Margin of Victory in Tournaments: Structural and Experimental Results

Markus Brill,[1] Ulrike Schmidt-Kraepelin,[1] Warut Suksompong[2]

[1] Research Group Efficient Algorithms, TU Berlin, Germany
[2] School of Computing, National University of Singapore, Singapore
{brill,u.schmidt-kraepelin}@tu-berlin.de, warut@comp.nus.edu.sg

Abstract

Tournament solutions are standard tools for identifying winners based on pairwise comparisons between competing alternatives. The recently studied notion of *margin of victory (MoV)* offers a general method for refining the winner set of any given tournament solution, thereby increasing the discriminative power of the solution. In this paper, we reveal a number of structural insights on the MoV by investigating fundamental properties such as monotonicity and consistency with respect to the covering relation. Furthermore, we provide experimental evidence on the extent to which the MoV notion refines winner sets in tournaments generated according to various stochastic models.

1 Introduction

Tournaments serve as a practical tool for modeling scenarios involving a set of alternatives along with pairwise comparisons between them. Perhaps the most common example of a tournament is a round-robin sports competition, where every pair of teams play each other once and there is no tie in match outcomes. Another application, typical especially in the social choice literature, concerns elections: here, alternatives represent election candidates, and pairwise comparisons capture the majority relation between pairs of candidates. In order to select the "winners" of a tournament in a consistent manner, numerous methods—known as *tournament solutions*—have been proposed. Given the ubiquity of tournaments, it is hardly surprising that tournament solutions have drawn substantial interest from researchers in the past few decades (Laslier 1997; Woeginger 2003; Hudry 2009; Aziz et al. 2015; Dey 2017; Brandt, Brill, and Harrenstein 2018; Brandt et al. 2018; Han and van Deemen 2019).

While tournament solutions are useful for selecting the best alternatives according to various desiderata, several solutions suffer from the setback that they tend to choose large winner sets. For instance, Fey (2008) showed that the top cycle, the uncovered set, and the Banks set are likely to include all alternatives in a large random tournament. To address this issue, we recently introduced the notion of *margin of victory (MoV)* for tournaments (Brill, Schmidt-Kraepelin, and Suksompong 2020b). The MoV of a winner is defined as the minimum number of pairwise comparisons that need to be reversed in order for the winner to drop out of the winner set. Analogously, the MoV of a non-winner is defined as the negative of the minimum number of comparisons that must be reversed for it to enter the winner set.[1] In addition to refining tournament solutions, the notion also has a natural interpretation in terms of bribery and manipulation: the MoV of an alternative reflects the cost of bribing voters or manipulating match outcomes so that the status of the alternative changes with respect to the winner set. For a number of common tournament solutions, we studied the complexity of computing the MoV and provided bounds on its values for both winners and non-winners (Brill et al. 2020b).

Our previous results paint an initial picture on the properties of the MoV in tournaments. Nevertheless, several important questions about the notion remain unanswered from that work. For each tournament solution, how many different values does the MoV take on average? How large is the set of alternatives with the highest MoV in a random tournament? If two alternatives dominate the same number of other alternatives, for which tournament solutions is it the case that the MoV of both alternatives must be equal? If an alternative "covers" another (i.e., the former alternative dominates the latter along with all alternatives that the latter dominates), for which tournament solutions is it always true that the MoV of the former alternative is at least that of the latter? In this paper, we provide a comprehensive answer to these questions for several common tournament solutions using axiomatic and probabilistic analysis (Section 3) as well as through experiments (Section 4).

1.1 Related Work

Despite their origins in social choice theory, tournament solutions have found applications in a wide range of areas including game theory (Fisher and Ryan 1995), webpage ranking (Brandt and Fischer 2007), dueling bandit problems (Ramamohan, Rajkumar, and Agarwal 2016), and philosophical decision theory (Podgorski 2020). As is the case for social choice theory in general, early studies of tournament solutions were primarily based on the axiomatic approach. With

[1] In our previous paper (Brill et al. 2020b), we considered a more general setting where each pairwise comparison can have a weight representing the cost of reversing it, but here we will focus on the unweighted setting.

the rise of computational social choice in the past fifteen years or so, tournament solutions have also been thoroughly examined from an algorithmic perspective. For an overview of the literature, we refer to the surveys of Laslier (1997), Hudry (2009), and Brandt, Brill, and Harrenstein (2016).

While we introduced the MoV concept for tournament solutions (Brill et al. 2020b), similar concepts have been applied to a large number of settings, perhaps most notably voting. In addition, various forms of bribery and manipulation have been considered for both elections and sports tournaments. We refer to our previous paper for relevant references, but note here that the MoV continues to be a popular concept in recent work, for example in the context of sports modeling (Kovalchik 2020), election control (Castiglioni, Ferraioli, and Gatti 2020), and political and educational districting (Stoica et al. 2020). Yang and Guo (2017) gave a parameterized complexity result for the decision version of computing the MoV with respect to the uncovered set.

The discriminative power of tournament solutions has been studied both analytically and experimentally. As we mentioned earlier, Fey (2008) showed that in a large tournament drawn uniformly at random, the top cycle, the uncovered set, and the Banks set are unlikely to exclude any alternative. Scott and Fey (2012) established an analogous result for the minimal covering set, while Fisher and Ryan (1995) proved that the bipartisan set selects half of the alternatives on average. Saile and Suksompong (2020) extended some of these results to more general probability distributions, and Brandt and Seedig (2016) performed experiments using both stochastic models and empirical data.

