair of teams can improve the probability that one of them wins by more than α (Definition 7). [14, 15] design simple Condorcet-consistent tournament rules (defined in Section 3) which are 2-SNM-1/3, and also show this is the best possible *worst-case* guarantee.

1.1 Our Results: Probabilistic Tournaments

The lower bounds in previous works assume that the *deterministic future outcome* of all matches is known at the time of manipulation. While this is certainly a plausible scenario, most competitions worth watching have some element of uncertainty, even for matches between strong and weak teams. Indeed, [14, Open Problem 2] explicitly asks whether improved guarantees are possible if instead the teams have a common Bayesian prior about the possible outcomes of future matches which is *bounded away from deterministic*.

For example, consider the case where the outcome of all matches are uniformly at random. Then it is not hard to design a Condorcet-consistent tournament rule which is 2-SNM in this case. One example is a simple single-elimination bracket: when two manipulating teams face each other, each of them is equally likely to continue on and win the tournament, so manipulating doesn't help.

What if instead the outcome of all matches are not uniformly random, but close? More specifically, what if the match results are independent, and no team wins any match with probability more than $1/2 + \varepsilon$? When $\varepsilon = 0$, the previous paragraph establishes that rules exist where profitable manipulation is impossible. When $\varepsilon = 1/2$, [14, 15] establish that 2-SNM-1/3 tournaments exist, but no better. What about when $\varepsilon \in (0, 1/2)$? How do the achievable guarantees vary as a function of ε ?

Our main result resolves this question, and nails down the guarantees precisely as a function of ε . Moreover, we show that *the same tournament rule achieves the optimal guarantee for all ε* . Below, Randomized Single Elimination Bracket (Definition 12, henceforth RSEB) randomly seeds all teams, then runs a single-elimination bracket to determine the winner, and was shown to be 2-SNM-1/3 in [14]. Randomized Death Match (Definition 14, henceforth RDM) repeatedly picks two uniformly random teams to play a match, and eliminates the loser (and is first analyzed in this paper).

► **Informal Theorem 1** (See Theorems 30, 31). *For all $\varepsilon \in [0, 1/2]$, Randomized Death Match and Randomized Single Elimination Bracket are 2-SNM- $(\varepsilon/3 + 2\varepsilon^2/3)$ when match outcomes are independent, and no team wins any match with probability more than $1/2 + \varepsilon$. Moreover, for all $\varepsilon \in [0, 1/2]$, no Condorcet-consistent tournament rule is 2-SNM- α for any $\alpha < \varepsilon/3 + 2\varepsilon^2/3$, on the set of all tournaments with independent match outcomes that no team wins with probability more than $1/2 + \varepsilon$.*

1.2 Technical Highlights

We prove our main result by finding a strong structural similarity between tournaments like RDM and RSEB: they can be defined recursively. Specifically, RDM could be alternatively defined as “Pick two teams uniformly at random, and eliminate the loser of their match. Then, recurse on the remaining teams.” Similarly, RSEB can be alternatively defined as “Pick a uniformly random perfect matching between the teams, and eliminate all teams which

lose their match. Then, recurse on the remaining teams.” Even the Randomized King of the Hill rule (Definition 13, henceforth RKotH) defined by [15] fits this framework as well: “Pick a uniformly random team to play all other teams. Eliminate all teams who lose a match, and recurse on the remaining teams.”

In Section 3, we give a formal definition of what it means to be a recursive tournament rule. And in Section 4 (specifically, Theorem 21), we provide a general framework to analyze the manipulability of recursive tournament rules on probabilistic tournaments. This provides a fairly clean outline to analyze recursive tournament rules, and our main result then applies this framework to RDM and RSEB. In the $\varepsilon = 1/2$ case, our analysis of RSEB in isolation is perhaps not much simpler than that of [14], but our proof is arguably more structured. Indeed, a substantial fraction of our proof can be applied verbatim to other tournament rules like RDM, or applied verbatim to the $\varepsilon < 1/2$ case.

It is worth noting that our framework does face some technical barriers in accommodating RKotH (and we leave open whether RKotH achieves the same guarantees as RDM and RSEB). But, the technical barrier is easy to describe: the matches played in each round of RDM and RSEB form a matching – no team plays more than one match. In RKotH, some team plays multiple matches. It seems likely that our analysis would extend (perhaps with messier calculations) to any recursive rule where each round’s matches form a matching. But we highlight the aspects of our analysis which rely on this aspect of RDM/RSEB (and therefore don’t hold for RKotH), and believe this is a genuine barrier.

1.3 Further Related Work

We’ve already discussed the most related work [2, 1, 14, 15]. The model is first posed in [2], and [1] design tournaments which are 2-SNM and approximately Condorcet-consistent (e.g. pick a uniformly random match and declare the winner of that match to win the tournament). [14] first proposed to instead consider rules which are Condorcet-consistent and approximately strategyproof, and establishes that RSEB is 2-SNM-1/3 and that this is optimal. [15] considers larger manipulating sets (not relevant to this paper) and also designs RKotH, showing that it too is 2-SNM-1/3 and satisfies a stronger notion of fairness termed “cover-consistent”. In relation to these works, our main contribution is going beyond the worst-case to derive improved bounds when match outcomes are more uncertain. A technical contribution is our framework of recursive tournament rules.

Other recent works within TCS focus specifically on single-elimination brackets and manipulation in the form of a bracket designer trying to get a certain team to win [18, 17, 9, 8], or manipulability of particular tournament rules such as the World Cup qualifying procedure [13, 5]. Aside from being thematically related, there is no significant technical overlap with our work.

1.4 Roadmap

Section 2 immediately follows with definitions and preliminaries. Section 3 provides definitions concerning recursive tournament rules, and formally defines RDM and RSEB. Section 4 provides our framework for analyzing the manipulability of recursive tournament rules. Section 5 applies this framework, as a warmup, to rederive the main result of [14] and analyze RDM/RSEB in the deterministic case. Section 6 proves our main results, and Section 7 provides a brief conclusion.

2 Preliminaries

2.1 Tournament Rule Basics

In this section, we introduce notation consistent with prior work [1, 14, 15].

► **Definition 2** (Deterministic 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 3** (Tournament Rule). *A tournament rule r is a function that maps (deterministic) tournaments T to a distribution over teams, where $r_i(T) := \Pr[r(T) = i]$ denotes the probability that team i is declared the winner of tournament T under rule r . We use the shorthand $r_S(T) := \sum_{i \in S} r_i(T)$ to denote the probability that a team in S is declared the winner of tournament T under rule r .*

Finally, we are interested in tournament rules which satisfy basic notions of fairness. Importantly, note that Condorcet-consistency is a minimal notion of fairness, and in particular does not constrain the behavior of r on any tournament without a Condorcet winner.

► **Definition 4** (Condorcet-Consistent). *Team i is a Condorcet winner of a tournament T if i beats every other team (under T). A tournament rule r is Condorcet-consistent if for every tournament T with a Condorcet winner i , $r_i(T) = 1$ (whenever T has a Condorcet winner, that team wins with probability 1).*

2.2 Independent Probabilistic Tournaments

In this work, we study probabilistic tournaments. That is, we are interested in tournaments where the outcome of each match is not known to teams “in advance”, but teams share a Bayesian prior about the likelihood of each possible outcome. In particular, we consider when match outcomes are independent.

► **Definition 5** (Independent Probabilistic Tournament). *A probabilistic tournament T is just a distribution over deterministic tournaments. For notational convenience, we slightly abuse notation and refer by $r_i(T)$ to $\mathbb{E}[r_i(T)]$ (that is, $r_i(T)$ is the probability that team i wins when rule r is applied to T , over the randomness in r and the randomness in drawing T). A probabilistic tournament T is independent if all match outcomes in T are independent events. Observe that a probabilistic tournament T is fully defined by probabilities p_{ij}^T for all $i < j$, where p_{ij}^T denotes the probability that i beats j in tournament T .*

Observe that deterministic tournaments are also independent probabilistic tournaments, with each $p_{ij}^T \in \{0, 1\}$. Like prior work, we study tournament rules which are “as strategyproof as possible”. Because of our focus on independent probabilistic tournaments, we first refine previous definitions of non-manipulability.

► **Definition 6** (S -adjacent). *Two independent probabilistic tournaments T, T' are S -adjacent if $p_{ij}^T = p_{ij}^{T'}$ whenever $\{i, j\} \not\subseteq S$. That is, two independent probabilistic tournaments are S -adjacent when all (probabilistic) match outcomes are identical, except possibly for matches between two teams in S .*

Intuitively, two tournaments T, T' are S -adjacent if the set of teams S can manipulate the outcomes of matches between them in advance and cause the resulting (probabilistic) tournament to go from T to T' .

