Generalized Kings and Single-Elimination Winners in Random Tournaments

Pasin Manurangsi[1] and **Warut Suksompong**[2]
[1]Google Research, USA
[2]School of Computing, National University of Singapore, Singapore
pasin@google.com, warut@comp.nus.edu.sg

Abstract

Tournaments can be used to model a variety of practical scenarios including sports competitions and elections. A natural notion of strength of alternatives in a tournament is a generalized king: an alternative is said to be a k-*king* if it can reach every other alternative in the tournament via a directed path of length at most k. In this paper, we provide an almost complete characterization of the probability threshold such that all, a large number, or a small number of alternatives are k-kings with high probability in two random models. We show that, perhaps surprisingly, all changes in the threshold occur in the range of constant k, with the biggest change being between $k = 2$ and $k = 3$. In addition, we establish an asymptotically tight bound on the probability threshold for which all alternatives are likely able to win a single-elimination tournament under some bracket.

1 Introduction

Social choice theory is the study of how to aggregate individual preferences and opinions of agents on a set of alternatives in order to reach a collective decision. In many practical situations, the relationship between the alternatives is represented by a *dominance relation*, which specifies the relative strength of the alternatives in any pairwise comparison. For example, in sports competitions the dominance relation signifies the match outcome when two players or teams play each other, while in elections the relation represents the pairwise majority comparisons among the candidates. The structure consisting of the alternatives and their dominance relation is called a *tournament*, and the analysis of tournament winner selection methods—also known as *tournament solutions*—has received significant attention from researchers in the past few decades [Laslier, 1997; Brandt *et al.*, 2016; Suksompong, 2021].

Among the vast array of tournament solutions proposed in the literature, two of the earliest and best-known ones are the *top cycle* [Good, 1971; Schwartz, 1972; Miller, 1977] and the *uncovered set* [Fishburn, 1977; Miller, 1980]. An alternative belongs to the top cycle if it can reach every other alternative via a directed path in the tournament. Note that if the tournament contains n alternatives, any such path has length $n-1$ or less (the length of a path refers to the number of edges in the path).[1] Similarly, the uncovered set—also known as the set of *kings* [Maurer, 1980]—consists of the alternatives that can reach every other alternative via a path of length at most two. It is clear from the definitions that the uncovered set is always a subset of the top cycle. Moreover, both tournament solutions can be viewed as special cases of a generalized notion of kings called k-*kings*, which correspond to the alternatives that can reach every other alternative via a path of length at most k. Indeed, the uncovered set is the set of 2-kings, while the top cycle contains precisely the $(n-1)$-kings.

Given that tournament solutions are meant to distinguish the best alternatives from the rest, it is natural to ask how selective each tournament solution is. Moon and Moser [1962] and Fey [2008] addressed this question and showed that the top cycle and the uncovered set are likely to include all alternatives when the tournament is large. In particular, their results hold under the *uniform random model*, wherein each edge is oriented in one direction or the other with equal probability independently of other edges. Saile and Suksompong [2020] extended these results to the *generalized random model*, in which the orientation of each edge is determined by probabilities within the range $[p, 1-p]$ for some parameter $p \leq 1/2$, and these probabilities may vary across edges. The generalized random model allowed these authors to demonstrate a difference between the two tournament solutions—while the top cycle almost never excludes any alternative as long as $p \in \omega(1/n)$, the uncovered set is likely to select all alternatives only when $p \in \Omega(\sqrt{\log n/n})$, so the two thresholds differ by roughly $\Theta(\sqrt{n})$. This raises the following question: How does the probability threshold change as we transition from $k = 2$ to $k = n-1$? Does it already decrease at around $k = \sqrt{n}$, or does it remain the same until, say, $k \approx n/2$?

In this paper, we show that, perhaps surprisingly, all of the changes in the probability threshold occur when k is constant. In fact, when $k = 6$, all alternatives are already likely

[1]The bound $n-1$ cannot be improved. To see this, consider a tournament with alternatives $x_1, \dots, x_n$ such that x_i dominates x_j if $i - j \geq 2$ or $j - i = 1$. Alternative x_1 can reach every other alternative, but it cannot reach x_n via a path of length $n-2$ or less.

Tournament solution		Condorcet random model	Generalized random model
k-kings	$k = 2$	$\Omega(\sqrt{\log n/n})$ [Saile and Suksompong, 2020]	$\Omega(\sqrt{\log n/n})$ [Saile and Suksompong, 2020]
	$3 \leq k \leq 4$	$\Omega(\log n/n)$ (Thm. 3.1, 3.3)	$\Omega(\log n/n)$ (Thm. 3.1, 3.3)
	$k = 5$	$\omega(1/n)$ (Thm. 3.4)	$\Omega(\log\log n/n)$ (Thm. 3.8)
	$6 \leq k \leq n-1$	$\omega(1/n)$ (Thm. 3.5)	$\omega(1/n)$ (Thm. 3.5)
Single-elimination winners		$\Omega(\log n/n)$ [Kim *et al.*, 2017]	$\Omega(\log n/n)$ (Thm. 4.3)