2 Preliminaries

A *tournament* $T = (V, E)$ is a directed graph in which exactly one directed edge exists between any pair of vertices. The vertices of T, denoted by $V(T)$, are often referred to as *alternatives*, and their number $n := |V(T)|$ is referred to as the *size* of T. The set of directed edges of T, denoted by $E(T)$, represents an asymmetric and connex *dominance relation* between the alternatives. An alternative x is said to *dominate* another alternative y if $(x, y) \in E(T)$ (i.e., there is a directed edge from x to y). When the tournament is clear from the context, we often write $x \succ y$ to denote $(x, y) \in E(T)$. By definition, for each pair x, y of distinct alternatives, either x dominates y ($x \succ y$) or y dominates x ($y \succ x$), but not both. The dominance relation can be extended to sets of alternatives by writing $X \succ Y$ if $x \succ y$ for all $x \in X$ and all $y \in Y$.

For a given tournament T and an alternative $x \in V(T)$, the *dominion* of x, denoted by $D(x)$, is the set of alternatives y such that $x \succ y$. Similarly, the set of *dominators* of x, denoted by $\overline{D}(x)$, is the set of alternatives y such that $y \succ x$. The *outdegree* of x is denoted by $\text{outdeg}(x) = |D(x)|$, and the *indegree* of x by $\text{indeg}(x) = |\overline{D}(x)|$. For any $x \in V(T)$, it holds that $\text{outdeg}(x) + \text{indeg}(x) = n - 1$. An alternative $x \in V(T)$ is said to be a *Condorcet winner* in T if it dominates every other alternative (i.e., $\text{outdeg}(x) = n-1$), and a *Condorcet loser* in T if it is dominated by every other alternative (i.e., $\text{outdeg}(x) = 0$). A tournament is *regular* if all alternatives have the same outdegree. A regular tournament exists for every odd size, but not for any even size.

2.1 Tournament Solutions

A *tournament solution* is a function that maps each tournament to a nonempty subset of its alternatives, usually referred to as the set of *winners* or the *choice set*. A tournament solution must not distinguish between isomorphic tournaments; in particular, if there is an automorphism that maps an alternative x to another alternative y in the same tournament, any tournament solution must either choose both x and y or neither of them. The set of winners of a tournament T with respect to a tournament solution S is denoted by $S(T)$. The tournament solutions considered in this paper are as follows:

- The *Copeland set* (CO) is the set of alternatives with the largest outdegree. The outdegree of an alternative is also referred to as its *Copeland score*.
- The *top cycle* (TC) is the (unique) nonempty smallest set X of alternatives such that $X \succ V(T) \setminus X$. Equivalently, TC is the set of alternatives that can reach every other alternative via a directed path.
- The *uncovered set* (UC), is the set of alternatives that are not "covered" by any other alternative. An alternative x is said to *cover* another alternative y if $D(y) \subseteq D(x)$. Equivalently, UC is the set of alternatives reaching every other alternative via a directed path of length at most two.
- The set of *k-kings*, for an integer $k \geq 3$, is the set of alternatives that can reach every other alternative via a directed path of length at most k.
- The *Banks set* (BA) is the set of alternatives that appear as the Condorcet winner of some transitive subtournament that cannot be extended.[2]

All of these tournament solutions satisfy *Condorcet-consistency*, meaning that whenever a Condorcet winner exists, it is chosen as the unique winner.

It is clear from the definitions that UC (the set of "2-kings") is contained in the set of k-kings for any $k \geq 3$, which is in turn a subset of TC (the set of "$(n-1)$-kings", as any directed path has length at most $n - 1$). Moreover, both CO and BA are contained in UC (Laslier 1997).

Given a tournament T and an edge $e = (x, y) \in E(T)$, we let $\overline{e} := (y, x)$ denote its reversal. Denote by T^e the tournament that results from T when reversing e. A tournament solution S is said to be *monotonic* if for any edge $e = (y, x) \in E(T)$,

$$x \in S(T) \quad \text{implies} \quad x \in S(T^e).$$

In other words, a tournament solution is monotonic if a winner remains in the choice set whenever its dominion is enlarged (while everything else is unchanged). Equivalently, monotonicity means that a non-winner remains outside of the choice set whenever it becomes dominated by an additional alternative.

[2]We say that an alternative $x \in V(T) \setminus V(T')$ *extends* a transitive subtournament T' if x dominates all alternatives in T'.

2.2 Margin of Victory

For a set of edges $R \subseteq E(T)$ of a tournament T, we define $\overline{R} := \{\overline{e} : e \in R\}$. Denote by T^R the tournament that results from T when reversing all edges in R, i.e., $V(T^R) = V(T)$ and $E(T^R) = (E(T) \setminus R) \cup \overline{R}$.

Fix a tournament solution S and consider a tournament T. An edge set $R \subseteq E(T)$ is called a *destructive reversal set (DRS)* for $x \in S(T)$ if $x \notin S(T^R)$. Analogously, R is called a *constructive reversal set (CRS)* for $x \in V(T) \setminus S(T)$ if $x \in S(T^R)$. The *margin of victory* of $x \in S(T)$ is given by

$$\mathsf{MoV}_S(x, T) = \min\{|R| : R \text{ is a DRS for } x \text{ in } T\},$$

and for $x \notin S(T)$ it is given by

$$\mathsf{MoV}_S(x, T) = -\min\{|R| : R \text{ is a CRS for } x \text{ in } T\}.$$

By definition, $\mathsf{MoV}_S(x, T)$ is a positive integer if $x \in S(T)$, and a negative integer otherwise.[3]

It follows from the definition of MoV that edge reversals have limited effects on the MoV value of alternatives: If a single edge e of a tournament T is reversed, then $\mathsf{MoV}_S(x, T)$ and $\mathsf{MoV}_S(x, T^e)$ differ by at most 1, unless x is a winner in exactly one of the two tournaments T and T^e (in which case $|\mathsf{MoV}_S(x, T) - \mathsf{MoV}_S(x, T^e)| = 2$).

