

Robust Draws in Balanced Knockout Tournaments

Krishnendu Chatterjee[†]

Rasmus Ibsen-Jensen[†]

Josef Tkadlec[†]

[†] IST Austria

{krish.chat, ribsens, jtkadlec}@ist.ac.at

Abstract

Balanced knockout tournaments are ubiquitous in sports competitions and are also used in decision-making and elections. The traditional computational question, that asks to compute a draw (optimal draw) that maximizes the winning probability for a distinguished player, has received a lot of attention. Previous works consider the problem where the pairwise winning probabilities are known precisely, while we study how robust is the winning probability with respect to small errors in the pairwise winning probabilities. First, we present several illuminating examples to establish: (a) there exist deterministic tournaments (where the pairwise winning probabilities are 0 or 1) where one optimal draw is much more robust than the other; and (b) in general, there exist tournaments with slightly suboptimal draws that are more robust than all the optimal draws. The above examples motivate the study of the computational problem of robust draws that guarantee a specified winning probability. Second, we present a polynomial-time algorithm for approximating the robustness of a draw for sufficiently small errors in pairwise winning probabilities, and obtain that the stated computational problem is NP-complete. We also show that two natural cases of deterministic tournaments where the optimal draw could be computed in polynomial time also admit polynomial-time algorithms to compute robust optimal draws.

1 Introduction

Balanced knockout tournaments (BKTs). A Balanced knockout tournament (BKT) consists of $N = 2^n$ players (for some positive integer n) and is played in n rounds. In every round each remaining player plays a (win/lose) game with some other remaining player and the loser is eliminated (removed from the tournament). In the end, only one player remains and is declared the winner. The whole process can be visualised on a balanced binary tree (over N leaves) with players starting at the leaves and winners advancing to the parent. Such tournaments are used in many sports such as tennis, elections and other decision making processes.

Draws in BKT. For a given set of N players, the tournament can be played in many different ways based on how we draw the players in the first round (i.e., on which leaf they are placed in the balanced binary tree). Even if we consider the draws that correspond to isomorphic labeled binary trees as equivalent, the number of draws grows rapidly in N , precisely, there are $\frac{N!}{2^{N-1}}$ draws. As a numerical example, for a Grand Slam tournament in tennis with $N = 128$, the number of distinct draws is at least 10^{177} and even if we rank the players and require that, for any k , the top 2^k of them do not meet until the last k rounds, we have at least 10^{144} distinct draws.

Computational problems. The traditional computational problem that has been extensively studied is related to obtaining draws that are most favorable for a distinguished player [Vu *et al.*, 2009; Aziz *et al.*, 2014; Stanton and Williams, 2011a]. The input consists of an $N \times N$ comparison matrix P that specifies the probabilities P_{ij} of player i beating player j in a single match, and a distinguished player i^* . The input, along with a draw, defines the winning probability that the distinguished player wins the whole tournament. If the comparisons matrix contains 0's and 1's only, a draw gives a unique winner. The *probabilistic (resp. deterministic) tournament fixing problem* (PTFP) (resp. TFP) asks whether there exists a draw such that the winning probability for the distinguished player i^* is at least q (resp., player i^* is the winner). The TFP and PTFP problems are desired to be hard, as otherwise the tournament could be manipulated by choosing the draw that favors a specific player.

Previous results. It was shown in [Vu *et al.*, 2009] that given an input comparison matrix and a draw, the winning probability for a player can be computed by a recursive procedure in $O(N^2)$ time (i.e., in polynomial time). Since a candidate draw for the PTFP problem is a polynomial witness, an NP upper bound follows for both the PTFP and the TFP problem. An NP lower bound for the PTFP problem was shown in [Vu *et al.*, 2009; Williams, 2010], and finally the TFP problem was shown to be NP hard in [Aziz *et al.*, 2014]. Hence both the PTFP and the TFP problems are NP-complete. Moreover, in [Aziz *et al.*, 2014] two important special cases have been identified (namely, constant number of player types, and linear ordering among players with constant number of exceptions) where all the winning draws for player i^* can be computed in polynomial time for deterministic tournaments.

Robustness question. The previous works focused on the computational problems when the pairwise winning probabilities are precisely known. However, in most practical scenarios, the winning probabilities in the comparison matrix are only approximation of the real probabilities. Indeed, in all typical applications, either the probabilities are obtained from past samples (such as games played) or uniformly selected from a small subset of samples (such as in elections). The probabilities obtained in these ways are always at best an approximation of the real probabilities and subject to small errors. This leads to the natural question about robustness (or sensitivity) of optimal draws for probabilistic tournaments (or winning draws in deterministic tournaments) in the presence of small errors in the pairwise winning probabilities. Formally, given a comparison matrix P , and a small error term ε , we consider *all* comparison matrices P' where each entry differs from P by at most ε . We refer to the above set as the ε -perturbation matrices of P . For a draw, we consider the *drop* of the draw as the difference between the winning probability for the distinguished player i^* for P and the infimum of the winning probability of the ε -perturbation matrices.