Table 1: Summary of the bounds on the probability p at which the respective tournament solutions select all alternatives with high probability under the corresponding random model. All bounds are asymptotically tight except the bound for $k = 5$ with the generalized random model, where there is a gap between $\Omega(\log\log n/n)$ and $\omega(1/n)$. The results with $\Omega(\cdot)$ hold when the associated constant term is sufficiently large.

to be k-kings provided that $p \in \omega(1/n)$, the same threshold as $k = n - 1$—this significantly strengthens the result of Saile and Suksompong [2020] on the top cycle. For $k = 3$ and 4, we establish an asymptotically tight bound of $p \in \Omega(\log n/n)$, while for $k = 5$ we leave the only (small) gap between $\Omega(\log\log n/n)$ and $\omega(1/n)$. Besides the generalized random model, we consider a more specific model which has nevertheless been studied in several papers called the *Condorcet random model*. In this model, there is a parameter p and a linear order of alternatives from strongest to weakest, and the probability that a stronger alternative dominates a weaker one is $1 - p$, independently of other pairs of alternatives.[2] For the Condorcet random model, we show that the threshold for $k = 5$ is $\omega(1/n)$, whereas the thresholds for other values of k remain tight. Our results are summarized in Table 1 and presented in **Section 3**. Taken together, they reveal the intriguing facts that (i) the uncovered set is distinctly more selective than k-kings for $k \geq 3$; (ii) 3-kings and 4-kings are slightly more selective than higher-order kings; and (iii) there is virtually no difference in discriminative power from $k = 5$ all the way to $k = n - 1$.

In addition to k-kings, we also consider the set of *single-elimination winners*, which are alternatives that can win a (balanced) single-elimination tournament under some bracket, where the match outcomes in the single-elimination tournament are determined according to the dominance relation in the original tournament.[3] Kim et al. [2017] showed that all alternatives are likely to be single-elimination winners in the Condorcet random model as long as $p \in \Omega(\log n/n)$, and this bound is tight.[4] For the generalized random model, they established an analogous statement in the range $p \in \Omega(\sqrt{\log n/n})$. We close this gap by proving that even for the generalized random model, $p \in \Omega(\log n/n)$ already suffices for all alternatives to be single-elimination winners with high probability; moreover, a winning bracket for each alternative can be computed in polynomial time. Our result, which can be found in **Section 4**, further lends credence to the observation that real-world tournaments can be easily manipulated [Mattei and Walsh, 2016].

Finally, in **Section 5**, we move beyond the question of when *all* alternatives are likely to be selected by a tournament solution, and instead ask when this is the case for a large or small number of alternatives. For the uncovered set, even though $p \in \Omega(\sqrt{\log n/n})$ is required in order for all alternatives to be chosen with high probability [Saile and Suksompong, 2020], we show that most of them are already likely to be included as long as $p \in \Omega(\log n/n)$. This threshold is exactly where the transition occurs: if $p \in O(\log n/n)$, we prove that the uncovered set almost surely contains only a small fraction of the alternatives. Furthermore, for any $k \geq 3$, we establish that a large fraction of alternatives are likely to be k-kings provided that $p \in \omega(1/n)$. These results illustrate the probability range under which each tournament solution is discriminative, and again exhibit a clear difference between the uncovered set and higher-order kings.

1.1 Related Work

Tournament solutions have been extensively studied for the past several decades; we refer to the surveys by Laslier [1997] and Brandt et al. [2016]. There are several containment relations among common tournament solutions. For example, the Copeland set, Slater set, Markov set, and Banks set are all contained in the uncovered set, which is in turn contained in the top cycle—this provides a range of options in terms of discriminative power and other properties. Even though k-kings admit a simple and elegant definition generalizing the top cycle as well as the uncovered set, and have attracted interest from graph theorists [Petrovic and Thomassen, 1991; Tan, 2006; Brcanov and Petrovic, 2010], as far as we know, they have not been studied in the social choice context until recently. Kim and Vassilevska Williams [2015] and Kim et al. [2017] identified conditions under which a 3-king can win a single-elimination tournament. Brill et al. [2020] showed that computing the "margin of victory" of k-kings can be done efficiently for $k \leq 3$ but becomes NP-hard for $k \geq 4$. Brill et al. [2021] illustrated through experiments that the margin of victory of 3-kings behaves much more similarly to

[2] The uniform random model corresponds to taking $p = 1/2$. For any p, the Condorcet random model with parameter p is a special case of the generalized random model with the same p. Hence, a positive result for the generalized random model carries over to the Condorcet random model, while a negative result transfers in the opposite direction.

[3] Following prior work on single-elimination tournaments, we assume that the number of alternatives is a power of two.