Furthermore, MoV values behave monotonically with respect to edge reversals, provided the underlying tournament solution is monotonic.

Proposition 1. *Let S be a monotonic tournament solution and consider two tournaments T and T^e, where $e = (y, x) \in E(T)$. Then, $\mathsf{MoV}_S(x, T^e) \geq \mathsf{MoV}_S(x, T)$.*

All omitted proofs can be found in the full version of this paper (Brill, Schmidt-Kraepelin, and Suksompong 2020a).

3 Structural Results

In this section we provide a number of results relating the MoV notion to structural properties of the tournament in question. In particular, we identify conditions on tournament solutions ensuring that the corresponding MoV values are consistent with the covering relation (Section 3.1) and we examine the relationship between MoV values and Copeland scores (Sections 3.2 and 3.3). Our results are summarized in Table 1.

3.1 Cover-Consistency

Recall from Section 2 that an alternative x covers another alternative y if $D(y) \subseteq D(x)$. In particular, this implies that x dominates y (as otherwise $x \in D(y)$). The covering relation, which forms the basis for defining the uncovered set UC, is transitive and has a close connection to Pareto dominance in voting settings (Brandt, Geist, and Harrenstein 2016).

Intuitively, if x covers y, there is a strong argument that x is preferable to y. We show that for all of the tournament solutions that we consider, their corresponding MoV values are indeed consistent with this intuition.

	cover-cons.	strong deg.-cons.	degree-cons.	equal-deg.-cons.
MoV_{CO}	✓	✗	✓	✗
MoV_{TC}	✓	✓	✓	✓
MoV_{UC}	✓	✗	✗	✗
$\mathsf{MoV}_{k\text{-kings}}$	✓	✗	✗	✗
MoV_{BA}	✓	✗	✗	✗

Table 1: Consistency properties of the margin of victory for the tournament solutions CO, TC, UC, k-kings, and BA.

Definition 2. For a tournament solution S, we say that MoV_S is *cover-consistent* if, for any tournament T and any alternatives $x, y \in V(T)$, x covers y implies $\mathsf{MoV}_S(x, T) \geq \mathsf{MoV}_S(y, T)$.

We introduce a new property that will be useful for showing that a tournament solution is cover-consistent.

Definition 3. A tournament solution S is said to be *transfer-monotonic* if for any edges $(y, z), (z, x) \in E(T)$,

$$x \in S(T) \quad \text{implies} \quad x \in S(T'),$$

where T' is the tournament obtained from T by reversing edges (y, z) and (z, x).

In other words, if an alternative x is chosen, then it remains chosen when an alternative z is "transferred" from the dominion $D(y)$ of another alternative y to its dominion $D(x)$.

We show that monotonicity and transfer-monotonicity together imply cover-consistency of the margin of victory.

Lemma 4. *If a tournament solution S is monotonic and transfer-monotonic, then MoV_S satisfies cover-consistency.*

Proof. Let S be a monotonic and transfer-monotonic tournament solution, and suppose that alternative x covers another alternative y in a tournament T. We will show that $\mathsf{MoV}_S(x, T) \geq \mathsf{MoV}_S(y, T)$.

If $x \in S(T)$ and $y \notin S(T)$, the statement holds trivially since $\mathsf{MoV}_S(x, T) > 0 > \mathsf{MoV}_S(y, T)$. Suppose for contradiction that $x \notin S(T)$ and $y \in S(T)$. Consider the tournament T' obtained from T by reversing the edge (x, y) as well as edges $(x, z), (z, y)$ for each $z \in D(x) \setminus (D(y) \cup \{y\})$. By monotonicity and transfer-monotonicity, $y \in S(T')$. However, tournaments T and T' are isomorphic, and there is an isomorphism that maps $x \in T$ to $y \in T'$. Since $x \notin S(T)$, we must have $y \notin S(T')$, a contradiction.

The remaining two cases are $x, y \in S(T)$ and $x, y \notin S(T)$; both can be handled in an analogous manner, so let us focus on the latter case. It suffices to show that given any CRS for y of minimum size, we can construct a CRS of smaller or equal size for x. Let R_y be a CRS for y of minimum size; we will construct a CRS R_x for x such that $|R_x| \leq |R_y|$.

Let $A = V(T) \setminus \{x, y\}$, and partition A into three sets $A_1 = D(y)$, $A_2 = D(x) \setminus (D(y) \cup \{y\})$, and $A_3 = \overline{D}(x)$;

[3] The only exception is the degenerate case where S selects all alternatives for all tournaments of some size n; in this case we define $\mathsf{MoV}_S(x, T) = \infty$ for all alternatives x and all tournaments T of that size. For ease of exposition, we will assume for the rest of the paper that the degenerate case does not occur, but all of our results still hold even when this case occurs.