► **Definition 7** (Manipulating a Tournament). *For a set S of teams, independent probabilistic tournament T , and tournament rule r , we define $\alpha_S^r(T)$ to be the maximum winning probability that S can possibly gain by manipulating T to an S -adjacent T' . That is: $\alpha_S^r(T) := \max_{T': T' \text{ is } S\text{-adjacent to } T} \{r_S(T') - r_S(T)\}$.*

For a class \mathcal{T} of (independent probabilistic) tournaments, we also define $\alpha_k^r(\mathcal{T}) := \max_{T \in \mathcal{T}, S: |S| \leq k} \{\alpha_S^r(T)\}$. If $\alpha_k^r(\mathcal{T}) \leq \alpha$, we say r is k -Strongly Non-Manipulable at probability α with respect to \mathcal{T} (k -SNM $_{\mathcal{T}}\text{-}\alpha$). To match notation of prior work, we also say a tournament is k -SNM- α if it is k -SNM $_{\mathcal{S}^{1/2}}\text{-}\alpha$.¹

Finally, we also define $\alpha_k(\mathcal{T}) = \inf_{\text{Condorcet consistent } r} \{\alpha_k^r(\mathcal{T})\}$.

Intuitively, r is k -SNM $_{\mathcal{T}}\text{-}\alpha$ if no colluding set of $\leq k$ teams can manipulate a tournament in \mathcal{T} to improve the probability the winner is in S by more than α . The refinement over prior work is that the condition only holds for tournaments in \mathcal{T} – prior work only considers guarantees that hold over *all* tournaments. The additional notation in Definition 7 are just terms that will be helpful for later exposition.

We focus on independent probabilistic tournaments that are close to uniformly random.

► **Definition 8** (ε -Bounded Tournaments). *An independent probabilistic tournament T is weakly ε -bounded if for all i, j $p_{ij}^T \in [1/2 - \varepsilon, 1/2 + \varepsilon]$. We refer to \mathcal{T}^ε as the set of all ε -bounded tournaments.*

An independent probabilistic tournament T is strictly ε -bounded if for all i, j $p_{ij}^T \in \{1/2 - \varepsilon, 1/2 + \varepsilon\}$. We refer to \mathcal{S}^ε as the set of all strictly ε -bounded tournaments. It will be helpful to define the notation $\mathcal{T}_{\leq n}^\varepsilon$ (respectively, $\mathcal{S}_{\leq n}^\varepsilon$) as the set of all weakly- (respectively, strictly-) bounded tournaments on $\leq n$ teams.

For example, every deterministic tournament is 1/2-bounded, and the uniformly random tournament is 0-bounded. Our main results study $\alpha_k(\mathcal{T}^\varepsilon)$ as a function of ε . We conclude with a brief lemma relating $\alpha_k(\mathcal{T}^\varepsilon)$ to $\alpha_k(\mathcal{S}^\varepsilon)$, as direct analysis of $\alpha_k(\mathcal{S}^\varepsilon)$ is significantly simpler than direct analysis of $\alpha_k(\mathcal{T}^\varepsilon)$.

► **Proposition 9.** *For all rules r , and all ε, k , $\alpha_k^r(\mathcal{T}^\varepsilon) = \alpha_k^r(\mathcal{S}^\varepsilon)$. Therefore, for all ε, k , $\alpha_k(\mathcal{T}^\varepsilon) = \alpha_k(\mathcal{S}^\varepsilon)$.*

While the proof is deferred to Appendix A, the high-level outline is fairly intuitive. First, the “Therefore, . . .” statement follows trivially from the first portion of the lemma. Also, it is trivial to see that $\alpha_k^r(\mathcal{T}^\varepsilon) \geq \alpha_k^r(\mathcal{S}^\varepsilon)$, as $\mathcal{S}^\varepsilon \subseteq \mathcal{T}^\varepsilon$. So the interesting step is establishing $\alpha_k^r(\mathcal{T}^\varepsilon) \leq \alpha_k^r(\mathcal{S}^\varepsilon)$. Intuitively, this follows because all (independent probabilistic) tournaments in \mathcal{T}^ε can be written as convex combinations of tournaments in \mathcal{S}^ε , and one might expect that any particular tournament rule is most manipulable on extreme points (indeed this is true). With Proposition 9, we may restrict our study to $\alpha_k(\mathcal{S}^\varepsilon)$.

3 Recursive Tournament Rules

In this section, we formalize a class of tournament rules which have a recursive aspect to them. This will help us streamline previous analysis of [14], and also easily design a new tournament rule with matching guarantees. In addition, this view will help give us a clean outline to analyze the performance of these rules on tournaments in \mathcal{T}^ε , rather than just in the worst case.

¹ Note that $\mathcal{S}^{1/2}$ is the set of all deterministic tournaments, defined shortly below.

► **Definition 10** (Elimination Rule). *An elimination rule E takes as input a number n of teams and selects (possibly randomly) a set $M := E(n)$ of matches to play, with $|M| \in [1, n - 1]$.*

An elimination rule is matching if M is a (not necessarily perfect) matching with probability one.

► **Definition 11** (Recursive Tournament Rule). *A recursive tournament rule is fully defined by its elimination rule E . The recursive tournament rule r^E takes as input a tournament T , samples matches $M := E(n)$ to play, and then eliminates any team which loses a match in M .² Specifically, $T|_M$ denotes the induced subgraph of T on teams not eliminated by the matches in M . The tournament then recursively executes $r^E(T|_M)$ to select a winner. As a base case, when there is only one team left, that team is the winner.*

We now give three examples of elimination rules, and the resulting tournament rule. Randomized Single Elimination Bracket was first studied in [14], Randomized King of the Hill was first studied in [15], and Randomized Death Match is first studied in this paper.

► **Definition 12** (Randomized Single Elimination Bracket). *For a tournament T on n teams, let $n' := 2^{\lceil \log_2 n \rceil}$. Create $n' - n$ dummy players who all lose to the original n teams. Let M be matches corresponding to a uniformly random perfect matching (i.e. exactly $n'/2$ matches are played, and every team plays in exactly one match). Eliminate the losers and recurse on the remaining (non-dummy) teams.*

► **Definition 13** (Randomized King of the Hill). *Pick a uniformly random team i , to play all others. Observe that if i is a Condorcet winner, then i will be the only remaining team. Otherwise, i and every team it beats will be eliminated. Recurse on the remaining teams.*

► **Definition 14** (Randomized Death Match). *Pick two uniformly random teams (without replacement) and play their match. Eliminate the loser and recurse on the remaining teams.*

Observe that our definition of Randomized Single Elimination Bracket (RSEB) differs semantically from that given in [14], where the n' teams are uniformly permuted into n' seeds, and then the resulting bracket is played (without re-randomizing at each round). Observe that the two definitions are equivalent (identically distributed), however, as our definition simply produces the seeding by first figuring out the first-round matches, then the second-round matches, etc. Our definition of Randomized King of the Hill (RKotH) is identical (semantically) to that given in [15].

Randomized Death Match (RDM) is similar to Randomized Voting Caterpillar (RVC) [14]. Like RDM, RVC picks two uniformly random teams (without replacement) and eliminates the loser. However, rather than a “pure recursion”, RVC proceeds by picking *one* uniformly random remaining team to play the previous winner. This subtle distinction causes RDM to be 2-SNM-1/3 (Theorem 23), but not RVC ([14, Theorem 3.14]). Observe also that RSEB and RDM have matching elimination rules, but RKotH doesn’t (this is the main technical barrier in extending our analysis to RKotH).

Finally, observe that all three rules above are *anonymous*: relabeling the teams simply relabels the distribution over winners. This will play a role in our analysis.

► **Definition 15** (Anonymous). *A tournament rule r is anonymous if for every tournament T , and every permutation σ , and all i , $r_{\sigma(i)}(\sigma(T)) = r_i(T)$.*

² Observe that because $|M| \in [1, n - 1]$ that at least one team is eliminated, but not all teams are eliminated.

4 Key Framework

In this section, we propose an outline to analyze the manipulability of recursive tournament rules for independent probabilistic tournaments. For notational convenience, if a team i is not present in tournament T (e.g. because they were eliminated in an earlier round), we abuse notation and denote by $r_i(T) := 0$. Additionally, we let $M(T)$ denote the outcome of matches M for tournament T (i.e. for all $(u, v) \in M$, whether u or v wins in tournament T). Note that if T is probabilistic, $M(T)$ is a random variable, even after conditioning on M .

Our first step simply observes that recursive tournament rules can be analyzed recursively due to linearity of expectation. A (short) proof of Lemma 16 appears in Appendix B.

► **Lemma 16.** *Let r^E be any recursive tournament rule. Let T, T' be independent probabilistic tournaments on $n > 1$ teams, and let S be any subset of teams. Then:*

$$r_S^E(T') - r_S^E(T) = \mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)].$$

To help parse notation: recall that if T, T' are not deterministic, then the notation $r_S^E(T')$ (respectively, $r_S^E(T), r_S^E(T'|_M), r_S^E(T|_M)$) is taking an expectation over T' (respectively, $T, T'|_M, T|_M$), as per Definition 5.³ On the right-hand side, we are taking an expectation first over the matches M which are played. Inside the expectation, the teams in $T|_M$ and $T'|_M$ are still random variables (because they depend on the *outcome* of matches in M in T , which are still random after conditioning on M). And after the teams are determined by $M(T), M(T')$, the tournaments $T'|_M, T|_M$ are still probabilistic tournaments.