[4] If $p \in o(\log n/n)$, the weakest alternative dominates $o(\log n)$ alternatives in expectation; this is insufficient since winning $\log_2 n$ matches is required to win a single-elimination tournament.

that of the top cycle than to the corresponding notion of the uncovered set; our results therefore complement theirs by exhibiting that analogous behavior can be observed with respect to discriminative power.

As we mentioned earlier, the study of tournament solutions under the probabilistic lens was initiated by Moon and Moser [1962], who showed that the top cycle is likely to choose all alternatives when the tournament is generated by the uniform random model. Fey [2008] and Scott and Fey [2012] established the same property for the uncovered set as well as two other tournament solutions, the *Banks set* and the *minimal covering set*, while Fisher and Ryan [1995] showed that the *bipartisan set* includes half of the alternatives on average.[5] The Condorcet random model has been analyzed, among others, by Frank [1968], Łuczak et al. [1996], Vassilevska Williams [2010], and Kim et al. [2017], with the last paper also proposing the generalized random model. As Saile and Suksompong [2020] pointed out, the Condorcet random model suffers from limitations such as using the same probability for all pairs of alternatives (regardless of the extent to which one alternative is stronger than the other) or not allowing for "bogey teams" (i.e., weak teams that often beat certain stronger teams). The generalized random model only assumes that each match is sufficiently random, and therefore does not have these limitations.

Finally, single-elimination tournaments have constituted a popular topic of study in the past decade; see the surveys by Vassilevska Williams [2016] and Suksompong [2021]. In particular, even though the problem of determining whether an alternative can win a single-elimination tournament is known to be NP-hard [Aziz *et al.*, 2018], a wide range of algorithmic and complexity results have been developed by this active line of work.

2 Preliminaries

A tournament T consists of a set $V = \{x_1, \ldots, x_n\}$ of vertices, also called *alternatives*, and a set E of directed edges. For any two alternatives $x_i, x_j \in V$, there exists either an edge from x_i to x_j or an edge from x_j to x_i, but not both. The edges represent a *dominance relation* between the alternatives: an edge from x_i to x_j means that x_i *dominates* x_j, a relation which we denote by $x_i \succ x_j$. The *outdegree* (resp., *indegree*) of an alternative x_i is the number of alternatives that x_i dominates (resp., that dominate x_i). We extend the dominance relation to sets of alternatives: for $V_1, V_2 \subseteq V$, we write $V_1 \succ V_2$ to mean that $x \succ x'$ for all $x \in V_1$ and $x' \in V_2$, and $V_1 \succ x'$ to mean that $x \succ x'$ for all $x \in V_1$. A set $V' \subseteq V$ is called a *dominating set* if for every $x \in V \setminus V'$, there exists $x' \in V'$ such that $x' \succ x$.

We can now define the key notions of this paper.

- For any integer $k \geq 2$, an alternative is said to be a *k-king* if it can reach every other alternative via a directed path of length at most k.
- Suppose that $n = 2^r$ for some nonnegative integer r. An alternative is said to be a *single-elimination winner* if it wins a (balanced) single-elimination tournament under some bracket, where the outcome of each match is determined according to the dominance relation.[6]

We will consider two random models for generating tournaments. In the *Condorcet random model*, there is a parameter $0 \leq p \leq 1/2$. For $i < j$, alternative x_j dominates x_i with probability p (so x_i dominates x_j with probability $1 - p$), independently of other pairs of alternatives. In the *generalized random model*, there is a parameter $p_{i,j}$ for each pair $i \neq j$, where $p_{i,j} + p_{j,i} = 1$. For any pair i, j, alternative x_i dominates x_j with probability $p_{i,j}$, independently of other pairs. We will generally allow each probability $p_{i,j}$ to be chosen from the range $[p, 1-p]$ for a given parameter $0 \leq p \leq 1/2$. Before a tournament is generated from the generalized random model, the *expected outdegree* of x_i is defined as $\sum_{j \neq i} p_{i,j}$. Following standard terminology in probability theory, we say that an event whose probability depends on n occurs "with high probability" if the probability that it occurs approaches 1 as $n \to \infty$.

We now list two well-known probabilistic statements that will be used multiple times in this paper. The first statement provides an upper bound on the probability that a sum of independent random variables is far from its expectation.

Lemma 2.1 (Chernoff bound)**.** *Let $X_1, \ldots, X_k$ be independent random variables taking values in $[0, 1]$, and let $X := X_1 + \cdots + X_k$. Then, for any $\delta \in [0, 1]$,*

$$\Pr[X \geq (1+\delta)\mathbb{E}[X]] \leq \exp\left(\frac{-\delta^2\mathbb{E}[X]}{3}\right)$$

and

$$\Pr[X \leq (1-\delta)\mathbb{E}[X]] \leq \exp\left(\frac{-\delta^2\mathbb{E}[X]}{2}\right).$$

The second statement is a simple upper bound on the expression $1 + x$.

Lemma 2.2. *For every real number x, we have $1 + x \leq e^x$.*

We end this section with a lemma on the degree of alternatives in a tournament.