Our results. We study the computational problems related to robust draws in TFP and PTFP. Our main contributions are:

1. *Examples.* We present several illuminating examples for TFP and PTFP related to robustness. First, we show that there exist deterministic tournaments where for one winning (or optimal draw) the drop is about $\frac{N \cdot \varepsilon}{2}$, whereas for another winning draw the drop is about $\log N \cdot \varepsilon$. In other words, one winning draw is much more robust than the other. Second, we show that there exist probabilistic tournaments with suboptimal draws that are more robust than all the optimal draws. Motivated by the above examples we study the computational problem of robust draws, i.e. determining the existence of draws that guarantee a specified winning probability with a drop below a specified threshold.
2. *Algorithm.* We present a polynomial-time algorithm to approximate the drop for a given draw for small $\varepsilon > 0$ (informally, for ε where higher order terms of ε such as $\varepsilon^2, \varepsilon^3$ etc. can be ignored).
3. *Consequences.* Our algorithm has a number of consequences. First, it establishes that the computational problem of robust draws for small ε is NP-complete. Note that while the PTFP and TFP are existential questions (existence of a draw), the robustness question has a quantifier alternation (existence of a draw such that for all ε -perturbation matrices the drop is small), yet we match the complexity of the PTFP and TFP problem. Second, our polynomial-time algorithm along with the result of [Aziz et al., 2014] implies that for the two natural cases of [Aziz et al., 2014], the most robust winning draw (if one exists) is polynomial-time computable.

Detailed proofs omitted due to lack of space is available in [Chatterjee et al., 2016].

Significance as risk-averse strategies. As mentioned, BKTs are often used as sports tournaments. After a draw is fixed, the winning probabilities determine the betting odds. Since the

comparison matrix can only be approximated, a risk-averse strategy (as typically employed by humans) corresponds to the notion of robustness. The notion of robustness has been studied in many different contexts, such as for sensitivity analysis in MDPs [Puterman, 1994; Filar and Vrieze, 1997] as well as for decision making in markets [Rockafellar, 2007]. Our algorithm provides a risk-averse approximation for balanced knockout tournaments, for low levels of uncertainty in the probabilities.

1.1 Related Work

The most related previous works are: (a) [Vu et al., 2009], who showed that for a fixed draw the probability distribution over the winners can be computed in $O(N^2)$ time; and (b) [Aziz et al., 2014], who determined the complexity of TFP and found special cases with polynomial-time algorithms. Besides that, [Williams, 2010] identified various sufficient conditions for a player to be a winner of a BKT; [Stanton and Williams, 2011b] considered the case when weak players can possibly win a BKT; [Stanton and Williams, 2011a] studied the conditions under which the tournament can be fixed with high probability. The problem of fair draws was studied in [Vu and Shoham, 2011]; and the problem of determining the winner in unbalanced voting trees was considered in [Lang et al., 2007; 2012]. If only incomplete information on the preferences is known, then computing the winner with various voting rules has been studied in [Xia and Conitzer, 2011; Aziz et al., 2012]. The problem of checking whether a round-robin competition can be won when all the matches are not yet played has been studied in [Kern and Paulusma, 2004; Gusfield and Martel, 2002].

As compared to the existing works we consider the problem of robustness for TFP and PTFP which has not been studied before.

2 Preliminaries

In this section we present the formal definitions and previous results.

2.1 Definitions

For the basic definitions we very closely follow the notations of [Aziz et al., 2014].

Definition 1 (Comparison matrices). For $N \in \mathbb{N}$, let $[N] := \{1, \dots, N\}$. Consider N players numbered 1 to N . A *comparison matrix* is an $N \times N$ matrix P such that for all $1 \leq i \neq j \leq N$ we have $0 \leq P_{ij} \leq 1$ and $P_{ij} + P_{ji} = 1$. The entry P_{ij} of the comparison matrix expresses the probability that if players i, j play a match, player i wins. Note that the entries P_{ii} are not defined. A comparison matrix is *deterministic* if $P_{ij} \in \{0, 1\}$ for all $i, j \in [N]$ with $i \neq j$.

Definition 2 (Draws). Let $N = 2^n$ for some integer n . For σ a permutation of $[N]$, an *n-round ordered balanced knockout tournament* $T([N], \sigma)$ is a binary tree with N leaf nodes labelled from left to right by σ . All ordered balanced knockout tournaments that are isomorphic to each other are said to have the same *draw*. They are represented by a single (unordered) *balanced knockout tournament* $T([N], \sigma)$, where σ is again a permutation of $[N]$. The set of all draws is denoted by Σ .




Figure 1: A 2-round complete tournament with comparison matrix P . The draw σ gives $\text{wp}(1, P, \sigma) = 0.9^3$.

Definition 3 (Complete tournaments). A *complete tournament* $C(P, \sigma)$ is a balanced knockout tournament $T([N], \sigma)$ together with an $N \times N$ comparison matrix P . A complete tournament $C(P, \sigma)$ is called *deterministic* if the comparison matrix P is deterministic. The complete tournament $C(P, \sigma)$ is conducted in the following fashion. If two nodes with labels i, j have the same parent in $T([N], \sigma)$ then players i, j play a *match*. The winner then labels the parent (i.e. the parent is labeled by i with probability P_{ij} and by j otherwise). The *winner* of $C(P, \sigma)$ is the player who labels the root node.

Definition 4 (Winning probabilities). Given a complete tournament $C(P, \sigma)$, each player $i \in [N]$ has a probability, denoted $\text{wp}(i, P, \sigma)$, of being the winner of $C(P, \sigma)$. This probability can be computed in time $O(N^2)$ via a recursive formulation [Vu *et al.*, 2009]. We denote by $\text{mwp}(i, P) := \max_{\sigma \in \Sigma} \{\text{wp}(i, P, \sigma)\}$ the maximum possible winning probability of i in $C(P, \sigma)$ taken over all draws $\sigma \in \Sigma$. Given P , and $\delta \geq 0$, a draw σ is called δ -*optimal* for player i provided that $\text{wp}(i, P, \sigma) \geq \text{mwp}(i, P) - \delta$. A draw is *optimal* if it is δ -optimal for $\delta = 0$.