Importantly, observe that when $S = \{u, v\}$ and T, T' are S -adjacent, the tournaments $T|_M$ and $T'|_M$ may differ for two reasons. First, perhaps $(u, v) \in M$. In this case, perhaps $M(T) \neq M(T')$ (because T and T' can differ on the (u, v) match), and then $T|_M, T'|_M$ may have different sets of teams. Second, perhaps $(u, v) \notin M$, implying that $M(T) = M(T')$ (because T and T' are identically distributed outside of the (u, v) match). $T|_M$ and $T'|_M$ therefore have the same sets of teams, but it can still be that $p_{uv}^T \neq p_{uv}^{T'}$, so the tournaments $T|_M$ and $T'|_M$ can still differ due to this match (if both u and v are not eliminated).

The second step in our framework simply splits the recursive analysis into cases based on M , and the results of the matches $M(T)$. We define these cases clearly below, and then state our main framework.

► **Definition 17 (Base Case).** *We say that tournament T is a base case if $\alpha_S^r(T) = 0$. That is, it is not possible for S to gain by manipulating tournament T under rule r .*

One clear base case occurs if $S = \{u, v\}$, but u is not even in T . Lemma 22 later identifies another for anonymous tournament rules (if relabeling u and v doesn't change T except for the (u, v) match, then $\{u, v\}$ cannot gain by manipulating an anonymous tournament rule).

► **Definition 18 (Bad Terminal Event).** *M is a bad terminal event if $(u, v) \in M$. That is, a bad terminal event occurs when the (u, v) match is played this round. We denote the occurrence of this event by B .*

► **Definition 19 (Good Terminal Event).** *$M, M(T)$ is a good terminal event if $(u, v) \notin M$, and $M(T)$ is such that $T|_M$ is a base case.⁴ That is, the (u, v) match is not played now (so no gains from manipulation are possible), and no future gains are possible either. We denote the occurrence of this event by G .*

³ In particular, note that $T|_M$ and $T'|_M$ are not independent probabilistic tournaments, but are still probabilistic tournaments, as they are distributions over independent probabilistic tournaments.

⁴ Observe that because $(u, v) \notin M$, that if T' is S -adjacent to T then $M(T)$ and $M(T')$ are identically distributed.

► **Definition 20** (Recursive Event). $M, M(T)$ is a recursive event if $(u, v) \notin M$, but $M(T)$ is not such that $T|_M$ is a base case. That is, a recursive event occurs when the (u, v) match is not played now, but may be played later (and manipulating it may be beneficial). We denote the occurrence of this event by R .

We now state the theorem which drives our analysis. Observe that the bound claimed by Theorem 21 could in principle be (very) loose, but our subsequent sections show that this framework suffices to nail down $\alpha_2^r(\mathcal{T}^\varepsilon)$ for RDM and RSEB, and also that these are the best possible guarantees for any Condorcet-consistent tournament rule.

► **Theorem 21.** Let r^E be a recursive tournament rule, $\varepsilon \in [0, 1/2]$, and $S := \{u, v\}$. Let $b, g, c \geq 0$ (with $b + g > 0$) be so that for any $T \in \mathcal{S}^\varepsilon$, and any T' which is S -adjacent to T :

- $\Pr_{M \leftarrow E(n)}[B] \leq b$.
- $\Pr_{M \leftarrow E(n), M(T)}[G] \geq g$.
- $\mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)|B] \leq c$.

Then: $\alpha_2^{r^E}(\mathcal{T}^\varepsilon) \leq \frac{bc}{b+g}$.

Proof. We prove the theorem by induction, and focus on \mathcal{S}^ε first (extending to \mathcal{T}^ε using Proposition 9). As a base case, observe that when there are at most two teams remaining, no gains from manipulation are possible and therefore $\alpha_2^{r^E}(\mathcal{S}_{\leq 2}^\varepsilon) \leq 0 \leq \frac{bc}{b+g}$ as desired.

For the inductive hypothesis, assume that $\alpha_2^{r^E}(\mathcal{S}_{\leq n-1}^\varepsilon) \leq \frac{bc}{b+g}$ and consider now any tournament $T \in \mathcal{S}^\varepsilon$ on n teams, and an S -adjacent T' . We have the following chain of equalities, which essentially just breaks down $r_S^E(T') - r_S^E(T)$ based on the three events R, B, G (below, $I(X)$ denotes the indicator random variable for event X , which is 1 when event X occurs and 0 otherwise):

$$\begin{aligned} r_S^E(T') - r_S^E(T) &= \mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)] \\ &= \mathbb{E}_{M \leftarrow E(n), M(T)}[(r_S^E(T'|_M) - r_S^E(T|_M)) \cdot (I(R) + I(B) + I(G))] \\ &= \Pr_{M \leftarrow E(n), M(T)}[R] \cdot \mathbb{E}_{M \leftarrow E(n), M(T)}[r_S^E(T'|_M) - r_S^E(T|_M)|R] \\ &\quad + \Pr_{M \leftarrow E(n)}[B] \cdot \mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)|B] \\ &\quad + \Pr_{M \leftarrow E(n), M(T)}[G] \cdot \mathbb{E}_{M \leftarrow E(n), M(T)}[r_S^E(T'|_M) - r_S^E(T|_M)|G] \end{aligned}$$

The first line restates Lemma 16. The second line simply observes that exactly one of the events R, B, G occur, and also uses linearity of expectation to take an expectation also over $M(T)$. The third line just observes that $\mathbb{E}[Y \cdot I(X)] = \Pr[X] \cdot \mathbb{E}[Y|X]$ for any random variable Y and event X , breaks the sum into three parts (again using linearity of expectation), and observes that the event B is determined entirely by M and is independent of $M(T)$.

We now want to analyze the three terms separately. First, observe that by bullet three:

$$\mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)|B] \leq c. \quad (1)$$

Next, observe that in either a good terminal event or a recursive event, $(u, v) \notin M$. Because T, T' are S -adjacent, this means that $M(T), M(T')$ are identically distributed (and can therefore be coupled so that $M(T) = M(T')$ with probability one), so the teams in $T|_M$ and $T'|_M$ are therefore the same. Finally, because T is an *independent* probabilistic tournament, and at least one team that participates in every match in M is eliminated, $p_{ij}^{T|M} = p_{ij}^T$ for all teams i, j which are present in $T|_M$ (and also $p_{ij}^{T'|_M} = p_{ij}^{T'}$). This means that $T|_M \in \mathcal{S}_{\leq n-1}^\varepsilon$, and also that $T'|_M, T|_M$ are S -adjacent.

In a good terminal event, because $\alpha_2^{r^E}(T|_M) = 0$ by definition, we must therefore have:

$$\mathbb{E}_{M \leftarrow E(n), M(T)}[r_S^E(T'|_M) - r_S^E(T|_M)|G] \leq 0. \quad (2)$$

In a recursive event, we instead have:

$$\mathbb{E}_{M \leftarrow E(n), M(T)}[r_S^E(T'|_M) - r_S^E(T|_M)|R] \leq \alpha_2^{r^E}(\mathcal{S}_{\leq n-1}). \quad (3)$$

Finally, we may now use our inductive hypothesis and Equations (1), (2), (3) to conclude:

$$\begin{aligned} r_S^E(T') - r_S^E(T) &= \Pr_{M \leftarrow E(n), M(T)}[R] \cdot \mathbb{E}_{M \leftarrow E(n), M(T)}[r_S^E(T'|_M) - r_S^E(T|_M)|R] \\ &\quad + \Pr_{M \leftarrow E(n)}[B] \cdot \mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)|B] \\ &\quad + \Pr_{M \leftarrow E(n), M(T)}[G] \cdot \mathbb{E}_{M \leftarrow E(n), M(T)}[r_S^E(T'|_M) - r_S^E(T|_M)|G] \\ &\leq \Pr_{M \leftarrow E(n), M(T)}[R] \cdot \frac{bc}{b+g} + \Pr_{M \leftarrow E(n)}[B] \cdot c + \Pr_{M \leftarrow E(n), M(T)}[G] \cdot 0 \\ &= \frac{bc}{b+g} + \Pr_{M \leftarrow E(n)}[B] \cdot \left(c - \frac{bc}{b+g}\right) + \Pr_{M \leftarrow E(n), M(T)}[G] \cdot \left(0 - \frac{bc}{b+g}\right) \\ &\leq (1-b-g) \cdot \frac{bc}{b+g} + b \cdot c + g \cdot 0 = \frac{bc}{b+g} \end{aligned}$$

The first inequality just combines the work of Equations (1), (2), (3), and upper bounds the expected gains in all three cases using either the inductive hypothesis (R), direct hypothesis (B), or definition of good terminal event (G). The final inequality observes that $c \geq \frac{bc}{b+g} \geq 0$, so the bound is maximized when B occurs as often as possible, while G occurs as little as possible (consistent with the hypotheses).