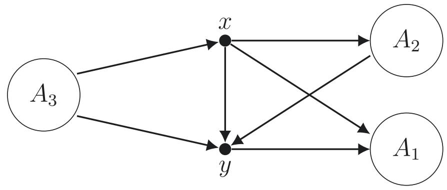

Figure 1: Illustration of the proof of Lemma 4.

see Figure 1 for an illustration. For any edge in R_y between two alternatives of A, we add the same edge to R_x. We do not add the edge (x, y) regardless of whether it is present in R_y. Each remaining edge in R_y is between an alternative in A and one of x, y. Note that $(y, a) \notin R_y$ for any $a \in A$—otherwise, by monotonicity, removing such an edge would keep R_y a CRS for y, contradicting the minimality of R_y.

For each $a \in A$, we add further edges to R_x as follows.

- For $a \in A_1$:
  - If $(x, a) \in R_y$, add (y, a) to R_x.
- For $a \in A_2$:
  - If $(x, a) \in R_y$ but $(a, y) \notin R_y$, add (x, a) to R_x.
  - If $(x, a) \notin R_y$ but $(a, y) \in R_y$, add (a, y) to R_x.
- For $a \in A_3$:
  - If $(a, x) \in R_y$, add (a, y) to R_x.
  - If $(a, y) \in R_y$, add (a, x) to R_x.

Clearly, $|R_x| \leq |R_y|$, and we have $y \in S(T^{R_y})$ by definition of R_y. From T^{R_y}, we reverse the edge (x, y) if it is present, and for $a \in A_2$ such that both $(x, a), (a, y) \notin R_y$, we reverse (x, a) and (a, y). Let T' be the resulting tournament. By monotonicity and transfer-monotonicity, we have $y \in S(T')$. However, one can verify that there exists an isomorphism from T' to T^{R_x} that maps x to y, y to x, and every other alternative a to itself. Since $y \in S(T')$, we must have $x \in S(T^{R_x})$, meaning that R_x is indeed a CRS for x. □

In the full version (Brill et al. 2020a), we show that neither monotonicity nor transfer-monotonicity can be dropped from the condition of Lemma 4. This also means that neither of the two properties implies the other.

We now show that all tournament solutions we consider in this paper satisfy both monotonicity and transfer monotonicity, thereby implying that their MoV functions are cover-consistent.

Proposition 5. *CO, UC, TC, k-kings, and BA satisfy monotonicity.*

Proof. It is already known that CO, UC, TC, and BA are monotonic (Laslier 1997; Brandt, Brill, and Harrenstein 2016); hence, it remains to establish the monotonicity of k-kings. Let x be a k-king in tournament T, and suppose that T' is the tournament obtained by reversing an edge (y, x). Since any path of length at most k from x to another alternative in T cannot contain the edge (y, x), the same path is also present in T'. Hence x is also a k-king in T'. □

Proposition 6. *CO, UC, TC, k-kings, and BA satisfy transfer-monotonicity.*

Proof. We prove the statement for CO; the other tournament solutions are handled in the full version. If $x \in CO(T)$ and edges (y, z) and (z, x) are reversed, then the outdegree of x increases by 1, that of y decreases by 1, while all other alternatives have the same outdegree as before. Hence x is in the Copeland set of the new tournament. □

Lemma 4 and Propositions 5 and 6 together imply the following:

Theorem 7. *For each $S \in \{CO, TC, UC, k\text{-kings}, BA\}$, MoV_S satisfies cover-consistency.*

In light of Theorem 7, one may wonder whether a stronger property, in which x covers y implies the strict inequality $\mathsf{MoV}_S(x) > \mathsf{MoV}_S(y)$, can also be achieved. However, the answer is negative for all Condorcet-consistent tournament solutions, including all solutions that we consider. Indeed, in a transitive tournament $x \succ y \succ z$ of size 3, such a solution only selects x. But since all three alternatives are chosen when they form a cycle (due to symmetry), both y and z can be brought into the winner set by reversing only one edge, so $\mathsf{MoV}_S(y) = -1 = \mathsf{MoV}_S(z)$ even though y covers z.

3.2 Degree-Consistency

Given a tournament solution S and a tournament T, the MoV_S values yield a natural *ranking* (possibly including ties) of the alternatives in T, where alternative x is ranked higher than y whenever $\mathsf{MoV}_S(x, T) > \mathsf{MoV}_S(y, T)$. We are interested in how closely this ranking by MoV values resembles the ranking by Copeland scores, according to which x is ranked higher than y if $\text{outdeg}(x) > \text{outdeg}(y)$.

Definition 8. For a tournament solution S, we say that MoV_S is

- *degree-consistent* if, for any tournament T and any alternatives $x, y \in V(T)$, $\text{outdeg}(x) > \text{outdeg}(y)$ implies $\mathsf{MoV}_S(x, T) \geq \mathsf{MoV}_S(y, T)$;
- *equal-degree-consistent* if, for any tournament T and any alternatives $x, y \in V(T)$, $\text{outdeg}(x) = \text{outdeg}(y)$ implies $\mathsf{MoV}_S(x, T) = \mathsf{MoV}_S(y, T)$; and
- *strong degree-consistent* if, for any tournament T and any alternatives $x, y \in V(T)$, $\text{outdeg}(x) \geq \text{outdeg}(y)$ implies $\mathsf{MoV}_S(x, T) \geq \mathsf{MoV}_S(y, T)$.

It follows from the definitions that MoV_S is strong degree-consistent if and only if it is both degree-consistent and equal-degree-consistent. Observe also that cover-consistency is implied by degree-consistency.