Example 1. Consider a complete 2-round tournament with comparison matrix P and draw $\sigma = (1, 2, 3, 4)$ as in Figure 1. Then the winning probabilities are

$$\begin{aligned} \text{wp}(1, P, \sigma) &= 0.729, & \text{wp}(2, P, \sigma) &= 0, \\ \text{wp}(3, P, \sigma) &= 0.171, & \text{wp}(4, P, \sigma) &= 0.1 \end{aligned}$$

and the draw σ is optimal for player 1, as the draws $\sigma' = (1, 3, 2, 4)$, $\sigma'' = (1, 4, 2, 3)$ give $\text{wp}(1, P, \sigma') = \text{wp}(1, P, \sigma'') = 0$.

2.2 Previous Results

We now describe the main computational problems and previous results related to them.

Problem 1 (TFP – Tournament Fixing Problem). Given a player set N , a deterministic comparison matrix P , and a distinguished player $i^* \in N$, does there exist a draw σ for the player set N for which i^* is the winner of $C(P, \sigma)$?

Problem 2 (PTFP – Probabilistic TFP). Given a player set $[N]$, a comparison matrix P , a distinguished player $i^* \in [N]$ and target probability $q \in [0, 1]$, does there exist a draw σ for the player set $[N]$ such that $\text{wp}(i^*, P, \sigma) \geq q$?

Theorem 1. ([Aziz *et al.*, 2014; Vu *et al.*, 2009]) The TFP and the PTFP problems are NP-complete.

3 Robustness: Examples and Questions

In practical applications of PTFP and TFP, the exact values of the entries in the comparison matrix are not known, but only approximations are obtained. Hence it is of interest to find out how sensitive is the winning chance in optimal draws (or near optimal draws) with respect to minor changes in the comparison matrix. We first present the basic definition and then our examples. Finally, we present the computational problems.

Definition 5 (ε -perturbation and the ε -worst drop). Given a comparison matrix P and $\varepsilon > 0$, an ε -perturbation of P is any comparison matrix P' such that $|P_{ij} - P'_{ij}| \leq \varepsilon$ for all $1 \leq i \neq j \leq N$. The set of all ε -perturbations of P is denoted $\mathcal{P}(P, \varepsilon)$. Given a complete tournament $C(P, \sigma)$, a distinguished player i^* , and $\varepsilon > 0$, define the ε -*guaranteed winning probability* $\text{wp}_\varepsilon(i^*, P, \sigma)$ by

$$\text{wp}_\varepsilon(i^*, P, \sigma) = \inf_{P' \in \mathcal{P}(P, \varepsilon)} \{\text{wp}(i^*, P', \sigma)\}$$

and the ε -*worst drop*, denoted $d_\varepsilon(i^*, P, \sigma)$, by $\text{wp}(i^*, P, \sigma) - \text{wp}_\varepsilon(i^*, P, \sigma)$. The smaller the drop $d_\varepsilon(i^*, P, \sigma)$ the more robust (or more risk-averse) is the draw σ .

3.1 Examples

We consider the robustness problem for sufficiently small $\varepsilon > 0$. Intuitively, “sufficiently small” will allow us to ignore all higher order terms of ε such as $\varepsilon^2, \varepsilon^3, \dots$. A formal definition comes at the end of this section.

First we construct a deterministic tournament with a unique and nonrobust optimal draw. We will use Proposition 1 in Proposition 2. For brevity, the distinguished player will always be the first one, i.e. $i^* = 1$.

Proposition 1 (Only Nonrobust Optimal Draws). For any $2^n = N \in \mathbb{N}$, there exists a deterministic n -round tournament (called a *hard n -round tournament* and denoted \mathcal{H}_n) with a comparison matrix P such that for every draw σ which makes player 1 win we have

$$\text{wp}_\varepsilon(1, P, \sigma) \leq 1 - (N - 1)\varepsilon + \varepsilon^2 \cdot Q(\varepsilon)$$

for some integer polynomial $Q(\varepsilon)$.

Intuitively, the polynomial Q stores higher order terms (from ε^2 on). The proof uses a construction from [Aziz *et al.*, 2014, Lemma 1]. The tournament, illustrated in Figure 2, has the property that exactly one draw makes player 1 win and if any single match changes outcome then player 1 loses. All these facts follow by easy induction arguments.

Next we present a concrete numerical example to illustrate Proposition 1.

Example 2. There exists a deterministic 6-round tournament with a comparison matrix P such that any draw σ , which makes player 1 win, satisfies

$$\begin{aligned} \text{wp}_{0.01}(1, P, \sigma) &< 0.54, & \text{wp}_{0.05}(1, P, \sigma) &< 0.07, \\ \text{wp}_{0.1}(1, P, \sigma) &< 0.02. \end{aligned}$$

By Proposition 1, the 6-round hard tournament with comparison matrix P has only one draw $\sigma = (1, \dots, 64)$ that makes player 1 win. Let P' be the ε -perturbation where we adjusted every $(0,1)$ match into an $(\varepsilon, 1 - \varepsilon)$ match. We recursively




Figure 2: A hard 4-round tournament \mathcal{H}_4 . Each broken line connects nodes that are eventually labelled by the same player. Except for the depicted exceptions, each player beats all players to their left.




Figure 3: An unbalanced 4-round tournament \mathcal{U}_4 . The players in the right subtree beat everyone from the left subtree but the player 1.

compute $\text{wp}(1, P', \sigma) = 1 - 63\varepsilon + \dots - 2\,147\,483\,648\varepsilon^{63}$ and plug in $\varepsilon \in \{0.01, 0.05, 0.1\}$ to get numbers less than 0.54, 0.07, 0.02, respectively.