We have now shown that for any $T \in \mathcal{S}^\varepsilon$, and any T' which is S -adjacent, that $r_S^E(T') - r_S^E(T) \leq \frac{bc}{b+g}$. This establishes that $\alpha_2^{r^E}(\mathcal{S}^\varepsilon) \leq \frac{bc}{b+g}$. Proposition 9 extends this to \mathcal{T}^ε . \blacktriangleleft

Theorem 21 is our main framework for analysis. The remainder of this paper now computes the bounds required for the three bullets for the two tournament rules of interest as a function of ε , and substitutes to obtain tight bounds. Finally, it is worth briefly noting that the definition of $\alpha_2^r(\mathcal{T}^\varepsilon)$ semantically assumes that teams must decide how to manipulate the outcome of their match *in advance*, before any matches are played. Still, any analysis that follows from Theorem 21 applies *even to manipulations which are decided upon as the match is played*. This is because the bound in bullet three of Theorem 21 must hold over *all* S -adjacent T' , not just one which was decided in advance. This claim is not central to our main results, and making it formal would be notationally cumbersome, so we provide only this brief discussion.

5 Warmup: Rederiving Bounds for Deterministic Tournaments

In this section, we rederive the main result of [14] (RSEB is 2-SNM-1/3), and analyze a novel tournament rule (RDM is 2-SNM-1/3) using a simple application of Theorem 21. This will require one further lemma about anonymous tournament rules to find additional base cases. The proof is in Appendix C.

► **Lemma 22.** *Let r be any anonymous tournament rule, $S := \{u, v\}$, and T, T' be independent probabilistic tournaments which are S -adjacent and satisfy $p_{uw}^T = p_{vw}^T$ for all $w \notin \{u, v\}$. Then $r_S(T) = r_S(T')$.*

We now prove that RDM is 2-SNM-1/3. Recall that $\alpha_2(\mathcal{S}^{1/2}) = 1/3$, so this is optimal.

► **Theorem 23.** *RDM is 2-SNM-1/3. Or in our language, $\alpha_2^{\text{RDM}}(\mathcal{T}^{1/2}) = 1/3$.*

Proof. Recall that we need to lower bound the probability of a good terminal event, upper bound the probability of a bad terminal event, and upper bound the gains from manipulation in case of a bad terminal event. For the deterministic case, the latter bound is particularly simple, and we will just observe that clearly $\mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)|B] \leq 1$. This is simply because the maximum probability that any coalition can win in any tournament is 1, and the minimum is 0. So we have established that RDM satisfies $c = 1$.

To bound the probability of a bad terminal event, observe that a bad terminal event occurs only when u plays v in this round, which happens with probability exactly $1/\binom{n}{2}$. So RDM satisfies $b = 1/\binom{n}{2}$.

For the good terminal events, we claim that $\Pr[G] \geq 2/\binom{n}{2}$. First, observe that once we show this, we can plug into Theorem 21 and conclude $\alpha_2^{\text{RDM}}(\mathcal{T}^{1/2}) \leq \frac{1-1/\binom{n}{2}}{1/\binom{n}{2}+2/\binom{n}{2}} = 1/3$.

To see this bound, let ℓ_u denote the number of teams which beat u but not v , and ℓ_v denote the number of teams which beat v but not u . Without loss of generality let $\ell_u \geq \ell_v$.

If $\ell_u + \ell_v = 0$, then Lemma 22 already establishes that u and v can gain nothing by manipulating. If $\ell_u + \ell_v \geq 2$, then whenever u or v play a team which beat them, we have a good terminal event (because u or v is already eliminated before having ever played the (u, v) match, resulting in a base case). This happens with probability at least $2/\binom{n}{2}$, as desired.

If $\ell_u + \ell_v = 1$, then $\ell_u = 1, \ell_v = 0$. Let w be the unique team which beats u but not v . We claim that if w plays *either* u or v that we are in a good terminal event. Indeed, if w plays u , then u is eliminated having never played the (u, v) match. If instead w plays v , then w is eliminated, but now there are no remaining teams which beat u but not v (or vice versa) and Lemma 22 asserts that there are no further gains from manipulation. The probability that w plays u or v is $2/\binom{n}{2}$, as desired.

This handles all possible cases, and establishes that $g \geq 2/\binom{n}{2}$ in all cases. Plugging into Theorem 21 as described above completes the proof. ◀

The analysis of RSEB is extremely similar to RDM, and requires only slightly more calculations to bound the probability of a good terminal event. This proof structure is fairly different than the original analysis in [14], and highlights the similarities to other recursive tournament rules. A complete proof of Theorem 24 appears in Appendix C.

► **Theorem 24** ([14]). *RSEB is 2-SNM-1/3. Or in our language, $\alpha_2^{\text{RSEB}}(\mathcal{T}^{1/2}) = 1/3$.*

We wrap up our warmup by highlighting key points of the analysis which will be relevant for our main results. First, observe that our analysis of both rules succeeded by simply upper-bounding c by 1. Improving this as a function of ε is the biggest technical difference between our warmup and main results. Second, observe that our analysis required Lemma 22 in case $\ell_u < 2$. In particular, we needed good terminal events *even when both u and v were not eliminated* (and this need continues in main results).

6 Optimal Bounds for Independent Probabilistic Tournaments

Before analyzing our two tournament rules, we extend the simple 3-team lower bound of [14] for $\alpha_2(\mathcal{T}^{1/2})$ to $\alpha_2(\mathcal{T}^\varepsilon)$. The proof is in Appendix D.

► **Lemma 25.** $\alpha_2(\mathcal{T}^\varepsilon) \geq \frac{1}{3}\varepsilon + \frac{2}{3}\varepsilon^2$.

6.1 Gauntlets: Upper Bounding Gains from Bad Terminal Events

As previously noted, the main difference between our warmup and main results is bounding gain from bad terminal events. We provide a short, but key, structural insight about recursive tournament rules. Intuitively, a team u wins under rule r^E as long as they survive all elimination matches. Our key observation is that this defines a *gauntlet* of teams such that u wins if and only if they defeat every team in the gauntlet.

► **Definition 26** (Gauntlet). *For deterministic tournament T , recursive tournament rule r , and team u , let T' be such that the outcome of the (v, w) match is the same for all $v, w \neq u$, but u is a Condorcet winner. The gauntlet for u in tournament T under recursive rule r , $G_u^r(T)$, is the set of teams that u plays in elimination matches when r is executed on T' . If r is randomized, then $G_u^r(T)$ is a random variable.*

If T is a probabilistic tournament, we extend the notation $G_u^r(T)$ to be the random variable which first samples T , then outputs $G_u^r(T)$ (again over randomness in r).

For intuition, consider RSEB. u wins RSEB if and only if they win each of their $\lceil \log_2(n) \rceil$ matches, so their gauntlet is a list of $\lceil \log_2(n) \rceil$ teams, one per round. For RDM, however, the set of matches that u plays is itself a random variable (depending on how many times u is selected to play), but u still must win all these matches in order to win. For RKotH, the size of the gauntlet is a random variable as well. Importantly, however, observe that for all three rules (and any elimination rule), as soon as u loses a match, they are eliminated, so their gauntlet opponents can be set assuming that u won all previous matches.

Importantly, observe that u 's gauntlet does *not* depend on the outcome of any of its own matches (because T' immediately causes u to win all its matches anyway). This lets us make the following key observation, which requires E to be a matching elimination rule.

► **Lemma 27.** *Let r^E be an anonymous recursive tournament rule with a matching elimination rule, T be an independent probabilistic tournament on n teams, $M \leftarrow E(n)$, and u, v be two teams. Let also $(u, v) \in M$ and w denote the winner of the (u, v) match. Then the random variable $G_w^{r^E}(T|_M)$ is independent of w .*

Proof. Observe first that because M is a matching which contains (u, v) , that exactly one of $\{u, v\}$ (namely, w) is present in $T|_M$. Moreover, because T is independent probabilistic, the remaining teams in $T|_M$ are independent of w . Because the definition of w 's gauntlet immediately considers a tournament T' which replaces the outcome of all matches involving w by having w be a Condorcet winner, the tournament T' is independent of w *except for whether w is labeled as “ u ” or “ v ”*. But because $r^E(\cdot)$ is anonymous, its behavior on T' is independent of w 's label. This completes the proof. ◀

Note that Lemma 27 fails when M is not a matching. This is because: (a) perhaps both u and v are eliminated, and w is undefined, but also (b) even conditioned on w being defined, the set of teams in $T|_M$ can depend on w . This is the main technical challenge in extending beyond matching elimination rules.⁵

⁵ For example, in RKotH: conditioned on $w := u$, we know either that u was a Condorcet winner, or that v was selected to play everyone, so all teams which lose to v are eliminated. If instead $w := u$, we know either that v was a Condorcet winner, or that u was selected to play everyone. The teams which lose to u vs. v could be different, so the Lemma fails to hold.