We remark that these properties are not necessarily desirable from a normative perspective: Whereas the ranking implied by a strongly degree-consistent MoV function merely represents a coarsening of the straightforward ranking by outdegree, we are often interested in tournament solutions that take more structure of the tournament into account and, as a consequence, have MoV functions that may violate (equal-)degree-consistency. Indeed, since degree-consistent

MoV functions are in line with Copeland scores, their significance is somewhat limited and there would be little additional value derived from the MoV computations, which in some cases are much more involved than simply calculating Copeland scores.

We start by assessing the degree-consistency of MoV_{CO}.

Proposition 9. MoV_{CO} *satisfies degree-consistency, but not equal-degree-consistency.*

Next, we consider the top cycle. Recall that, for a given tournament T of size n, TC coincides with k-kings for $k = n-1$. In order to show that MoV_{TC} satisfies strong degree-consistency, we need two lemmas (one of which is already known). We first introduce some notation.

Given a tournament T and distinct alternatives $x, y \in V(T)$, an edge set $R \subseteq E(T)$ is said to be a *k-length bounded x-y-cut* if, once R is removed, every path from x to y has length strictly greater than k. Denote by $\text{min-cut}_k(x, y)$ the size of a smallest k-length bounded x-y-cut. A set R is said to be a *k-length bounded x-cut* if it is a k-length bounded x-y-cut for some $y \in T$.

Lemma 10 (Lemma 4 by Brill et al. (2020b))**.** *For any $k \in \{2, 3, \ldots, n-1\}$, a set $R \subseteq E(T)$ is a minimum DRS for x with respect to k-kings if and only if R is a minimum k-length bounded x-cut in T.*

The next lemma establishes a surprisingly succinct relation between the sizes of the minimum cuts with respect to a pair of alternatives, and can be shown using the max-flow min-cut theorem. Define $\text{min-cut}(x, y) = \text{min-cut}_{n-1}(x, y)$.

Lemma 11. *Let T be a tournament and $x, y \in V(T)$. Then,*

$$\text{min-cut}(x, y) - \text{min-cut}(y, x) = \text{outdeg}(x) - \text{outdeg}(y).$$

Theorem 12. MoV_{TC} *satisfies strong degree-consistency.*

Proof. Fix a tournament T and let $x, y \in V(T)$ with $\text{outdeg}(x) \geq \text{outdeg}(y)$. First, we show that $x, y \in TC(T)$ constitutes the only non-trivial case. Since all alternatives in $TC(T)$ dominate all alternatives outside, it cannot be that $x \notin TC(T)$ and $y \in TC(T)$. If $x, y \notin TC(T)$, Brill et al. (2020b) showed that $\mathsf{MoV}_{TC}(x, T) = -1 = \mathsf{MoV}_{TC}(y, T)$. If $x \in TC(T)$ and $y \notin TC(T)$, then $\mathsf{MoV}_{TC}(x) > 0 > \mathsf{MoV}_{TC}(y)$.

Assume now that $x, y \in TC(T)$. Let R be a minimum DRS for x. By Lemma 10 with $k = n-1$, we know that R is a minimum x-t-cut for some $t \in V(T)$. We consider two cases. First, assume that R is also a y-t-cut. Then, a minimum y-t-cut $R' \subseteq E(T)$ satisfies $|R'| \leq |R|$, proving that $\mathsf{MoV}_{TC}(x, T) = |R| \geq |R'| \geq \mathsf{MoV}_{TC}(y, T)$. For the second case, assume that R is not a y-t-cut. Then, R needs to be an x-y-cut (since otherwise x can reach t via y), and therefore it must be a minimum x-y-cut. By Lemma 11, since $\text{outdeg}(x) \geq \text{outdeg}(y)$, for a minimum y-x-cut R' it holds that $|R| \geq |R'|$. Hence $\mathsf{MoV}_{TC}(x, T) = |R| \geq |R'| \geq \mathsf{MoV}_{TC}(y, T)$. $\square$

On the other hand, we show in the next three propositions that UC, BA, and k-kings do not satisfy any of the degree-consistency properties.

Proposition 13. MoV_{UC}, MoV_{BA}, *and* $\mathsf{MoV}_{k\text{-kings}}$ *(for constant $k \geq 3$) satisfy neither degree-consistency nor equal-degree-consistency.*

Corollary 14. $\mathsf{MoV}_{CO}, \mathsf{MoV}_{UC}, \mathsf{MoV}_{k\text{-kings}}$ *(for constant $k \geq 3$), and* MoV_{BA} *do not fulfill strong degree-consistency.*

3.3 A Probabilistic Result

In this section, we establish a simple formula for the MoV of TC and k-kings for $k \geq 4$ that works "with high probability", i.e., the probability that the formula holds converges to 1 as n grows. We assume that the tournament is generated using the *uniform random model*, where each edge is oriented in either direction with equal probability independently of other edges; this model has been studied, among others, by Fey (2008) and Scott and Fey (2012).