Lemma 2.3. *Let $1 \leq r \leq n$. For any tournament T, the average outdegree and the average indegree of the alternatives in any subset of size r are at least $(r-1)/2$. Similarly, the average expected outdegree of the alternatives in such a subset is at least $(r-1)/2$.*

Proof. We prove the statement for the average outdegree; the proofs for the average indegree as well as the average expected outdegree are similar. In a subset of alternatives of size r, there are a total of $r(r-1)/2$ edges. Hence, the sum of the outdegrees of the r alternatives is at least $r(r-1)/2$, implying that their average is at least $(r-1)/2$. □

Unless a base is explicitly specified, log refers to the natural logarithm. All omitted proofs can be found in the full version of our paper [Manurangsi and Suksompong, 2021].

[5] Brandt et al. [2018] showed that any tournament solution that satisfies an attractive property called *stability*, including the top cycle, the minimal covering set, and the bipartisan set, must choose at least half of the alternatives on average.

[6] See, e.g., [Aziz *et al.*, 2018] for formal definitions.

3 Generalized Kings

Recall the result of Saile and Suksompong [2020] that in the generalized random model, all alternatives are 2-kings (i.e., belong to the uncovered set) with high probability only if $p \in \Omega(\sqrt{\log n/n})$. For our first result, we show that 3-kings are not as selective: even when $p \in \Omega(\log n/n)$, it is already likely that none of the alternatives is excluded by this set.

Theorem 3.1. *Assume that a tournament T is generated according to the generalized random model, and that $p_{i,j} \in [30\log n/n, 1-30\log n/n]$ for all $i \neq j$. Then with high probability, all alternatives in T are 3-kings.*

To prove this theorem, we establish a rather general lemma on the probability of one set dominating another, which will also be useful in our analysis of 5-kings later.

Lemma 3.2. *Let T_0 be a tournament consisting of n_0 alternatives $V_0 := \{x_1,\dots,x_{n_0}\}$, and let $q_{1,1}, q_{1,2},\dots,q_{n_0,1},q_{n_0,2} \in \left[\frac{10\lambda}{n_0}, 1\right]$ for some $1 \le \lambda \le \frac{n_0}{10}$. Suppose that we randomly create a set S_1 by including each alternative x_i independently with probability $q_{i,1}$, and a set S_2 by including each alternative x_i independently with probability $q_{i,2}$. Then, $\Pr[S_1 \cap S_2 = \emptyset$ and $S_1 \succ S_2] \le e^{-\lambda}$.*

Lemma 3.2 allows for a short proof of Theorem 3.1.

Proof of Theorem 3.1. Fix a pair of distinct alternatives x_i, x_j. We first bound the probability that x_j cannot reach x_i via a directed path of length at most three.

Consider the tournament T^0 defined by restricting T to $V^0 := V \setminus \{x_i, x_j\}$. Let S_1 denote the set of alternatives in V^0 that dominate x_i with respect to T, and let S_2 denote the set of alternatives in V^0 that are dominated by x_j with respect to T. Notice that if $S_1 \cap S_2 \neq \emptyset$ or $S_1 \not\succ S_2$, then there is a path of length at most three from x_j to x_i. Furthermore, from the assumption of the theorem, each alternative belongs to each of S_1 and S_2 independently with probability at least $\frac{30\log n}{n}$, which is at least $\frac{25\log n}{|V^0|}$ for any sufficiently large n. As a result, we may apply Lemma 3.2 with $\lambda = 2.5\log n$, which gives

$$\begin{aligned}&\Pr[\text{there is no path of length at most three from } x_j \text{ to } x_i]\\ &\qquad \le \Pr[S_1 \cap S_2 = \emptyset \text{ and } S_1 \succ S_2]\\ &\qquad \le e^{-\lambda} = 1/n^{2.5}.\end{aligned}$$

Finally, applying the union bound over all (ordered) pairs of alternatives $x_i \neq x_j$, the probability that some alternative cannot reach some other alternative via a directed path of length at most three is no more than $1/n^{0.5}$, which converges to 0 as n goes to infinity. □

Next, we show that in the Condorcet random model, if $p \in \Theta(\log n/n)$ and the associated constant is low enough, there is likely to be an alternative that is not a 4-king. Combined with Theorem 3.1, this implies that the bound $\Theta(\log n/n)$ is asymptotically tight for both 3- and 4-kings in both random models that we consider.

Theorem 3.3. *Assume that a tournament T is generated according to the Condorcet random model, and that $p \le 0.1\log n/n$. Then with high probability, there exists an alternative in T that is not a 4-king.*

Our results so far demonstrate that k-kings for any $k \ge 3$ are much closer to the top cycle (i.e., $(n-1)$-kings) than to the uncovered set (i.e., 2-kings) in terms of discriminative power. In the remainder of this section, we show that there is virtually no difference in selectiveness between k-kings for $k \ge 6$ and the top cycle. We begin by showing that in the Condorcet random model, all alternatives are likely to be 5-kings provided that $p \in \omega(1/n)$—this gives a complete characterization of the probability threshold for the Condorcet random model.