14:12 Approximately Strategyproof Tournament Rules in the Probabilistic Setting

► **Corollary 28.** Let r^E be an anonymous tournament rule with a matching elimination rule, $T \in \mathcal{T}^\varepsilon$, $S = \{u, v\}$ be any two teams, and T' be S -adjacent to T . Then:

$$\mathbb{E}_{M \leftarrow E(n)}[r_S^E(T'|_M) - r_S^E(T|_M)|B] \leq 2\varepsilon(\frac{1}{2} + \varepsilon).$$

Proof. Because we are in a bad terminal event, this means that $(u, v) \in M$. If we let w denote the winner of the (u, v) match in T , and w' denote the winner in T' , then we know that $G_w^{r^E}(T|_M)$ and $G_{w'}^{r^E}(T'|_M)$ are identically distributed by Lemma 27. This means that:

$$\begin{aligned} r_S^E(T|_M) &= \mathbb{E}_{G_w^{r^E}(T|_M)} \left[\prod_{x \in G_w^{r^E}(T|_M)} p_{wx}^T \right] \\ r_S^E(T'|_M) &= \mathbb{E}_{G_{w'}^{r^E}(T'|_M)} \left[\prod_{x \in G_{w'}^{r^E}(T'|_M)} p_{w'x}^{T'} \right] = \mathbb{E}_{G_w^{r^E}(T|_M)} \left[\prod_{x \in G_w^{r^E}(T|_M)} p_{w'x}^T \right] \\ \Rightarrow r_S^E(T'|_M) - r_S^E(T|_M) &= \mathbb{E}_{G_w^{r^E}(T|_M)} \left[\prod_{x \in G_w^{r^E}(T|_M)} p_{w'x}^T - \prod_{x \in G_w^{r^E}(T|_M)} p_{wx}^T \right] \\ &\leq \mathbb{E}_{G_w^{r^E}(T|_M)} \left[\prod_{x \in G_w^{r^E}(T|_M)} (\frac{1}{2} + \varepsilon) - \prod_{x \in G_w^{r^E}(T|_M)} (\frac{1}{2} - \varepsilon) \right] \\ &= \mathbb{E}_{G_w^{r^E}(T|_M)} \left[(\frac{1}{2} + \varepsilon)^{|G_w^{r^E}(T|_M)|} - (\frac{1}{2} - \varepsilon)^{|G_w^{r^E}(T|_M)|} \right] \\ &\leq 2\varepsilon \end{aligned}$$

The first two lines follow by definition of the gauntlet, and Lemma 27. The third line is basic algebra. The fourth line follows as $T \in \mathcal{T}^\varepsilon$. The fifth line is again basic algebra, and the final line invokes Lemma 29, which is stated below (proof omitted from this version). The above calculations hold for any $T|_M, T'|_M$.

► **Lemma 29.** For all $n \in \mathbb{N}_{\geq 0}$, and $\varepsilon \in [0, 1/2]$: $(\frac{1}{2} + \varepsilon)^i - (\frac{1}{2} - \varepsilon)^n \leq 2\varepsilon$.

To see where the additional factor of $(\frac{1}{2} + \varepsilon)$ comes from, recall that $p_{uv}^T \in [1/2 - \varepsilon, 1/2 + \varepsilon]$. Therefore, $p_{uv}^{T'} - p_{uv}^T \leq \frac{1}{2} + \varepsilon$. Consider now coupling the tournaments T and T' so the the result of every match except for the one between u and v is identical (this is possible because T and T' are independent probabilistic). Under this coupling, T and T' are identical with probability at least $\frac{1}{2} - \varepsilon$ (and gains from manipulation are clearly only possible when $T \neq T'$, which occurs with probability at most $\frac{1}{2} + \varepsilon$, and independently of M). Therefore, $T \neq T'$ with probability at most $(\frac{1}{2} + \varepsilon)$ and conditioned on this, $r_S^E(T'|_M) - r_S^E(T|_M) \leq 2\varepsilon$. ◀

6.2 RDM and RSEB

Corollary 28 is the main technical lemma to extend from $\mathcal{T}^{1/2}$ to \mathcal{T}^ε . We begin with RDM.

► **Theorem 30.** $\alpha_2^{RDM}(\mathcal{T}^\varepsilon) = \alpha_2(\mathcal{T}^\varepsilon) = \frac{1}{3}\varepsilon + \frac{2}{3}\varepsilon^2$.

Proof. Let $T \in \mathcal{S}^\varepsilon$. Recall that we need to lower bound the probability of a good terminal event, upper bound the probability of a bad terminal event, and upper bound the gains from manipulation in case of a bad terminal event. We have already upper bounded the gains from a bad terminal event using Corollary 28 and can therefore take $c = 2\varepsilon(\frac{1}{2} + \varepsilon)$.

To bound the probability of a bad terminal event, observe that a bad terminal event occurs only when u plays v in this round, which happens with probability exactly $1/\binom{n}{2}$. So RDM satisfies $b = 1/\binom{n}{2}$.

For the good terminal events, we claim that $\Pr[G] \geq 2/\binom{n}{2}$. First, observe that once we show this, we can plug into Theorem 21 and conclude $\alpha_2^{\text{RDM}}(\mathcal{T}^\varepsilon) \leq \frac{2\varepsilon(\frac{1}{2}+\varepsilon) \cdot 1/\binom{n}{2}}{1/\binom{n}{2} + 2/\binom{n}{2}} = \frac{1}{3}\varepsilon + \frac{2}{3}\varepsilon^2$.

To see this bound, let ℓ_u denote the number of teams which beat u with probability $(1/2 + \varepsilon)$ but v with probability $(1/2 - \varepsilon)$, and ℓ_v denote the number of teams which beat v with probability $(1/2 + \varepsilon)$ but u with probability $(1/2 - \varepsilon)$. Without loss of generality let $\ell_u \geq \ell_v$.

If $\ell_u = 0$, then Lemma 22 already establishes that u and v can gain nothing by manipulating. If $\ell_u + \ell_v \geq 2$, then whenever u or v plays a team (not in $\{u, v\}$) which beats them, we have a good terminal event (because u or v is already eliminated before having ever played the (u, v) match, so manipulating the match has no impact). Observe that any team which beats u (respectively, v) with probability $(1/2 + \varepsilon)$ beats v (respectively, u) with probability at least $(1/2 - \varepsilon)$. Therefore, we have a good terminal event with probability at least $2(1/2 + \varepsilon)/\binom{n}{2} + 2(1/2 - \varepsilon)/\binom{n}{2} = 2/\binom{n}{2}$, as desired.⁶

If $\ell_u = 1, \ell_v = 0$, then let w be the unique team which beats u with probability $(1/2 + \varepsilon)$ but v with probability $(1/2 - \varepsilon)$. We claim that if w plays *either* u or v that we have a good terminal event. Indeed, if w wins, then either u or v are eliminated having never played the (u, v) match. If instead w loses this match, then w is eliminated, but now $p_{ux}^T = p_{vx}^T$ for all remaining teams x , and Lemma 22 asserts that there are no further gains from manipulation. The probability that w plays u or v is $2/\binom{n}{2}$, as desired.

This handles all possible cases, and establishes that $g \geq 2/\binom{n}{2}$ in all cases. Plugging into Theorem 21 as described above completes the proof. \blacktriangleleft

The analysis for RSEB is again similar to RDM, but some calculations are more involved.

► **Theorem 31.** $\alpha_2^{\text{RSEB}}(\mathcal{T}^\varepsilon) = \alpha_2(\mathcal{T}^\varepsilon) = \frac{1}{3}\varepsilon + \frac{2}{3}\varepsilon^2$.

Proof. Let $T \in \mathcal{S}^\varepsilon$. We have already upper bounded the gains from a bad terminal event using Corollary 28 and can therefore take $c = 2\varepsilon(\frac{1}{2} + \varepsilon)$.

To bound the probability of a bad terminal event, observe that a bad terminal event occurs only when u plays v this round, which happens with probability exactly $1/(n' - 1)$.⁷ So RSEB satisfies $b = 1/(n' - 1)$.

For the good terminal events, we claim that $\Pr[G] \geq 2/(n' - 1)$. First, observe that once we show this, we can plug into Theorem 21 and conclude $\alpha_2^{\text{RSEB}}(\mathcal{T}^\varepsilon) \leq \frac{2\varepsilon(\frac{1}{2}+\varepsilon) \cdot 1/(n'-1)}{1/(n'-1) + 2/(n'-1)} = \frac{1}{3}\varepsilon + \frac{2}{3}\varepsilon^2$.

To see this bound (which requires more calculations than previous proofs), let ℓ_u denote the number of teams which beat u with probability $(1/2 + \varepsilon)$ but v with probability $(1/2 - \varepsilon)$, and ℓ_v denote the number of teams which beat v with probability $(1/2 + \varepsilon)$ but u with probability $(1/2 - \varepsilon)$. Without loss of generality let $\ell_u \geq \ell_v$.

Case One: $\ell_u + \ell_v \leq 1$. If $\ell_u = 0$, then Lemma 22 already establishes that u and v can gain nothing by manipulating. If $\ell_u = 1, \ell_v = 0$, then let w be the unique team which beats u with probability $(1/2 + \varepsilon)$ but v with probability $(1/2 - \varepsilon)$. We claim that if w

⁶ To be extra clear, the two matches $(w, u), (w, v)$ together contribute probability $1/\binom{n}{2}$ to the probability of a good terminal event, as long as w beats either u or v with probability $(1/2 + \varepsilon)$. Because there are two such teams, we get this twice.

⁷ Recall in RSEB that n' denotes the total number of teams plus dummy teams, and is $2^{\lceil \log_2(n) \rceil}$

plays either u or v that we are in a good terminal event. Indeed, if w wins this match, then either u or v are eliminated having never played the (u, v) match. If instead w loses this match, then w is eliminated, but now $p_{ux}^T = p_{vx}^T$ for all remaining teams x , and Lemma 22 asserts that there are no further gains from manipulation. The probability that w plays u or v is $2/(n' - 1)$, as desired.