Theorem 15. *Let $S \in \{TC, k\text{-kings}\}$, where $4 \leq k \leq n-1$. Assume that a tournament T is generated according to the uniform random model. Then, with high probability, the following holds for all $x \in V(T)$ simultaneously:*

$$\mathsf{MoV}_S(x, T) = \min\left(\text{outdeg}(x), \min_{y \in V(T): y \neq x} \text{indeg}(y)\right).$$

Theorem 15 suggests that when tournaments are generated according to the uniform random model, MoV_{TC} and $\mathsf{MoV}_{k\text{-kings}}$ for $k \geq 4$ can likely be computed by a simple formula based on the degrees of the alternatives. In particular, even though the problem is computationally hard for $\mathsf{MoV}_{k\text{-kings}}$ for any constant $k \geq 4$ (Brill et al. 2020b), there exists an efficient heuristic that correctly computes the MoV value in most cases. In the full version (Brill et al. 2020a), we give an example showing that the heuristic is not always correct. More precisely, for any positive integer ℓ, we construct a tournament such that $\{\mathsf{MoV}_{TC}(x, T) \mid x \in V(T)\}$ contains the values $1, 2, \ldots, \ell$ whereas the formula in Theorem 15 predicts that all alternatives have the same (arbitrarily large) MoV_{TC} value.

At a high level, to prove this theorem, we first observe that by a result of Fey (2008), it is likely that $S(T) = V(T)$, i.e., all alternatives are chosen by S. In order to remove alternative x from the winner set, one option is to make it a Condorcet loser—this requires $\text{outdeg}(x)$ reversals—while another option is to make another alternative y a Condorcet winner—this requires $\text{indeg}(y)$ reversals. Hence, the left-hand side is at most the right-hand side. To establish that both sides are equal with high probability, we need to show that the aforementioned options are the best ones for making x a non-winner—by Lemma 10, this requires making some y unreachable from x in four steps. The intuition behind this claim is that the tournament resulting from the uniform random model is highly connected, with many paths of length at most four from x to y. Thus, if we want to make y unreachable from x, it is unlikely to be beneficial to destroy intermediate edges instead of edges adjacent to x or y.

4 Experiments

In order to better understand how MoV values of tournament solutions behave in practice, we conducted computational experiments using randomly generated tournaments.

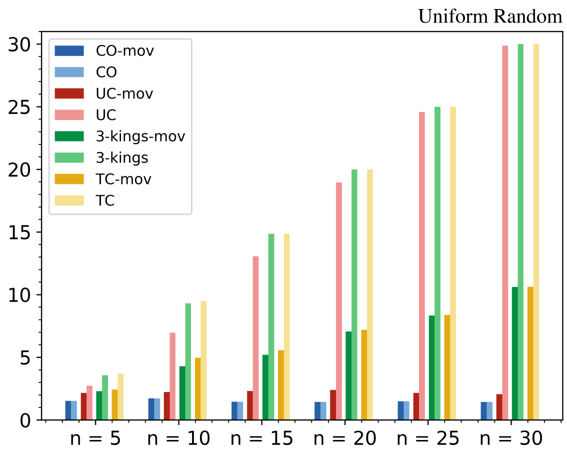

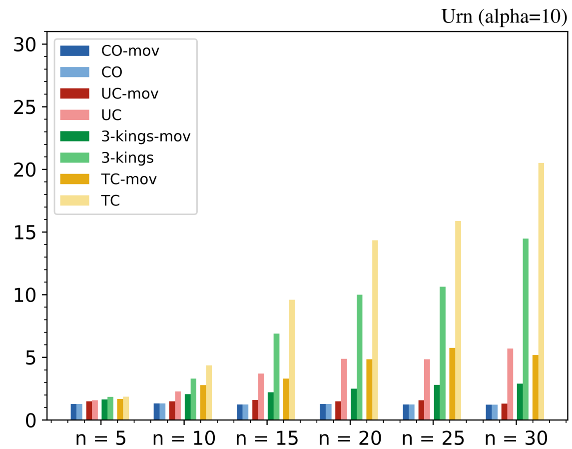

Figure 2: The illustrations show the average number of alternatives with maximum MoV value for the uniform random model (left) and the urn model (right), for different tournament solutions and sizes. For comparison, the average size of the entire winning set of the corresponding original tournament solution is depicted by a lighter shade.

For the sake of diversity of the generated instances, we implemented six different stochastic models to generate tournaments. To make our study comparable to the experiments presented by Brandt and Seedig (2016), we selected a similar set of stochastic models and parameterizations.

Given a tournament solution S and a tournament T, we are interested in

- the number $|\arg\max_{x\in V(T)} \mathsf{MoV}_S(x,T)|$ of alternatives with maximum MoV_S value, and
- the number $|\{\mathsf{MoV}_S(x,T) : x \in V(T)\}|$ of different MoV values taken by all alternatives in the tournament.

The first value directly measures the discriminative power of the refinement of S that only selects alternatives with a maximal MoV_S value, whereas the second value measures more generally the ability of the MoV notion to distinguish between the alternatives in a tournament.

Set-up We used six stochastic models to generate preferences: the uniform random model (which was used in Section 3.3), two variants of the *Condorcet noise model* (with and without voters), the *impartial culture* model, the Pólya-Eggenberger *urn model*, and the *Mallows* model. Detailed descriptions of these models can be found in the full version of our paper (Brill et al. 2020a).

For each stochastic model and each number of alternatives $n \in \{5, 10, 15, 20, 25, 30\}$, we sampled 100 tournaments. Using the methods described by Brill et al. (2020b), we implemented algorithms to calculate the MoV values for CO, UC, 3-kings, and TC. Due to their computational intractability, we did not implement procedures to calculate the MoV values for BA and k-kings for $k \geq 4$.

The experiments were carried out on a system with 1.4 GHz Quad-Core Intel Core i5 CPU, 8GB RAM, and macOS 10.15.2 operating system. The software was implemented in Python 3.7.7 and the libraries networkx 2.4, matplotlib 3.2.1, numpy 1.18.2, and pandas 1.0.3 were used. For implementing the Mallows and urn models, we utilized implementations contributed by Mattei and Walsh (2013). The code for our implementation can be found at http://github.com/uschmidtk/MoV.