We now use Proposition 1 to show that there exist deterministic tournaments where for one winning draw the drop is approximately $\frac{N\varepsilon}{2}$, whereas for a different winning draw the drop is approximately $\varepsilon(2 \log N - 1)$, for sufficiently small ε . In other words, some winning draws can be much more robust than other winning draws.

Proposition 2 (Robust and Nonrobust Optimal Draws). For any $2^n = N \in \mathbb{N}$, $n \geq 2$, there exists a deterministic n -round tournament (called an *unbalanced n -round tournament* and denoted by \mathcal{U}_n) with a comparison matrix P and draws σ, σ' such that $\text{wp}(1, P, \sigma) = 1 = \text{wp}(1, P, \sigma')$ and

$$\begin{aligned} \text{wp}_\varepsilon(1, P, \sigma) &\leq 1 - \varepsilon \cdot N/2 + \varepsilon^2 \cdot Q(\varepsilon), \\ \text{wp}_\varepsilon(1, P, \sigma') &> 1 - \varepsilon \cdot (2n - 1) \end{aligned}$$

for some integer polynomial $Q(\varepsilon)$.

Proof. Take an $(n - 1)$ -round hard tournament with players $1, 2, \dots, 2^{n-1}$. Add 2^{n-1} more players numbered $2^{n-1} + 1$ through 2^n who all lose to player 1 but beat all the players 2 through 2^{n-1} (see Figure 3). Then we claim that $\sigma = (1, 2, \dots, 2^n)$ and $\sigma' = (2^n, 2, 3, \dots, 2^n - 1, 1)$ fulfil the statement of the proposition.

Clearly, player 1 wins both $C(P, \sigma)$ and $C(P, \sigma')$.

To upper bound $\text{wp}_\varepsilon(1, P, \sigma)$ it is enough to find one ε -perturbation that reduces the winning chance of player 1 sufficiently. Let P' again be the comparison matrix obtained by

replacing each $(0, 1)$ match by an $(\varepsilon, 1 - \varepsilon)$ match. By Proposition 1, the probability of player 1 winning her first $n - 1$ rounds is $p_{n-1} = 1 - (2^{n-1} - 1)\varepsilon + \varepsilon^2 Q'(\varepsilon)$ for some polynomial $Q'(\varepsilon)$. The probability of winning the final is then always $1 - \varepsilon$ which altogether gives

$$\begin{aligned} \text{wp}_\varepsilon(1, P, \sigma) &\leq [1 - (2^{n-1} - 1)\varepsilon + \varepsilon^2 Q'(\varepsilon)] \cdot (1 - \varepsilon) \\ &= 1 - 2^{n-1}\varepsilon + \varepsilon^2 Q(\varepsilon) \end{aligned}$$

for some polynomial $Q(\varepsilon)$.

For $\text{wp}_\varepsilon(1, P, \sigma')$, no matter which $P'' \in \mathcal{P}(P, \varepsilon)$ we take, players 2^n and 1 win their first $n - 1$ matches with probability at least $(1 - \varepsilon)^{n-1}$ each and then player 1 beats player 2^n with probability at least $1 - \varepsilon$ again, giving

$$\text{wp}_\varepsilon(1, P, \sigma') \geq (1 - \varepsilon)^{2n-1} > 1 - (2n - 1)\varepsilon. \quad \square$$

We now illustrate Proposition 2 with a numerical example.

Example 3. There exists a deterministic 6-round tournament with a comparison matrix P and draws σ, σ' such that $\text{wp}(1, P, \sigma) = 1 = \text{wp}(1, P, \sigma')$ and

$$\begin{aligned} \text{wp}_{0.01}(1, P, \sigma) &< 0.73 < 0.89 < \text{wp}_{0.01}(1, P, \sigma'), \\ \text{wp}_{0.05}(1, P, \sigma) &< 0.23 < 0.56 < \text{wp}_{0.05}(1, P, \sigma'), \\ \text{wp}_{0.1}(1, P, \sigma) &< 0.08 < 0.31 < \text{wp}_{0.1}(1, P, \sigma'). \end{aligned}$$

Indeed, take the unbalanced 6-round tournament with comparison matrix P and the draws $\sigma = (1, 2, \dots, 64)$, $\sigma' = (64, 2, 3, \dots, 63, 1)$. By Proposition 2, player one wins both $C(P, \sigma)$ and $C(P, \sigma')$. Consider the ε -perturbation P' of P that replaces each $(0, 1)$ match by an $(\varepsilon, 1 - \varepsilon)$ match. Then for any $\varepsilon > 0$ we have

$$\begin{aligned} \text{wp}_\varepsilon(1, P, \sigma) &\leq \text{wp}(1, P', \sigma) = p_5 \cdot (1 - \varepsilon) = \\ &= 1 - 64\varepsilon + \dots + 2\,147\,483\,648\varepsilon^{64} \end{aligned}$$

which, after plugging in $\varepsilon \in \{0.01, 0.05, 0.1\}$, gives the desired inequalities.

On the other hand, the proof of the Proposition 2 implies that for any $\varepsilon > 0$ we have $\text{wp}_\varepsilon(1, P, \sigma') \geq (1 - \varepsilon)^{11}$ which gives the remaining inequalities.

Probabilistic setting. Proposition 1 and Proposition 2 concern the deterministic setting. In a deterministic setting, if a draw is not optimal (i.e., winning), then the player loses with probability 1 (hence the notion of δ -optimality, for $\delta > 0$ is not relevant). We now consider the probabilistic setting and show that there are combinations of δ and ε such that a certain δ -optimal draw gives a better ε -guarantee than any optimal draw. In other words, near optimal draws can be more robust than all optimal draws.