Case Two: $\ell_u + \ell_v = 2$. Next, consider the case where $\ell_u + \ell_v = 2$, and call the two relevant teams x, w . Observe first that if u plays x and v plays w , or if u plays w and v plays x , then we are surely in a good terminal event. This is because either (a) u or v is eliminated without having played the (u, v) match, or (b) both x and w are eliminated (allowing us to invoke Lemma 22). This occurs with probability $\frac{2}{(n' - 1)(n' - 3)}$. There are also the cases where exactly one of $\{w, x\}$ plays a team in $\{u, v\}$. For any given pair (a, b) , with $a \in \{w, x\}$ and $b \in \{u, v\}$, this case occurs with probability $\frac{(n' - 4)}{(n' - 1)(n' - 3)}$.⁸ Two of these cases contribute at least a $(1/2 + \varepsilon)$ probability of eliminating the team in $\{u, v\}$, and the other two contribute at least a $(1/2 - \varepsilon)$ probability. So in total, all four cases contribute at least $\frac{2(n' - 4)}{(n' - 1)(n' - 3)}$, and together we get that a good terminal event occurs with probability at least:

$$\frac{2}{(n' - 1)(n' - 3)} + \frac{2(n' - 4)}{(n' - 1)(n' - 3)} = \frac{2(n' - 3)}{(n' - 1)(n' - 3)} = \frac{2}{n' - 1}.$$

Case Three: $\ell_u + \ell_v \geq 3$. Next, consider the case where $\ell_u + \ell_v \geq 3$ (observe that this implies $n \geq 5$). We will show that either u or v are eliminated with probability at least $2/(n' - 1)$. Indeed, let L_u denote the set of teams which beat u with probability $(1/2 + \varepsilon)$ (but not v), and L_v denote the set of teams which beat v with probability $(1/2 + \varepsilon)$ (but not u). Then consider the case where u plays a team in L_u , or v plays a team in L_v . Conditioned on this, both u and v survive with probability at most $(1/2 + \varepsilon)(1/2 - \varepsilon) \leq 1/4$, so one of $\{u, v\}$ is eliminated with probability at least $3/4$. So one sufficient condition would be to establish that u plays a team in L_u or v plays a team in L_v with probability at least $\frac{8/3}{n' - 1}$ (because conditioned on this, one of $\{u, v\}$ is eliminated with probability $3/4$, for a total probability of at least $\frac{2}{n' - 1}$ that one of $\{u, v\}$ is eliminated). In the subcase that $\ell_u \geq 3$, then the probability is in fact at least $3/(n' - 1)$, as desired. Clearly, this probability is monotone in ℓ_u, ℓ_v , so this leaves the only remaining subcase as $\ell_u = 2, \ell_v = 1$.

Subcase Three-A: $\ell_u = 2, \ell_v = 1, n \geq 9$. For the subcase of $\ell_u = 2, \ell_v = 1$, the probability that u plays a team in L_u or v plays a team in L_v is $\frac{2}{n' - 1} + \frac{(n' - 5)}{(n' - 1)(n' - 3)}$.⁹ As $n' \geq 16$ (because $n \geq 9$ and n' is a power of 2), we have that $(n' - 5)/(n' - 3) \geq 11/13$, meaning that $\frac{2}{n' - 1} + \frac{(n' - 5)}{(n' - 1)(n' - 3)} \geq \frac{37/13}{n' - 1} > \frac{8/3}{n' - 1}$, which resolves this case by the work above.

Subcase Three-B: $\ell_u = 2, \ell_v = 1, \varepsilon \geq \frac{1}{2\sqrt{13}}, n \in [5, 8]$. When $n \in [5, 8]$, we have $n' = 8$. This means that $(n' - 5)/(n' - 3) = 3/5$, and therefore $\frac{2}{n' - 1} + \frac{(n' - 5)}{(n' - 1)(n' - 3)} = \frac{13/5}{n' - 1}$. Unfortunately, this is $< \frac{8/3}{n' - 1}$, meaning that this case doesn't immediately resolve by the method above. In this range, we do an explicit case analysis. First, observe that u and v are least likely to be eliminated when there are more dummy teams (because u and v

⁸ Because the probability that a plays b is $1/(n' - 1)$, and the probability that the other two teams do not play, conditioned on this, is $(n' - 4)/(n' - 3)$.

⁹ To see this, observe that the probability that u plays a team in L_u is $2/(n' - 1)$. The probability that u plays a team not in $L_u \cup L_v \cup \{v\}$ is $(n' - 5)/(n' - 1)$. Conditioned on this, the probability that v plays the team in L_v is $1/(n' - 3)$ (and this is the only way that v can possibly play the team in L_v without u playing L_u).

beat dummy teams with probability one, but real teams with probability at most $1/2 + \varepsilon$). So the worst case to consider is when $n = 5$: two real teams beat only u , one beats only v , and there are three dummy teams. We just have to compute several cases:

- Perhaps u plays one of the ℓ_u teams, and v plays one of the ℓ_v teams. This occurs with probability $\frac{2}{7} \cdot \frac{1}{5}$, and eliminates u or v with probability $1 - (1/2 - \varepsilon)^2 = 3/4 + \varepsilon - \varepsilon^2$.
- Perhaps u plays one of the ℓ_u teams, and v plays the other. This occurs with probability $\frac{2}{7} \cdot \frac{1}{5}$, and eliminates u or v with probability $1 - (1/2 + \varepsilon)(1/2 - \varepsilon) = 3/4 + \varepsilon^2$.
- Perhaps u plays one of the ℓ_u teams, and v plays a dummy team. This occurs with probability $\frac{2}{7} \cdot \frac{3}{5}$, and eliminates u with probability $1/2 + \varepsilon$.
- Perhaps u plays one of the ℓ_v teams, and v plays one of the ℓ_u teams. This occurs with probability $\frac{1}{7} \cdot \frac{2}{5}$, and eliminates u or v with probability $1 - (1/2 + \varepsilon)^2 = 3/4 - \varepsilon - \varepsilon^2$.
- Perhaps u plays one of the ℓ_v teams, and v plays a dummy team. This occurs with probability $\frac{1}{7} \cdot \frac{3}{5}$, and eliminates u with probability $1/2 - \varepsilon$.
- Perhaps u plays a dummy team, and v plays one of the ℓ_v teams. This occurs with probability $\frac{3}{7} \cdot \frac{1}{5}$, and eliminates v with probability $1/2 + \varepsilon$.
- Perhaps u plays a dummy team, and v plays one of the ℓ_u teams. This occurs with probability $\frac{3}{7} \cdot \frac{2}{5}$, and eliminates v with probability $1/2 - \varepsilon$.
- Perhaps both u and v play dummy teams. This happens with probability $\frac{3}{7} \cdot \frac{2}{5}$, but eliminates neither team.

So either u or v is eliminated (without playing each other) with probability:

$$\begin{aligned} & \frac{2}{35} \cdot (3/4 + \varepsilon - \varepsilon^2) + \frac{2}{35} \cdot (3/4 + \varepsilon^2) + \frac{6}{35} \cdot (1/2 + \varepsilon) + \frac{2}{35} \cdot (3/4 - \varepsilon - \varepsilon^2) \\ & + \frac{3}{35} \cdot (1/2 - \varepsilon) + \frac{3}{35} \cdot (1/2 + \varepsilon) + \frac{6}{35} \cdot (1/2 - \varepsilon) \\ & = \frac{27}{70} - \frac{2}{35} \varepsilon^2 \geq \frac{26}{70} > 2/7. \end{aligned}$$

Note that when $n \in [5, 8]$, the probability of a bad event is exactly $1/7$, and the above work establishes that in case Three-B, the probability of a good terminal event is at least twice that of a bad terminal event.

The arguments above handle all possible cases, and establishes that $g \geq 2/(n' - 1)$ in all cases. Plugging into Theorem 21 as described above completes the proof. \blacktriangleleft

7 Conclusion and Open Problems

We take a beyond worst-case view on manipulability of tournament rules, and nail down optimal guarantees as a function of the uncertainty of match outcomes. Specifically, our main result shows that $\alpha_2(\mathcal{T}^\varepsilon) = \varepsilon/3 + 2\varepsilon^2/3$, and this is achieved *for all* ε by Randomized Death Match and Randomized Single Elimination Bracket. Our main technical contribution is a framework to analyze recursive tournament rules.

There are two natural directions for future work. The first concerns further work in the probabilistic setting: does $\alpha_2^{\text{RKotH}}(\mathcal{T}^\varepsilon) = \varepsilon/3 + 2\varepsilon^2/3$? The main technical barrier is replacing Lemma 27, which only holds for matching elimination rules.¹⁰ In addition, it is interesting to analyze probabilistic tournaments which are not independent. Here there are technical barriers to overcome (many of our steps do require independence), but also conceptual

¹⁰ There are other barriers to using our precise definitions, but these barriers seem semantic rather than substantial.

ones: if tournament match outcomes are correlated, should we consider manipulations which are correlated with external outcomes as well? If so, is there a natural way to consider manipulations which are “not more correlated than the original tournament itself”?