Theorem 3.4. *Assume that a tournament T is generated according to the Condorcet random model, and that $p \in \omega(1/n)$. Then with high probability, all alternatives in T are 5-kings.*

For the generalized random model, we show that as long as $p \in \omega(1/n)$, all alternatives are likely to be 6-kings.

Theorem 3.5. *Assume that a tournament T is generated according to the generalized random model, and that $p_{i,j} \in \omega(1/n)$ for all $i \neq j$. Then with high probability, all alternatives in T are 6-kings.*

In light of Theorem 3.5, the only remaining gap in our probability threshold characterization is for 5-kings in the generalized random model. We conjecture that the true threshold is $\omega(1/n)$, but our proof of Theorem 3.4 relies on the ordering of the alternatives in the Condorcet random model and cannot be easily extended. Instead, we present a slightly weaker bound of $\Omega(\log\log n/n)$—this shows that 5-kings are closer to 6-kings than to 4-kings with respect to selective power. To establish this bound, we need a lemma on generalized dominating sets.

Definition 3.6. Given a positive integer r, a set of alternatives D is said to be an *r-dominating set* of a tournament T if every alternative $x \notin D$ is dominated by at least r alternatives in D.

Lemma 3.7. *For any tournament T and any positive integer r, there exists an r-dominating set of T of size at most $r\lceil\log_2 n\rceil$.*

Proof. Let us start with $D = \emptyset$ and repeat the following procedure r times: find a minimum dominating set S of the restriction of T on $V \setminus D$, and update D to $D \cup S$. It is clear that the final set D is an r-dominating set of T. Furthermore, it is well-known [Megiddo and Vishkin, 1988, Fact 2.5] that any tournament has a dominating set of size at most $\lceil\log_2 n\rceil$, which implies that the final set D is of size at most $r\lceil\log_2 n\rceil$. □

We are now ready to establish our result on 5-kings.

Theorem 3.8. *Assume that a tournament T is generated according to the generalized random model, and that $p_{i,j} \in [50(\log\log n)/n, 1-50(\log\log n)/n]$ for all $i \neq j$. Then with high probability, all alternatives in T are 5-kings.*

Proof. Define a tournament T' on alternatives $x_1,\dots,x_n$ so that for each pair $i \neq j$, we have $x_i \succ x_j$ if $p_{i,j} > 1/2$. Then, our tournament T is generated by reversing the edges of T' so that the edge between x_i and x_j is reversed with probability $q_{i,j} := \min\{p_{i,j}, 1-p_{i,j}\} \le 1/2$, independently of other edges. For each alternative x, let $I(x)$ denote the set of alternatives that dominate x in T'.

Let $r = \lceil 1.1 \log_2 n \rceil$, and let D and D_{inv} be a minimum r-dominating set of the tournament T' and its "inverse" constructed by reversing all edges in T', respectively. From Lemma 3.7, we have $|D|, |D_{\text{inv}}| \leq r\lceil \log_2 n \rceil$, which is at most $2.3(\log n)^2$ for any sufficiently large n.

We define the following three events in T:

- E_1: For every $x_m \notin D$, there exists $x_i \in D$ such that $x_i \succ x_m$.
- E_2: For every $x_\ell \notin D_{\text{inv}}$, there exists $x_j \in D_{\text{inv}}$ such that $x_\ell \succ x_j$.
- E_3: For every $x_i \in D$ and $x_j \in D_{\text{inv}}$, there exists a directed path of length at most three from x_j to x_i.

Similarly to the proof of Theorem 3.4, when E_1, E_2, and E_3 all occur, every alternative is a 5-king. As a result, it suffices to show that each of the three events occurs with high probability.

For E_1, we have

$$\begin{aligned}
\Pr[\neg E_1] &= \Pr[\exists x_m \notin D, \forall x_i \in D, x_m \succ x_i] \\
&\leq \sum_{x_m \notin D} \Pr[\forall x_i \in D, x_m \succ x_i] \\
&\leq \sum_{x_m \notin D} \Pr[\forall x_i \in D \cap I(x_m), x_m \succ x_i] \\
&= \sum_{x_m \notin D} \prod_{x_i \in D \cap I(x_m)} \Pr[x_m \succ x_i] \\
&= \sum_{x_m \notin D} \prod_{x_i \in D \cap I(x_m)} q_{i,m} \\
&\leq \sum_{x_m \notin D} (1/2)^{|D \cap I(x_m)|} \\
&\leq \sum_{x_m \notin D} (1/2)^r \\
&\leq \sum_{x_m \notin D} 1/n^{1.1} \\
&\in o(1),
\end{aligned}$$

where the first inequality follows the union bound and the fourth inequality from the fact that D is an r-dominating set in T'. An analogous argument shows that E_2 also occurs with high probability.