Proposition 3. There exists a tournament with comparison matrix P and a distinguished player i and there exist $\delta, \varepsilon > 0$ and a δ -optimal draw σ' such that for all optimal draws σ we have $\text{wp}_\varepsilon(i, P, \sigma) < \text{wp}_\varepsilon(i, P, \sigma')$.

For proving Proposition 3 we use an auxiliary lemma.

Let a *big vs. small tournament* with players $[2^n]$ have the comparison matrix P such that each of the players in $\{2^{n-1} + 1, \dots, 2^n\}$ (“big” players) beats each of the players in $[2^{n-1}]$ (“small” players) with the same probability $p \in (0.5, 1)$.

Lemma 1. The draws that maximize the probability of a big player winning the big vs. small tournament (the *big-optimal* draws) are precisely those that pair up big players and small players in the first round (so-called *mixed* draws). Moreover, for $p = 0.5 + \varepsilon$, the probability $b(P)$ of a big player winning in such a draw equals $b(P) = 0.5 + \frac{1}{2}(n+1)\varepsilon + \varepsilon^2 \cdot Q(\varepsilon)$ for some polynomial $Q(\varepsilon)$.

Proof. The proof of the lemma is technical but straightforward induction argument. \square

Proof. (of Proposition 3) Consider an $(n+1)$ -round tournament with players $[2^{n+1}]$. Call players $\{2, 3, \dots, 2^n\}$ *small*, players $2^n + 1, \dots, 2^n + 2^{n-1}$ *medium* and players $2^n + 2^{n-1} + 1, \dots, 2^{n+1}$ *big*. For s small, m medium, and b big, let the comparison matrix P satisfy:

- $P_{1s} = 1, P_{1m} = 0.4, P_{1b} = 0.6,$
- $P_{sm} = P_{sb} = 0,$
- $P_{mb} = 0.5 - \frac{1}{2}\varepsilon.$

We claim that:

1. The draws optimal for player 1 are those that have all the small players in one half and the medium players paired up with the big players in the other half.
2. For any optimal draw σ , the draw $\sigma' = (1, \dots, 2^n)$ satisfies $\text{wp}_\varepsilon(1, P, \sigma) < \text{wp}_\varepsilon(1, P, \sigma')$ for all sufficiently small $\varepsilon > 0$.

Indeed:

1. Let σ be a draw of the desired form. By Lemma 1, we have

$$\begin{aligned} \text{wp}(1, P, \sigma) &= \left[0.5 + \frac{1}{4}(n+1)\varepsilon + \varepsilon^2 Q(\varepsilon) \right] \cdot 0.6 + \\ &\quad \left[0.5 - \frac{1}{4}(n+1)\varepsilon - \varepsilon^2 Q(\varepsilon) \right] \cdot 0.4 \\ &= 0.5 + \frac{1}{20}(n+1)\varepsilon + \varepsilon^2 R(\varepsilon). \end{aligned}$$

for some polynomials $Q(\varepsilon), R(\varepsilon)$.

Since there are $2^n - 1 < 2^n$ small players, player 1 has to face a medium or big player in the final (provided he made it that far). Now let τ be a draw optimal for player 1 and suppose it contains a medium or big player in the left half. Then the player 1 has to face at least one more such player before the final. Hence his winning chance is at most $0.6^2 < \text{wp}(1, P, \sigma)$ for all sufficiently small $\varepsilon > 0$, and τ is not optimal. All the small players are therefore in the left half. For the right half, we want to maximize the probability of a big player winning it. By Lemma 1, this is accomplished by the mixed draws.

2. Let σ be optimal. Note that since σ and σ' only differ in the right subtree, the inequality

$$\text{wp}_\varepsilon(1, P, \sigma) < \text{wp}_\varepsilon(1, P, \sigma')$$

holds if and only if the minimum possible probability $b_\varepsilon(\sigma)$ of a big player winning the right half of $C(P^*, \sigma)$, taken over all $P^* \in \mathcal{P}(P, \varepsilon)$, is less than the minimum possible probability

$b_\varepsilon(\sigma')$ of a big player winning the right half of $C(P', \sigma')$, taken over all $P' \in \mathcal{P}(P, \varepsilon)$.

For σ , consider the ε -perturbation that replaces each $(0.5 - \frac{1}{2}\varepsilon, 0.5 + \frac{1}{2}\varepsilon)$ match by a $(0.5 + \frac{1}{2}\varepsilon, 0.5 - \frac{1}{2}\varepsilon)$ match. Then by Lemma 1 we have

$$b_\varepsilon(\sigma) = 0.5 - \frac{1}{4}(n+1)\varepsilon - \varepsilon^2 Q(\varepsilon)$$

for some polynomial $Q(\varepsilon)$.

On the other hand, for all $P' \in \mathcal{P}(P, \varepsilon)$ a big player certainly wins the right half of the right subtree of $C(P', \sigma')$ (since that subtree only consists of big players), hence

$$b_\varepsilon(\sigma') \geq \left(0.5 + \frac{1}{2}\varepsilon \right) - \varepsilon = 0.5 - \frac{1}{2}\varepsilon.$$

For $n \geq 2$ and all sufficiently small $\varepsilon > 0$ we then get the desired $b_\varepsilon(\sigma) < b_\varepsilon(\sigma')$ and thus in turn $\text{wp}_\varepsilon(1, P, \sigma) < \text{wp}_\varepsilon(1, P, \sigma')$. \square

We again illustrate this with a numerical example.