A second direction concerns applications of our recursive framework towards other problems in approximately strategy-proof tournament design. For example, it is still an open question following [15] what is $\alpha_k(\mathcal{T}^{1/2})$ for *any* $k > 2$. Our recursive framework may prove useful for analyzing this, or at least determining achievable guarantees for recursive rules.

References

- 1 Alon Altman and Robert Kleinberg. Nonmanipulable randomized tournament selections. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*. AAAI Press, 2010. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1703>.
- 2 Alon Altman, Ariel D. Procaccia, and Moshe Tennenholtz. Nonmanipulable selections from a tournament. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI'09*, pages 27–32, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc. URL: <http://dl.acm.org/citation.cfm?id=1661445.1661451>.
- 3 J. S. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 1(4):295–306, December 1985. doi:10.1007/BF00649265.
- 4 Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- 5 Laszlo Csato. European qualifiers 2018 FIFA world cup qualification can be manipulated, September 2017. URL: <https://mpra.ub.uni-muenchen.de/id/eprint/83437>.
- 6 Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977. doi:10.1137/0133030.
- 7 David C. Fisher and Jennifer Ryan. Optimal strategies for a generalized “scissors, paper, and stone” game. *The American Mathematical Monthly*, 99(10):935–942, 1992. doi:10.1080/00029890.1992.11995957.
- 8 Michael P. Kim, Warut Suksompong, and Virginia Vassilevska Williams. Who can win a single-elimination tournament? In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 516–522, 2016. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12194>.
- 9 Michael P. Kim and Virginia Vassilevska Williams. 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*, pages 561–567, 2015. URL: <http://ijcai.org/Abstract/15/085>.
- 10 G. Laffond, J.F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993. doi:10.1006/game.1993.1010.
- 11 J.F. Laslier. *Tournament Solutions and Majority Voting*. Studies in Economic Theory (Berlin, Germany), 7. Springer, 1997. URL: <https://books.google.com/books?id=vYbGAAAAIAAJ>.
- 12 Nicholas R. Miller. A new solution set for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, 24(1):68–96, 1980. URL: <http://www.jstor.org/stable/2110925>.
- 13 Marc Pauly. Can strategizing in round-robin subtournaments be avoided? *Social Choice and Welfare*, 43(1):29–46, 2014. URL: <http://EconPapers.repec.org/RePEc:spr:sochwe:v:43:y:2014:i:1:p:29-46>.
- 14 Jon Schneider, Ariel Schwartzman, and S. Matthew Weinberg. Condorcet-consistent and approximately strategyproof tournament rules. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, pages 35:1–35:20, 2017. doi:10.4230/LIPIcs.ITCS.2017.35.

- 15 Ariel Schwartzman, S. Matthew Weinberg, Eitan Zlatin, and Albert Zuo. Approximately strategyproof tournament rules: On large manipulating sets and cover-consistence. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPICS*, pages 3:1–3:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs. ITCS.2020.3.
- 16 Kenneth A. Shepsle and Barry R. Weingast. Uncovered sets and sophisticated voting outcomes with implications for agenda institutions. *American Journal of Political Science*, 28(1):49–74, 1984. URL: <http://www.jstor.org/stable/2110787>.
- 17 Isabelle Stanton and Virginia Vassilevska Williams. 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*, pages 357–364, 2011. doi:10.5591/978-1-57735-516-8/IJCAI11-069.
- 18 Virginia Vassilevska Williams. Fixing a tournament. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*, 2010. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1726>.

A Omitted Proofs from Section 2

Proof of Proposition 9. The second statement of the lemma clearly follows from the first. It is trivial to see that $\alpha_k^r(\mathcal{T}^\varepsilon) \geq \alpha_k^r(\mathcal{S}^\varepsilon)$, as $\mathcal{S}^\varepsilon \subseteq \mathcal{T}^\varepsilon$. To get intuition for why $\alpha_k^r(\mathcal{T}) \leq \alpha_k^r(\mathcal{S}^\varepsilon)$, observe that every probabilistic tournament in \mathcal{T}^ε can be written as a convex combination of probabilistic tournaments in \mathcal{S}^ε . Indeed, this intuition drives the proof.

To see this formally, let T be an arbitrary (independent probabilistic) tournament in \mathcal{T}^ε , and let T' be any S -adjacent (independent probabilistic) tournament for some $|S| \leq k$. Then $p_{ij}^T \in [1/2 - \varepsilon, 1/2 + \varepsilon]$ for all i, j , by definition. Consider the following procedure to jointly sample tournaments from T, T' . The procedure first constructs another pair of probabilistic tournaments U, U' , and then samples from these.

1. For all (i, j) , set $q_{ij} := \frac{p_{ij}^T - (1/2 - \varepsilon)}{2\varepsilon}$. Observe that $q_{ij} \cdot (1/2 + \varepsilon) + (1 - q_{ij}) \cdot (1/2 - \varepsilon) = p_{ij}^T$, and that $q_{ij} \in [0, 1]$ as $p_{ij}^T \in [1/2 - \varepsilon, 1/2 + \varepsilon]$.
2. For each (i, j) , independently, set p_{ij} equal to $(1/2 + \varepsilon)$ with probability q_{ij} and equal to $(1/2 - \varepsilon)$ with probability $(1 - q_{ij})$.
3. For all (i, j) , set $p_{ij}^U := p_{ij}$.
4. For all (i, j) such that $\{i, j\} \not\subseteq S$, set $p_{ij}^{U'} := p_{ij}^U$. For other (i, j) , set $p_{ij}^{U'} := p_{ij}^{T'}$.
5. Draw a pair of tournaments according to U, U' (independently, say).

Observe that this procedure correctly draws two tournament according to T, T' . Indeed, it is easy to see for each output tournament that the outcome of each match is independent, simply because they are independent in U (resp. U'), and because the random variables p_{ij}^U (resp. $p_{ij}^{U'}$) are independent. Moreover, we claim that the probability that i beats j is exactly p_{ij}^T . Indeed, the probability that i beats j in U is: $\mathbb{E}[p_{ij}^U] = q_{ij} \cdot (1/2 + \varepsilon) + (1 - q_{ij}) \cdot (1/2 - \varepsilon) = p_{ij}^T$. In U' , when $\{i, j\} \subseteq S$, the probability that i beats j in U' is clearly $p_{ij}^{T'}$ by definition. When $\{i, j\} \not\subseteq S$, the probability that i beats j is also p_{ij}^T , which is equal to $p_{ij}^{T'}$ as T, T' are S -adjacent. Additionally, observe that: (a) $U \in \mathcal{S}^\varepsilon$ and (b) U and U' are S -adjacent. Now, consider any tournament rule $r(\cdot)$:

$$\begin{aligned} \text{For all } i: r_i(T) &= \mathbb{E}[r_i(U)], \text{ and } r_i(T') = \mathbb{E}[r_i(U')] \\ \Rightarrow r_S(T') - r_S(T) &= \mathbb{E}[r_S(U') - r_S(U)] \leq \alpha_k^r(\mathcal{S}^\varepsilon) \end{aligned}$$

Indeed, the first line is simply linearity of expectation, once the previous work confirms that going through U, U' is a valid way to draw tournaments from T, T' . The third line then also follows by linearity of expectation. The final line follows as $U \in \mathcal{S}^\varepsilon$, and U, U' are S -adjacent. This completes the proof, as we have now shown that $\alpha_k^r(\mathcal{T}^\varepsilon) \leq \alpha_k^r(\mathcal{S}^\varepsilon)$. \blacktriangleleft

B Omitted Proofs from Section 3

Proof of Lemma 16. This follows immediately from linearity of expectation. For all i , the probability that i wins in tournament T under rule $r^E(\cdot)$ is the expected probability that i wins in $T|_M$. Summing over $i \in S$, and repeating this for T' yields the lemma. \blacktriangleleft

C Omitted Proofs from Section 5

Proof of Lemma 22. Let $\sigma(\cdot)$ denote the permutation which swaps u and v . Consider any two deterministic tournaments $U, \sigma(U)$. Then because $r(\cdot)$ is anonymous, we have:

$$\begin{aligned} r_u(U) + r_v(U) &= r_{\sigma(u)}(\sigma(U)) + r_{\sigma(v)}(\sigma(U)) = r_v(\sigma(U)) + r_u(\sigma(U)) \\ \Rightarrow r_S(U) &= r_S(\sigma(U)) \end{aligned}$$

Indeed, the first line simply applies anonymity, and the second line simply applies σ . Now let's return to T, T' (which are independent probabilistic tournaments, rather than deterministic). Consider the following process to draw T and T' jointly:

1. To emphasize that $p_{uw}^T = p_{vw}^T$, for all $w \notin \{u, v\}$, denote by $p_w^T := p_{uw}^T$.
2. Without loss of generality, let $p_{uv}^T \leq p_{uv}^{T'}$.
3. Draw the outcome of all matches involving two teams both $\notin \{u, v\}$. Set the outcome of these matches the same for T and T' .
4. For all $w \notin \{u, v\}$, draw q_{w1} and q_{w2} iid and uniformly from $[0, 1]$. Draw q_{uv} independently and uniformly from $[0, 1]$.
5. For all $w \notin \{u, v\}$, have u beat w if and only if $q_{w1} < p_w^T$. Have v beat w if and only if $q_{w2} < p_w^T$. Have u beat v if and only if $q_{uv} < p_{uv}^T$.
6. If $q_{uv} \notin [p_{uv}^T, p_{uv}^{T'}]$, set $T' := T$.
7. If $q_{uv} \in [p_{uv}^T, p_{uv}^{T'}]$, then set $T' := \sigma(T)$.