Finally, consider E_3. Since both D and D_{inv} are of size $O((\log n)^2)$, by the union bound, it suffices to show that for each fixed pair $x_i \in D$ and $x_j \in D_{\text{inv}}$, a path of length at most three from x_j to x_i exists with probability $1 - o(1/(\log n)^4)$.

To prove this, consider the tournament T^0 defined by restricting T to $V^0 := V \setminus (D \cup D_{\text{inv}})$. Let S_1 denote the set of alternatives in V^0 that dominate x_i with respect to T, and let S_2 denote the set of alternatives in V^0 that are dominated by x_j with respect to T. Notice that if $S_1 \cap S_2 \neq \emptyset$ or $S_1 \not\succ S_2$, then there is a path of length at most three from x_j to x_i. Furthermore, from the assumption of the theorem, each alternative belongs to each of S_1 and S_2 independently with probability at least $\frac{50 \log\log n}{n}$, which is at least $\frac{45 \log\log n}{|V^0|}$ for any sufficiently large n. As a result, we may apply Lemma 3.2 with $\lambda = 4.5 \log\log n$, which gives

$$\begin{aligned}
&\Pr[\text{there is no path of length at most three from } x_j \text{ to } x_i] \\
&\qquad \leq \Pr[S_1 \cap S_2 = \emptyset \text{ and } S_1 \succ S_2] \\
&\qquad \leq e^{-\lambda} \\
&\qquad = 1/(\log n)^{4.5} \\
&\qquad \in o(1/(\log n)^4),
\end{aligned}$$

which concludes our proof. □

4 Single-Elimination Winners

In this section, we consider single-elimination winners and derive a tight bound of $\Omega(\log n/n)$ for the generalized random model, thereby strengthening the bound $\Omega(\sqrt{\log n/n})$ of Kim et al. [2017] and matching their bound for the Condorcet random model. As in previous work on this subject, we assume for simplicity that $n = 2^r$ for some positive integer r, so $r = \log_2 n$. In order to construct a winning bracket, a useful notion is that of a "superking", introduced by Vassilevska Williams [2010].

Definition 4.1. Given a tournament T, an alternative x is said to be a *superking* if for every alternative x' such that $x' \succ x$, there exist at least $\log_2 n$ alternatives x'' such that $x \succ x''$ and $x'' \succ x'$.

Lemma 4.2 (Vassilevska Williams [2010]). *In any tournament, every superking is a single-elimination winner, and its winning bracket can be computed in polynomial time.*

Before we proceed to our result, let us briefly recap the proofs of the two aforementioned results by Kim et al. [2017], and explain why they cannot be used to establish our desired strengthening. In order to show that all alternatives are single-elimination winners with high probability when $p \in \Omega(\sqrt{\log n/n})$, these authors showed that all alternatives are likely to be superkings in this range; it is not difficult to verify that this condition no longer holds when $p \in o(\sqrt{\log n/n})$. For the $\Omega(\log n/n)$ bound in the Condorcet random model, they argued that the weakest alternative x is likely to dominate one of the top $n/6$ alternatives, and constructed a winning bracket for this latter alternative y among the top half of the alternatives, so that x can play y in the final round and win the single-elimination tournament. Since there is no clear notion of strength in the generalized random model (indeed, all alternatives may be roughly equally strong, with no linear order), this approach also does not work for our purpose.

At a high level, our proof proceeds by choosing $r = \log_2 n$ alternatives that our desired winning alternative x dominates. In order to ensure that x can play against these alternatives in the r rounds, we partition the r alternatives along with the remaining alternatives into subsets of size $1, 2, \ldots, 2^{r-1}$, so that the r alternatives are superkings in the respective subtournament and can therefore win a single-elimination tournament with respect to their subset by Lemma 4.2.

Theorem 4.3. *Assume that a tournament T is generated according to the generalized random model, and that $p_{i,j} \in [80 \log n/n, 1 - 80 \log n/n]$ for all $i \neq j$. Then with high* *probability, all alternatives in T are single-elimination winners, and a winning bracket of each alternative can be computed in polynomial time.*

5 Number of Kings

So far, we have addressed the question of when each tournament solution is likely to select *all* alternatives, i.e., the case where the solution is decidedly not useful. In this section, we move beyond this question (which is also the focus of several previous works) and ask when a small or large number of alternatives are chosen. Our first result shows that even though $p \in \Omega(\sqrt{\log n/n})$ is required for the uncovered set to choose all alternatives with high probability [Saile and Suksompong, 2020], the smaller threshold of $p \in \Omega(\log n/n)$ suffices in order for most of the alternatives to be included.

Theorem 5.1. *Assume that a tournament T is generated according to the generalized random model, and that $p_{i,j} \in [50\log n/n, 1 - 50\log n/n]$ for all $i \neq j$. Then with high probability, at least $0.9n$ alternatives in T are 2-kings.*

Proof. Assume without loss of generality that the alternatives before the tournament T is generated are $x_1, \ldots, x_n$ in non-increasing order of expected outdegree. Let $r = \lceil 0.9n \rceil$, and consider any $1 \leq i \leq r$. We will show that the probability that x_i is a 2-king in T is at least $1 - o(1/n)$. The union bound then implies that with high probability, at least $0.9n$ alternatives in T are 2-kings.