Results For two representative stochastic models (uniform random and urn), Figure 2 depicts the average size of the set of alternatives with maximum MoV value, and Figure 3 shows the average number of unique MoV values. Results on the other four models (impartial culture, Mallows, and two variants of Condorcet noise) can be found in the full version (Brill et al. 2020a).

Observations The first observation we make is that $\mathsf{MoV}_{\text{3-kings}}$ behaves rather similarly to MoV_{TC}: the average number of alternatives with maximum MoV grows with increasing n, and this number is on average slightly less than half of the number of 3-kings and TC winners, respectively. However, this ratio becomes smaller for tournaments where the number of 3-kings or TC winners is already large. For example, when we only consider tournaments where the number of TC winners is greater than 10, only one-third of the TC winners have a maximum MoV_{TC} value on average; the same holds for 3-kings. However, a more detailed look at the experimental results show that for both 3-kings and TC, the set of alternatives with maximum MoV consists of only one alternative in around 73% of all instances, while in the remaining instances this set is typically large. This particular behavior for TC and the uniform random model can be explained by Theorem 15: With high probability, the MoV values for TC winners follow a specific formula based on the degrees, which leads to the set of alternatives with maximum MoV containing either a single al-

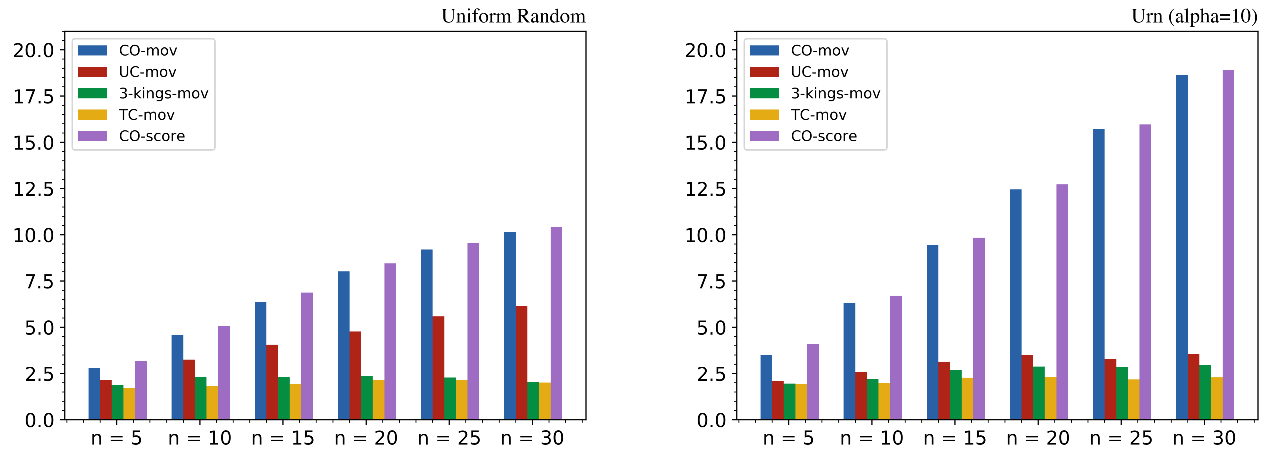

Figure 3: The illustrations show the average number of unique MoV values for the uniform random model (left) and the urn model (right), for different tournament solutions and sizes. For comparison, the average number of unique Copeland scores is shown in violet.

ternative or a large number of alternatives in most cases.[4] Our experiments show that this behavior is also present in tournaments generated by other stochastic models as well as for 3-kings; formalizing the behavior theoretically is an interesting future direction.

Our second main observation is that MoV$_{UC}$ behaves quite differently from MoV$_{\text{3-kings}}$ and MoV$_{TC}$. Most importantly, the number of UC winners with maximum MoV$_{UC}$ does *not* increase with a growing number of alternatives, but remains more or less constant for each stochastic model. For the uniform random model and the Condorcet noise models, this value is around 2, while it is roughly 1.4 for Mallows, the urn model, and the impartial culture model. As can be seen in Figure 2, the set of alternatives maximizing MoV$_{UC}$ is almost as discriminative as the Copeland set (all of whose alternatives maximize MoV$_{CO}$). However, we observe in Figure 3 that the number of unique values of the Copeland score is notably higher than that of MoV$_{UC}$. The latter is particularly low for models which tend to create tournaments with small UC, including Mallows, impartial culture and the urn model. Both of these effects can be explained by the observation that MoV$_{UC}$ is significantly better at distinguishing between UC winners than it is at distinguishing between UC non-winners.[5] As a consequence, tournaments with a small uncovered set generally give rise to a small number of unique MoV$_{UC}$ values.

[4]Indeed, if there is a unique Copeland winner, that winner will be the unique alternative with the largest MoV according to the formula. Otherwise, for several alternatives (including the Copeland winners), it can be the case that their MoV is equal to $\mathrm{indeg}(y)$ for a Copeland winner y.

[5]Brill et al. (2020b) showed that the smallest MoV$_{UC}$ value in a tournament is bounded below by $-\lceil \log_2(n) \rceil$, and that this bound is asymptotically tight. In our experiments, we observed that in most generated tournaments, the smallest MoV$_{UC}$ value is much higher than this lower bound, namely either -1 or -2.