This process clearly satisfies that with probability one, either $T' = T$ or $T' = \sigma(T)$. By the work above, this means that $r_S(T) = r_S(T')$ as desired, as long as we confirm that this process validly samples both T and T' . It is easy to see that the process is valid for T : the match outcomes are clearly independent, and any team w beats x if and only if a uniformly random draw from $[0, 1]$ is $< p_{wx}^T$ (which happens with probability p_{wx}^T , as desired).

To see that the process is valid for T' , observe first that u beats v with probability exactly $p_{uv}^{T'}$, because u beats v whenever $q_{uv} < p_{uv}^{T'}$. Moreover, after conditioning on q_{uv} , the outcome of each (u, w) match is either set according to q_{w1} (an independent, uniform draw from $[0, 1]$), or q_{w2} (also an independent, uniform draw from $[0, 1]$). In either case, these match results are all set independently and with the correct probability (because T, T' are S -adjacent, and because q_{w1}, q_{w2} are iid because $p_{uw}^T = p_{vw}^T$ for all w). This completes the proof. \blacktriangleleft

Proof of Theorem 24. We again simply let $c = 1$ and take the trivial bound on the gain in bad terminal events.

To bound the probability of a bad terminal event, observe that a bad terminal event occurs only when u plays v this round, which happens with probability exactly $1/(n' - 1)$.¹¹ So RSEB satisfies $b = 1/(n' - 1)$.

For the good terminal events, we claim that $\Pr[G] \geq 2/(n' - 1)$. First, observe that once we show this, we can plug into Theorem 21 and conclude $\alpha_2^{\text{RSEB}}(\mathcal{T}^{1/2}) \leq \frac{1 \cdot 1/(n' - 1)}{1/(n' - 1) + 2/(n' - 1)} = 1/3$.

¹¹ Recall that in RSEB, $n' := 2^{\lceil \log_2(n) \rceil}$.

To establish this bound, let ℓ_u denote the number of teams which beat u but not v , ℓ_v denote the number of teams which beat v but not u . Without loss of generality let $\ell_u \geq \ell_v$.

If $\ell_u + \ell_v = 0$, then Lemma 22 already establishes that u and v gain nothing by manipulating. If $\ell_u \geq 2$, then whenever u plays a team which beats them, we have a good terminal event (because u is already eliminated before having ever played the (u, v) match, so manipulating the match has no impact). This happens with probability at least $2/(n' - 1)$.

If $\ell_u = 1, \ell_v = 0$, then let w denote the unique team which beats u but not v . We now claim that if w plays *either* u or v that we are in a good terminal event. Indeed, if w plays u , then u is eliminated having never played the (u, v) match. If instead w plays v , then w is eliminated, but now there are no remaining teams which beat u but not v (or vice versa) and Lemma 22 asserts that there are no further gains from manipulation. The probability that w plays u or v is $2/(n' - 1)$, as desired.

Finally, if $\ell_u = \ell_v = 1$, let w denote the unique team which beats u but not v , and x denote the unique team which beats v but not u (this case requires slightly more calculations than RDM). If u plays w , *or* v plays x , then at least one of u, v is eliminated before the (u, v) match is played, and therefore no gains are possible. This happens with probability $2/(n' - 1) - \frac{1}{(n' - 1)(n' - 3)}$. In addition, if u plays x *and* v plays w , then both x and w will be eliminated, and Lemma 22 asserts that there are no further gains from manipulation. This occurs with probability $\frac{1}{(n' - 1)(n' - 3)}$. Therefore, a good terminal event happens with probability at least $2/(n' - 1)$, as desired. This handles all possible cases, and establishes that $g \geq 2/(n' - 1)$ in all cases. Plugging into Theorem 21 as described completes the proof. \blacktriangleleft

D Omitted Proofs from Section 6

Proof of Lemma 25. Consider a 3-player tournament T with teams u , v , and w , $p_{uv}^T = p_{vw}^T = p_{wu}^T = \frac{1}{2} + \varepsilon$.¹² Noting that T is randomized, there are two possible types of deterministic outcomes: outcomes where there is a Condorcet winner, and outcomes which form a cycle (either v beats u , u beats w , and w beats v , or vice versa).

In T , each of the three players has probability $(\frac{1}{2} + \varepsilon)(\frac{1}{2} - \varepsilon)$ of being a Condorcet winner.

Call the cycle where v beats u “cycle 1”; this occurs with probability $(\frac{1}{2} + \varepsilon)^3$. Call the opposing cycle (where u beats v) “cycle 2”; this occurs with probability $(\frac{1}{2} - \varepsilon)^3$.

Let r be any Condorcet-consistent tournament rule. Denote by γ_x the probability that the rule selects x as the winner when cycle 1 occurs (for any $x \in \{u, v, w\}$). Denote by β_x the probability that the rule selects x as the winner when cycle 2 occurs. Recall that r must select x as the winner with probability 1 when x is the Condorcet winner.

Suppose that u and v collude so that u throws their match to v . Specifically, let T^{uv} denote the $\{u, v\}$ -adjacent tournament to T where $p_{uv}^{T^{uv}} = 0$ (instead of $1/2 + \varepsilon$). In T^{uv} , v is a Condorcet winner with probability $\frac{1}{2} + \varepsilon$ (they need only beat w), cycle 2 occurs with probability $(\frac{1}{2} - \varepsilon)^2$ (v must lose to w , who must lose to u), and w is a Condorcet winner with probability $(\frac{1}{2} + \varepsilon)(\frac{1}{2} - \varepsilon)$. No other outcomes are possible. We then have:

$$\begin{aligned} r_u(T^{uv}) + r_v(T^{uv}) - r_u(T) - r_v(T) &= (\frac{1}{2} + \varepsilon) + (\beta_u + \beta_v) \cdot (\frac{1}{2} - \varepsilon)^2 \\ &\quad - 2(\frac{1}{2} + \varepsilon)(\frac{1}{2} - \varepsilon) - (\beta_u + \beta_v) \cdot (\frac{1}{2} - \varepsilon)^3 - (\gamma_u + \gamma_v)(\frac{1}{2} + \varepsilon)^3 \\ &= 2\varepsilon \cdot (\frac{1}{2} + \varepsilon) + (\beta_u + \beta_v) \cdot ((\frac{1}{2} - \varepsilon)^2 - (\frac{1}{2} - \varepsilon)^3) - (\gamma_u + \gamma_v)(\frac{1}{2} + \varepsilon)^3 \end{aligned}$$

¹²That is, if $\varepsilon = 1/2$, this is the same 3-cycle example from [14].

14:20 Approximately Strategyproof Tournament Rules in the Probabilistic Setting

Identical calculations hold for T^{vw}, T^{wu} . We can sum these three together to achieve:

$$\begin{aligned}
& (r_u(T^{uv}) + r_v(T^{uv}) - r_u(T) - r_v(T)) + (r_v(T^{vw}) + r_w(T^{vw}) - r_v(T) - r_w(T)) \\
& \quad + (r_w(T^{wu}) + r_u(T^{wu}) - r_w(T) - r_u(T)) \\
& = 6\varepsilon \cdot \left(\frac{1}{2} + \varepsilon\right) + 2(\beta_u + \beta_v + \beta_w) \cdot \left(\left(\frac{1}{2} - \varepsilon\right)^2 - \left(\frac{1}{2} - \varepsilon\right)^3\right) - 2(\gamma_u + \gamma_v + \gamma_w) \cdot \left(\frac{1}{2} + \varepsilon\right)^3 \\
& = 6\varepsilon \cdot \left(\frac{1}{2} + \varepsilon\right) + 2 \cdot \left(\frac{1}{2} - \varepsilon\right)^2 \cdot \left(\frac{1}{2} + \varepsilon\right) - 2 \cdot \left(\frac{1}{2} + \varepsilon\right)^3 \\
& = 6\varepsilon \cdot \left(\frac{1}{2} + \varepsilon\right) + 2 \cdot \left(\frac{1}{2} + \varepsilon\right) \cdot \left(\left(\frac{1}{2} - \varepsilon\right)^2 - \left(\frac{1}{2} + \varepsilon\right)^2\right) \\
& = 6\varepsilon \cdot \left(\frac{1}{2} + \varepsilon\right) + 2 \cdot \left(\frac{1}{2} + \varepsilon\right) \cdot (-2\varepsilon) = 2\varepsilon \cdot \left(\frac{1}{2} + \varepsilon\right)
\end{aligned}$$

This means that one of the three coalitions can gain at least $\frac{2}{3}\varepsilon\left(\frac{1}{2} + \varepsilon\right) = \frac{1}{3}\varepsilon + \frac{2}{3}\varepsilon^2$. \blacktriangleleft