To show that $\Pr[x_i \text{ is a 2-king in } T]$ is at least $1 - o(1/n)$, the union bound again implies that it suffices to prove that $\Pr[\text{there is no path of length at most two from } x_i \text{ to } x_j]$ is at most $o(1/n^2)$ for each alternative $x_j \neq x_i$.

We henceforth fix $1 \leq i \leq r$ and $x_j \neq x_i$. By Lemma 2.3 on $x_i, x_{i+1}, \ldots, x_n$, the expected outdegree of x_i is at least $(n-i)/2$, which is at least $0.045n$ for any sufficiently large n. As a result, for large enough n, we have

$$\begin{aligned}
&\Pr[\text{there is no path of length at most two from } x_i \text{ to } x_j] \\
&= \Pr[x_j \succ x_i] \cdot \Pr[\forall k \notin \{i,j\}, x_k \succ x_i \text{ or } x_j \succ x_k] \\
&= (1 - p_{i,j}) \cdot \prod_{k \notin \{i,j\}} \Pr[x_k \succ x_i \text{ or } x_j \succ x_k] \\
&= (1 - p_{i,j}) \cdot \prod_{k \notin \{i,j\}} (1 - p_{i,k}p_{k,j}) \\
&\leq \exp(-p_{i,j}) \cdot \prod_{k \notin \{i,j\}} \exp(-p_{i,k}p_{k,j}) \\
&\leq \exp\left(-\frac{50\log n}{n} \cdot p_{i,j}\right) \cdot \prod_{k \notin \{i,j\}} \exp\left(-\frac{50\log n}{n} p_{i,k}\right) \\
&= \exp\left(-\frac{50\log n}{n}\left(\sum_{k \neq i} p_{i,k}\right)\right) \\
&\leq \exp\left(-\frac{50\log n}{n} \cdot 0.045n\right) \\
&= 1/n^{2.25} \\
&\in o(1/n^2).
\end{aligned}$$

We use Lemma 2.2 for the first inequality, and the assumption that the expected outdegree of x_i is at least $0.045n$ for the last inequality. This concludes our proof. □

Our next result establishes the asymptotic tightness of the threshold in Theorem 5.1.

Theorem 5.2. *Assume that a tournament T is generated according to the Condorcet random model, and that $p \leq 0.1\log n/n$. Then with high probability, at most $\sqrt{n}\log n$ alternatives in T are 2-kings.*

Finally, we show that as long as $p \in \omega(1/n)$, a large number of alternatives are already likely to be 3-kings (and therefore k-kings for every $k \geq 3$). This bound is again tight: when $p \in O(1/n)$ and the tournament is generated from the Condorcet random model, there is at least a constant probability that the strongest alternative dominates all remaining alternatives (in which case it is the only k-king for each $k \geq 2$).

Theorem 5.3. *Assume that a tournament T is generated according to the generalized random model, and that $p_{i,j} \in \omega(1/n)$ for all $i \neq j$. Then with high probability, at least $0.9n$ alternatives in T are 3-kings.*

6 Concluding Remarks

In this paper, we have extensively investigated the behavior of generalized kings and single-elimination winners in random tournaments in view of their discriminative power. Our results reveal surprisingly clear distinctions between the uncovered set and k-kings for $k \geq 3$, and illustrate why manipulating a single-elimination tournament is often possible in practice despite the problem being NP-hard. All of the bounds that we obtained are asymptotically tight except for the bound for 5-kings in the generalized random model (Theorem 3.8); one could try to close this gap.

An exciting future direction in our view is to study k-kings and single-elimination winners with respect to axiomatic and computational properties, as is commonly done for other tournament solutions [Laslier, 1997; Brandt *et al.*, 2016]. For example, the uncovered set is known to be the finest tournament solution satisfying Condorcet consistency, neutrality, and expansion [Moulin, 1986]. Which axioms does the set of 3-kings satisfy, and can we characterize it by a collection of axioms? One could also study the relationship between these tournament solutions and traditional ones—this was partially done by Kim et al. [2017], who showed for instance that any alternative in the Copeland set or the Slater set can always win a single-elimination tournament. Another possible avenue is to extend our results to other stochastic models for tournaments—several interesting models have been studied experimentally by Brandt and Seedig [2016] and Brill et al. [2021]. Such questions illustrate the richness of tournaments and probabilistic models, which we expect to give rise to further fruitful research.

Acknowledgments

This work was partially supported by an NUS Start-up Grant. We would like to thank the anonymous reviewers for their valuable comments.

References

[Aziz *et al.*, 2018] Haris Aziz, Serge Gaspers, Simon Mackenzie, Nicholas Mattei, Paul Stursberg, and Toby Walsh. Fixing balanced knockout and double elimination tournaments. *Artificial Intelligence*, 262:1–14, 2018.