5 Discussion

The recently introduced notion of margin of victory (MoV) provides a generic framework for refining any tournament solution. In this paper, we have contributed to the understanding of the MoV by providing not only structural insights but also experimental evidence regarding the extent to which it refines winner sets in stochastically generated tournaments. We established that the MoV is consistent with the covering relation for all considered tournament solutions. Moreover, we have identified a number of tournament solutions, including the uncovered set and the Banks set, for which the corresponding MoV values give insights into the structure of the tournament that go beyond simply comparing the outdegrees of alternatives, as witnessed by the fact that these MoV functions do not satisfy degree-consistency.

In our experiments, the MoV function corresponding to the uncovered set (UC) stands out for its discriminative power: not only is the set max-MoV$_{UC}$ (containing all alternatives with maximal MoV$_{UC}$ score) consistently small, but the number of distinct MoV$_{UC}$ scores is also relatively high in general. It is consequently tempting to suggest max-MoV$_{UC}$ as a new tournament solution. Besides its discriminative power and structural appeal, it can be computed efficiently (Brill et al. 2020b) and inherits Pareto optimality from the uncovered set, which it refines (Brandt, Geist, and Harrenstein 2016). However, a thorough axiomatic analysis of max-MoV$_{UC}$, as well as max-MoV$_S$ for other tournament solutions S, is still outstanding.

For tournaments with several highest-scoring alternatives (i.e., several alternatives whose minimal destructive reversal sets are of the same maximal size), the *number* of distinct destructive reversal sets may serve as a further criterion for distinguishing between winners. It would therefore be interesting to determine the complexity of computing such numbers, and also to study the size of the resulting refined winner set experimentally in future work.

Acknowledgments

This work was partially supported by the Deutsche Forschungsgemeinschaft under grant BR 4744/2-1, by the European Research Council (ERC) under grant number 639945 (ACCORD), and by an NUS Start-up Grant. We would like to thank the anonymous reviewers for their valuable feedback.

References

Aziz, H.; Brill, M.; Fischer, F.; Harrenstein, P.; Lang, J.; and Seedig, H. G. 2015. Possible and necessary winners of partial tournaments. *Journal of Artificial Intelligence Research* 54: 493–534.

Brandt, F.; Brill, M.; and Harrenstein, P. 2016. Tournament solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 3, 57–84. Cambridge University Press.

Brandt, F.; Brill, M.; and Harrenstein, P. 2018. Extending tournament solutions. *Social Choice and Welfare* 51(2): 193–222.

Brandt, F.; Brill, M.; Seedig, H. G.; and Suksompong, W. 2018. On the structure of stable tournament solutions. *Economic Theory* 65(2): 483–507.

Brandt, F.; and Fischer, F. 2007. PageRank as a weak tournament solution. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE)*, 300–305.

Brandt, F.; Geist, C.; and Harrenstein, P. 2016. A note on the McKelvey uncovered set and Pareto optimality. *Social Choice and Welfare* 46(1): 81–91.

Brandt, F.; and Seedig, H. G. 2016. On the discriminative power of tournament solutions. In *Selected Papers of the International Conference on Operations Research, OR2014*, 53–58.

Brill, M.; Schmidt-Kraepelin, U.; and Suksompong, W. 2020a. Margin of victory in tournaments: Structural and experimental results. Technical report, arXiv:2012.02657.

Brill, M.; Schmidt-Kraepelin, U.; and Suksompong, W. 2020b. Refining tournament solutions via margin of victory. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 1862–1869.

Castiglioni, M.; Ferraioli, D.; and Gatti, N. 2020. Election control in social networks via edge addition or removal. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 1878–1885.

Dey, P. 2017. Query complexity of tournament solutions. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2992–2998.

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

Fisher, D. C.; and Ryan, J. 1995. Tournament games and positive tournaments. *Journal of Graph Theory* 19(2): 217–236.

Han, W.; and van Deemen, A. 2019. A refinement of the uncovered set in tournaments. *Theory and Decision* 86(1): 107–121.

Hudry, O. 2009. A survey on the complexity of tournament solutions. *Mathematical Social Sciences* 57(3): 292–303.

Kovalchik, S. 2020. Extension of the Elo rating system to margin of victory. *International Journal of Forecasting* 36(4): 1329–1341.

Laslier, J.-F. 1997. *Tournament Solutions and Majority Voting*. Springer-Verlag.

Mattei, N.; and Walsh, T. 2013. PrefLib: a library for preference http://www.preflib.org. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 259–270.

Podgorski, A. 2020. Tournament decision theory. *Noûs* Forthcoming.

Ramamohan, S.; Rajkumar, A.; and Agarwal, S. 2016. Dueling bandits: Beyond Condorcet winners to general tournament solutions. In *Proceedings of the 29th Annual Conference on Neural Information Processing Systems (NIPS)*, 1253–1261.

Saile, C.; and Suksompong, W. 2020. Robust bounds on choosing from large tournaments. *Social Choice and Welfare* 54(1): 87–110.

Scott, A.; and Fey, M. 2012. The minimal covering set in large tournaments. *Social Choice and Welfare* 38(1): 1–9.

Stoica, A.-A.; Chakraborty, A.; Dey, P.; and Gummadi, K. 2020. Minimizing margin of victory for fair political and educational districting. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1305–1313.

Woeginger, G. J. 2003. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare* 20(3): 523–528.

Yang, Y.; and Guo, J. 2017. Possible winner problems on partial tournaments: a parameterized study. *Journal of Combinatorial Optimization* 33(3): 882–896.