

Possible and Necessary Winners of Partial Tournaments*

Haris Aziz, Markus Brill, Felix Fischer,
Paul Harrenstein, Jérôme Lang, and Hans Georg Seedig

Abstract

We study the problem of computing possible and necessary winners for partially specified weighted and unweighted tournaments. This problem arises naturally in elections with incompletely specified votes, partially completed sports competitions, and more generally in any scenario where the outcome of some pairwise comparisons is not yet fully known. We specifically consider a number of well-known solution concepts—including the uncovered set, Borda, ranked pairs, and maximin—and show that for most of them possible and necessary winners can be identified in polynomial time. These positive algorithmic results stand in sharp contrast to earlier results concerning possible and necessary winners given partially specified preference profiles.

1 Introduction

Many multi-agent situations can be modeled and analyzed using weighted or unweighted tournaments. Prime examples are voting scenarios in which pairwise comparisons between alternatives are decided by majority rule and sports competitions that are organized as round-robin tournaments. Other application areas include webpage and journal ranking, biology, psychology, and AI (also see [6], and the references therein). More generally, tournaments and tournament solutions are used as a mathematical tool for the analysis of all kinds of situations where a choice among a set of alternatives has to be made exclusively on the basis of pairwise comparisons.

When choosing from a tournament, relevant information may only be partly available. This could be because some preferences are yet to be elicited, some matches yet to be played, or certain comparisons yet to be made. In such cases, it is natural to speculate which are the potential and inevitable outcomes on the basis of the information already at hand.

For complete tournaments, a number of attractive solution concepts have been proposed (see, e.g., [6, 17]). Given any such solution concept S , *possible winners* of a partial tournament G are defined as alternatives that are selected by S in *some* completion of G , and *necessary winners* are alternatives that are selected in *all* completions. By a completion we here understand a complete tournament extending G .

In this paper we address the computational complexity of identifying the possible and necessary winners for a number of solution concepts whose winner determination problem for complete tournaments is tractable. We consider four of the most common tournament solutions—namely, Condorcet winners (*COND*), the Copeland solution (*CO*), the top cycle (*TC*), and the uncovered set (*UC*)—and three common solutions for weighted tournaments—Borda (*BO*), maximin (*MM*) and ranked pairs (*RP*). For each of these solution concepts, we characterize the complexity of the following problems: deciding whether a given alternative is a possible winner (*PW*), deciding whether a given alternative is a necessary winner (*NW*), and deciding whether a given subset of alternatives equals the set

*A previous version of this paper has been accepted at AAMAS-2012. New results include speeding up pseudo-polynomial time algorithms to strongly polynomial time for PWS_{BO} (Thm. 8) and PWS_{MM} (Thm. 12).

S	PW_S		NW_S		PWS_S	
COND	in P	[16]	in P	[16]	in P	(Thm. 1)
CO	in P	(Thm. 2) ^a	in P	(Thm. 2) ^a	in P	(Thm. 2)
TC	in P	[16] ^a	in P	[16]	in P	(Thm. 3)
UC	in P	(Thm. 4)	in P	(Thm. 5)	NP-C	(Thm. 6)
BO	in P	(Thm. 7) ^a	in P	(Thm. 9)	in P	(Thm. 8)
MM	in P	(Thm. 10) ^a	in P	(Thm. 11)	in P	(Thm. 12)
RP	NP-C	(Thm. 13)	coNP-C	(Thm. 14)	NP-C	(Cor. 1)

^a This P-time result contrasts with the intractability of the same problem for partial preference profiles [16, 25].

Table 1: Complexity of computing possible winners (PW) and necessary winners (NW) and of checking whether a given subset of alternatives is a possible winning set (PWS) under different solution concepts given partial tournaments.

of winners in some completion (PWS). These problems can be challenging, as even unweighted partial tournaments may allow for an exponential number of completions. Our results are encouraging, in the sense that most of the problems can be solved in polynomial time. Table 1 summarizes our findings.