Example 4. Consider a 6-round tournament with players $[64]$. Call players $\{2, 3, \dots, 32\}$ *small*, players $33, \dots, 48$ *medium* and players $49, \dots, 64$ *big*. For s small, m medium, and b big, let the comparison matrix P satisfy:

- $P_{1s} = 1, P_{1m} = 0.4, P_{1b} = 0.6,$
- $P_{sm} = P_{sb} = 0,$
- $P_{mb} = 0.49.$

Consider $\sigma = (1, 2, \dots, 32, 33, 49, 34, 50, \dots, 48, 64)$ and $\sigma' = (1, \dots, 64)$. Then it is straightforward to compute $\text{mwp}(1, P) = \text{wp}(1, P, \sigma) > 0.506$ and $\text{wp}(1, P, \sigma') = 0.502 < \text{mwp}(1, P)$ so σ' is not optimal for player 1. However, $\text{wp}_{0.02}(1, P, \sigma) < 0.429 < 0.432 < \text{wp}_{0.02}(1, P, \sigma')$. Indeed, for σ , the 0.02-perturbation P^* of P satisfying

- $P_{1s}^* = 0.98, P_{1m}^* = 0.38, P_{1b}^* = 0.58,$
- $P_{sm}^* = P_{sb}^* = 0,$
- $P_{mb}^* = 0.51$

yields $\text{wp}_{0.02}(1, P, \sigma) < \text{wp}(1, P^*, \sigma) < 0.429$.

On the other hand, consider the draw σ' and take any $P' \in \mathcal{P}(P, 0.02)$. Player 1 wins his first five matches with probability at least 0.98^5 . Then he beats every medium (resp., big) player with probability at least 0.38 (resp., 0.58) and the probability that a big player wins the right half is at least $0.51 - 0.02 = 0.49$. Overall, $\text{wp}(1, P', \sigma) \geq 0.98^5 \cdot (0.49 \cdot 0.58 + 0.51 \cdot 0.38) > 0.432$.

3.2 Computational Questions

We extend the TFP and PTFP problem with robustness and we focus on the question of approximating the ε -worst drop $d_\varepsilon(i^*, P, \sigma)$.

Denote by $\xi(P) = \min\{P_{ij} | P_{ij} \neq 0\}$ the smallest nonzero value in the comparison matrix (if P is deterministic then $\xi(P) = 1$). It is straightforward to see that for $\varepsilon \leq \xi$, the ε -guaranteed winning probability $\text{wp}_\varepsilon(i^*, P, \sigma)$ is a polynomial in ε whose constant term is $\text{wp}(i^*, P, \sigma)$. For such ε , our approximation $\widehat{d}_\varepsilon(i^*, P, \sigma)$ of the drop $d_\varepsilon(i^*, P, \sigma)$ is the linear term of this polynomial. Since higher order terms of ε are

ignored, the number \widehat{d}_ε is really only an approximation of d_ε . However, for all $c < N$, if ε is sufficiently small (smaller than cN^{-2}), then \widehat{d}_ε is within $\pm c\varepsilon$ of d_ε , so the approximation is tight. Formally, we consider the following problems.

Problem 3 (RTFP). The Robust Tournament Fixing Problem (RTFP): Given a player set $[N]$, a deterministic comparison matrix P , a distinguished player $i^* \in N$ and $c \in \mathbb{N}$, does there exist a draw σ for the player set N such that (a) player i^* is the winner of $C(P, \sigma)$; and (b) $\widehat{d}_\varepsilon(i^*, P, \sigma) \leq c\varepsilon$, for all $\varepsilon > 0$ sufficiently small?

Problem 4 (RPTFP). The Robust Probabilistic Tournament Fixing Problem (RPTFP): Given a player set $[N]$, a comparison matrix P , a distinguished player $i^* \in [N]$, target probability $q \in [0, 1]$ and $s \in \mathbb{R}$, does there exist a draw σ for the player set $[N]$ such that (a) $\text{wp}(i^*, P, \sigma) \geq q$; and (b) $\widehat{d}_\varepsilon(i^*, P, \sigma) \leq s\varepsilon$, for all $\varepsilon > 0$ sufficiently small?

Note that the drop computation is an infimum over all ε -perturbation matrices, and thus represents a universal quantification. Thus in contrast to the TFP and the PTFP problems which have only existential quantification over draws, the robustness problem has a quantifier alternation of existential and universal quantifiers.

4 Algorithms

The aim of this section is to prove the following theorem.

Theorem 2. Given a comparison matrix P (deterministic or probabilistic), sufficiently small $\varepsilon > 0$, distinguished player i^* and a draw σ , the value $\widehat{d}_\varepsilon(i^*, P, \sigma)$ can be computed in polynomial time.

We start with a lemma stating that if we are allowed to perturb by ε , perturbing by less is not worth it.

Lemma 2. For every $C(P, \sigma)$ with a distinguished player i^* and any $\varepsilon > 0$ there exists an ε -perturbation P' of P that gives the worst possible winning probability $\text{wp}(i^*, P', \sigma) = \text{wp}_\varepsilon(i^*, P, \sigma)$ for player i^* and that for each $1 \leq i \neq j \leq N$ satisfies that either $|P'_{ij} - P_{ij}| = \varepsilon$ or $P'_{ij} \in \{0, 1\}$.