[Brandt and Seedig, 2016] Felix Brandt and Hans Georg Seedig. On the discriminative power of tournament solutions. In *Selected Papers of the International Conference on Operations Research, OR2014*, pages 53–58, 2016.

[Brandt *et al.*, 2016] Felix Brandt, Markus Brill, and Paul Harrenstein. Tournament solutions. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 3, pages 57–84. Cambridge University Press, 2016.

[Brandt *et al.*, 2018] Felix Brandt, Markus Brill, Hans Georg Seedig, and Warut Suksompong. On the structure of stable tournament solutions. *Economic Theory*, 65(2):483–507, 2018.

[Brcanov and Petrovic, 2010] Dejan Brcanov and Vojislav Petrovic. Toppling kings in multipartite tournaments by introducing new kings. *Discrete Mathematics*, 310(19):2550–2554, 2010.

[Brill *et al.*, 2020] Markus Brill, Ulrike Schmidt-Kraepelin, and Warut Suksompong. Refining tournament solutions via margin of victory. In *AAAI*, pages 1862–1869, 2020.

[Brill *et al.*, 2021] Markus Brill, Ulrike Schmidt-Kraepelin, and Warut Suksompong. Margin of victory in tournaments: Structural and experimental results. In *AAAI*, 2021.

[Fey, 2008] Mark Fey. Choosing from a large tournament. *Social Choice and Welfare*, 31(2):301–309, 2008.

[Fishburn, 1977] Peter C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

[Fisher and Ryan, 1995] David C. Fisher and Jennifer Ryan. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2):217–236, 1995.

[Frank, 1968] O. Frank. Stochastic competition graphs. *Review of the International Statistical Institute*, 36(3):319–326, 1968.

[Good, 1971] Irving John Good. A note on Condorcet sets. *Public Choice*, 10(1):97–101, 1971.

[Kim and Vassilevska Williams, 2015] Michael P. Kim and Virginia Vassilevska Williams. Fixing tournaments for kings, chokers, and more. In *IJCAI*, pages 561–567, 2015.

[Kim *et al.*, 2017] Michael P. Kim, Warut Suksompong, and Virginia Vassilevska Williams. Who can win a single-elimination tournament? *SIAM Journal on Discrete Mathematics*, 31(3):1751–1764, 2017.

[Laslier, 1997] Jean-François Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.

[Łuczak *et al.*, 1996] Tomasz Łuczak, Andrzej Ruciński, and Jacek Gruszka. On the evolution of a random tournament. *Discrete Mathematics*, 148(1–3):311–316, 1996.

[Manurangsi and Suksompong, 2021] Pasin Manurangsi and Warut Suksompong. Generalized kings and single-elimination winners in random tournaments. *CoRR*, abs/2105.00193, 2021.

[Mattei and Walsh, 2016] Nicholas Mattei and Toby Walsh. Empirical evaluation of real world tournaments. *CoRR*, abs/1608.01039, 2016.

[Maurer, 1980] Stephen B. Maurer. The king chicken theorems. *Mathematics Magazine*, 53(2):67–80, 1980.

[Megiddo and Vishkin, 1988] Nimrod Megiddo and Uzi Vishkin. On finding a minimum dominating set in a tournament. *Theoretical Computer Science*, 61(2–3):307–316, 1988.

[Miller, 1977] Nicholas R. Miller. Graph-theoretic approaches to the theory of voting. *American Journal of Political Science*, 21(4):769–803, 1977.

[Miller, 1980] 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.

[Moon and Moser, 1962] John W. Moon and Leo Moser. Almost all tournaments are irreducible. *Canadian Mathematical Bulletin*, 5(1):61–65, 1962.

[Moulin, 1986] Hervé Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.

[Petrovic and Thomassen, 1991] Vojislav Petrovic and Carsten Thomassen. Kings in k-partite tournaments. *Discrete Mathematics*, 98(3):237–238, 1991.

[Saile and Suksompong, 2020] Christian Saile and Warut Suksompong. Robust bounds on choosing from large tournaments. *Social Choice and Welfare*, 54(1):87–110, 2020.

[Schwartz, 1972] Thomas Schwartz. Rationality and the myth of the maximum. *Noûs*, 6(2):97–117, 1972.

[Scott and Fey, 2012] Alex Scott and Mark Fey. The minimal covering set in large tournaments. *Social Choice and Welfare*, 38(1):1–9, 2012.

[Suksompong, 2021] Warut Suksompong. Tournaments in computational social choice: Recent developments. In *IJCAI*, 2021.

[Tan, 2006] B. P. Tan. On the 3-kings and 4-kings in multipartite tournaments. *Discrete Mathematics*, 306(21):2703–2710, 2006.

[Vassilevska Williams, 2010] Virginia Vassilevska Williams. Fixing a tournament. In *AAAI*, pages 895–900, 2010.

[Vassilevska Williams, 2016] Virginia Vassilevska Williams. Knockout tournaments. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 19, pages 453–474. Cambridge University Press, 2016.