Proof. For a fixed draw σ , the winning chance $\text{wp}(i^*, P, \sigma)$ can be expressed using variables P_{ij} , $1 \leq i \neq j \leq N$. Since each combination of outcomes of the matches determines if i^* won or not, this expression is a polynomial and is linear in each of the P_{ij} 's. If we view $\mathcal{P}(P, \varepsilon)$ as a subset of $\mathbb{R}^{\binom{N-1}{2}}$ with ℓ_∞ metric, then the minimum of the function $\text{wp}(i^*, P', \sigma)$ over the set $\mathcal{P}(P, \varepsilon)$ is attained on the boundary. Hence the result follows. \square

4.1 Deterministic setting

First we focus on the deterministic case. If a player doesn't win a tournament we say he *lost*. Consider a complete deterministic tournament $C(P, \sigma)$ with winner i^* and let $\varepsilon \in (0, 1)$. Lemma 2 gives Corollary 1.

Corollary 1. The value $\text{wp}_\varepsilon(i^*, N, P, \sigma)$ is attained for a matrix P' with some $(0, 1)$ matches altered into $(\varepsilon, 1 - \varepsilon)$ matches and the remaining matches left intact.

Crucial matches. Denote by S_ε the set of $(0, 1)$ matches such that altering them to $(\varepsilon, 1 - \varepsilon)$ matches gives $P' \in \mathcal{P}(P, \varepsilon)$ satisfying $\text{wp}(i^*, P', \sigma) = \text{wp}_\varepsilon(i^*, P, \sigma)$. A $(0, 1)$ match is called *crucial* if replacing it (and it only) by a $(1, 0)$ match makes i^* lose. Denote the set of crucial matches by C . The next lemma says that the crucial matches are the only matches that matter for the sake of determining $\widehat{d}_\varepsilon(i^*, P, \sigma)$.

Lemma 3. Suppose $C(P, \sigma)$ is won by player i^* . Replace some subset S of $(0, 1)$ matches by $(\varepsilon, 1 - \varepsilon)$ matches to get $P' \in \mathcal{P}(P, \varepsilon)$. Suppose c of those $|S| = s$ matches are crucial. Then $\text{wp}(i^*, P', \sigma) = 1 - c \cdot \varepsilon + \varepsilon^2 Q(\varepsilon)$, for some polynomial $Q(\varepsilon)$.

Proof. We account for three scenarios: (A) All matches play out as before; (B) a single crucial match becomes an upset; (C) a single non-crucial match becomes an upset. For appropriate polynomials $Q_i(\varepsilon)$, $i = 1, 2, 3, 4$, the chances of this happening are

$$\begin{aligned} \text{(A)} \quad (1 - \varepsilon)^s &= 1 - s \cdot \varepsilon + \varepsilon^2 Q_1(\varepsilon) \\ \text{(B)} \quad c\varepsilon(1 - \varepsilon)^{s-1} &= c\varepsilon + \varepsilon^2 Q_2(\varepsilon) \\ \text{(C)} \quad (s - c)\varepsilon(1 - \varepsilon)^{s-1} &= (s - c)\varepsilon + \varepsilon^2 Q_3(\varepsilon) \end{aligned}$$

Since i^* wins in (A) and (C) and loses in (B), we have

$$1 - c\varepsilon - \varepsilon^2 Q_4(\varepsilon) \leq \text{wp}_\varepsilon(i^*, P', \sigma) \leq 1 - c\varepsilon - \varepsilon^2 Q_2(\varepsilon)$$

where $Q_4(\varepsilon) = Q_1(\varepsilon) + Q_3(\varepsilon)$. The result follows. \square

Lemma 3 yields the following corollary.

Corollary 2. Let $C(P, \sigma)$ be a deterministic tournament with winner i^* and c crucial matches and let $\varepsilon \in (0, 1)$. Then $\widehat{d}_\varepsilon(i^*, P, \sigma) = c\varepsilon$.

We now show that crucial matches can be found efficiently.

Lemma 4. Given a complete deterministic tournament $C(P, \sigma)$, there exists an $O(N \log N)$ algorithm that computes the number of crucial matches.

Proof. Going from below, at each node of the tournament (including the root) we store the winner of the corresponding subtournament. Since $N - 1$ matches are played in total, this takes $O(N)$.

Now imagine a single match corresponding to a subtree S was switched. In order to determine if the original winner still wins we might need to recompute the winners of all the subtournaments containing S . Since the depth of the tree is $\log(N)$, there are at most $\log(N)$ such tournaments and re-computing the winner of each takes constant time. Overall we get the running time $O(N + N \cdot \log N) = O(N \log N)$. \square

This proves the deterministic case of Theorem 2.

4.2 Probabilistic setting

The idea is similar to the deterministic case. By Lemma 2, we want to adjust each match in one of the two directions as much as possible. Since $\text{wp}(1, P, \sigma)$ is a polynomial linear in each P_{ij} , for any fixed P_{ij} we can express it as

$$\text{wp}(1, P, \sigma) = \alpha_{ij} \cdot P_{ij} + \beta_{ij}.$$

We show that in order to maximize $\widehat{d}_\varepsilon(1, P, \sigma)$,

- if $\alpha_{ij} > 0$, we adjust P_{ij} to $\max\{0, P_{ij} - \varepsilon\}$,
- if $\alpha_{ij} < 0$, we adjust it to $\min\{1, P_{ij} + \varepsilon\}$,
- if $\alpha_{ij} = 0$, the match isn't crucial (i.e. the ε -perturbation of P_{ij} doesn't influence $\widehat{d}_\varepsilon(1, P, \sigma)$),

to get $\widehat{d}_\varepsilon(1, P, \sigma) = \sum_{i,j} |\alpha_{ij}|$. Since all α_{ij} can be found in polynomial time by extending the recursive procedure for finding the winning probabilities, we establish Lemma 5 that covers the probabilistic case of Theorem 2.

Lemma 5. Let $C(P, \sigma)$ be a complete probabilistic tournament. Then $\widehat{d}_\varepsilon(1, P, \sigma)$ can be computed in polynomial time.

4.3 Consequences

Our main result has a number of important consequences which we present below.

Corollary 3. The RTFP and the RPTFP problems are NP-complete.

Proof. The hardness result is an easy consequence of the NP-hardness of the TFP problem [Aziz *et al.*, 2014], which is obtained as follows. Consider the RTFP problem (which is a special case of RPTFP problem) with $c = N + 1$ with sufficiently small $\varepsilon > 0$. Then for every draw σ we have $\widehat{d}_\varepsilon(i^*, P, \sigma) \leq c\varepsilon$; and hence the answer to the RTFP problem is yes iff the answer to the TFP problem is yes. Hence both RTFP and RPTFP problems are NP-hard. Since every draw is a polynomial witness, both conditions (a) and (b) for RPFTP can be checked in polynomial time by the results of [Vu *et al.*, 2009] and Theorem 2, respectively. \square

In [Aziz *et al.*, 2014] two important special cases are considered for TFP, namely, tournaments where there are constant number of types of players, and tournaments with linear ordering on players with constant number of exceptions. For both the cases, polynomial-time algorithms are presented in [Aziz *et al.*, 2014] to enumerate all winning draws. In combination with Theorem 2 it follows that the most robust winning draw can also be approximated in polynomial time.

Corollary 4. For the two special cases of TFP from [Aziz *et al.*, 2014], the most robust winning draw (if a winning draw exists) can be approximated in polynomial time.

5 Conclusion

We studied the problem of robustness related to fixing draws in BKT. We presented several illuminating examples related to the robustness properties of optimal and near optimal draws. We established polynomial-time algorithm to approximate the robustness of draws for sufficiently small $\varepsilon > 0$. As a consequence, for the robustness of draws, we establish NP-completeness in general and polynomial-time algorithms for special cases. Interesting directions of future work are (a) computation of robustness when $\varepsilon > 0$ is not small; (b) higher order uncertainties, such as uncertainty in ε or different ε 's for different matches; and (c) the robustness question in other forms of tournaments.

Acknowledgements. This research was partly supported by Austrian Science Fund (FWF) NFN Grant No S11407-N23 (RiSE/SHiNE), Vienna Science and Technology Fund (WWTF) through project ICT15-003, ERC Start grant (279307: Graph Games), and ERC Advanced Grant (267989: QUAREM).

References

- [Aziz *et al.*, 2012] Haris Aziz, Paul Harrenstein, Markus Brill, Jérôme Lang, Felix A. Fischer, and Hans Georg Seedig. Possible and necessary winners of partial tournaments. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012*, pages 585–592, 2012.
- [Aziz *et al.*, 2014] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Paul Stursberg, and Toby Walsh. Fixing a balanced knockout tournament. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014*, pages 552–558, 2014.
- [Chatterjee *et al.*, 2016] Krishnendu Chatterjee, Rasmus Ibsen-Jensen, and Josef Tkadlec. Robust draws in balanced knockout tournaments. *Arxiv CoRR 1604.05090* <http://arxiv.org/abs/1604.05090>, 2016.
- [Filar and Vrieze, 1997] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.
- [Gusfield and Martel, 2002] Dan Gusfield and Charles U. Martel. The structure and complexity of sports elimination numbers. *Algorithmica*, 32(1):73–86, 2002.
- [Kern and Paulusma, 2004] Walter Kern and Daniël Paulusma. The computational complexity of the elimination problem in generalized sports competitions. *Discrete Optimization*, 1(2):205–214, 2004.
- [Lang *et al.*, 2007] Jérôme Lang, Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Winner determination in sequential majority voting. In *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pages 1372–1377, 2007.
- [Lang *et al.*, 2012] Jérôme Lang, Maria Silvia Pini, Francesca Rossi, Domenico Salvagnin, Kristen Brent Venable, and Toby Walsh. Winner determination in voting trees with incomplete preferences and weighted votes. *Autonomous Agents and Multi-Agent Systems*, 25(1):130–157, 2012.
- [Puterman, 1994] M. L. Puterman. *Markov Decision Processes*. J. Wiley and Sons, 1994.
- [Rockafellar, 2007] R Tyrrell Rockafellar. Coherent approaches to risk in optimization under uncertainty. *Tutorials in operations research, INFORMS*, 2007.
- [Stanton and Williams, 2011a] Isabelle Stanton and Virginia Vassilevska Williams. Manipulating stochastically generated single-elimination tournaments for nearly all players. In *Internet and Network Economics - 7th International Workshop, WINE 2011*, pages 326–337, 2011.

- [Stanton and Williams, 2011b] 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*, pages 357–364, 2011.
- [Vu and Shoham, 2011] Thuc Vu and Yoav Shoham. Fair seeding in knockout tournaments. *ACM TIST*, 3(1):9, 2011.
- [Vu *et al.*, 2009] Thuc Vu, Alon Altman, and Yoav Shoham. On the complexity of schedule control problems for knockout tournaments. In *8th International Joint Conference on Autonomous Agents and Multiagent Systems AAMAS 2009*, pages 225–232, 2009.
- [Williams, 2010] Virginia Vassilevska Williams. Fixing a tournament. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010*, 2010.
- [Xia and Conitzer, 2011] Lirong Xia and Vincent Conitzer. Determining possible and necessary winners given partial orders. *J. Artif. Intell. Res. (JAIR)*, 41:25–67, 2011.