Similar problems have been considered before. For Condorcet winners, voting trees and the top cycle, it was already shown that possible and necessary winners are computable in polynomial time [16, 19, 20]. The same holds for computing possible Copeland winners that were considered in the context of sports tournaments [8].

A more specific setting that is frequently considered within the area of computational social choice differs from our setting in a subtle but important way that is worth being pointed out. There, tournaments are assumed to arise from pairwise majority comparisons on the basis of a profile of individual voters' preferences.¹ Since a *partial* preference profile R need not conclusively settle every majority comparison, it may give rise to a *partial* tournament only. There are two natural ways to define possible and necessary winners for a partial preference profile R and solution concept S . The first is to consider the completions of the incomplete tournament $G(R)$ corresponding to R and the winners under S in these. This is covered by our more general setting. The second is to consider the completions of R and the winners under S in the corresponding tournaments.² Since every tournament corresponding to a completion of R is also a completion of $G(R)$ but not necessarily the other way round, the second definition gives rise to a *stronger* notion of a possible winner and a *weaker* notion of a necessary winner. Interestingly, and in sharp contrast to our results, determining these stronger possible and weaker necessary winners is computationally hard for many voting rules [16, 25].

In the context of this paper, we do not assume that tournaments arise from majority comparisons in voting or from any other specific procedure. This approach has a number of advantages. Firstly, it matches the diversity of settings to which tournament solutions are applicable, which goes well beyond social choice and voting. For instance, our results also apply to a question commonly encountered in sports competitions, namely, which teams can still win the cup and which future results this depends on (see, e.g., [8, 14]). Secondly, (partial) tournaments provide an informationally sustainable way of representing the relevant aspects of many situations while maintaining a workable level of abstraction and concise-

¹See, e.g., [1, 2, 15, 24, 25] for the basic setting, [3] for parameterized complexity results, [12, 13] for probabilistic settings, and [7, 26] for settings with a variable set of alternatives.

²These two ways of defining possible and necessary winners are compared (both theoretically and experimentally) in [16, 20] for three solution concepts: Condorcet winners, voting trees and the top cycle.

ness. For instance, in the social choice setting described above, the partial tournament induced by a partial preference profile is a much more succinct piece of information than the preference profile itself. Finally, specific settings may impose restrictions on the feasible extensions of partial tournaments. The positive algorithmic results in this paper can be used to efficiently approximate the sets of possible and necessary winners in such settings, where the corresponding problems may be intractable. The voting setting discussed above serves to illustrate this point.

2 Preliminaries

A *partial tournament* is a pair $G = (V, E)$ where V is a finite set of alternatives and $E \subseteq V \times V$ an asymmetric relation on V , i.e., $(x, y) \in E$ implies $(y, x) \notin E$. If $(x, y) \in E$ we say that x *dominates* y . A (*complete*) *tournament* T is a partial tournament (V, E) for which E is also complete, i.e., either $(x, y) \in E$ or $(y, x) \in E$ for all distinct $x, y \in V$. We denote the class of complete tournaments by \mathcal{T} .

Let $G = (V, E)$ be a partial tournament. Another partial tournament $G' = (V', E')$ is called an *extension* of G , denoted $G \leq G'$, if $V = V'$ and $E \subseteq E'$. If E' is complete, G' is called a *completion* of G . We write $[G]$ for the set of completions of G , i.e., $[G] = \{T \in \mathcal{T} : G \leq T\}$.

For each $x \in V$, we define the *dominion* of x in G by $D_G^+(x) = \{y \in V : (x, y) \in E\}$, and the *dominators* of x in G by $D_G^-(x) = \{y \in V : (y, x) \in E\}$. For $X \subseteq V$, we let $D_G^+(X) = \bigcup_{x \in X} D_G^+(x)$ and $D_G^-(X) = \bigcup_{x \in X} D_G^-(x)$.

For given $G = (V, E)$ and $X \subseteq V$, we further write $E^{X \rightarrow}$ for the set of edges obtained from E by adding all missing edges from alternatives in X to alternatives not in X , i.e.,

$$E^{X \rightarrow} = E \cup \{(x, y) \in X \times V : y \notin X \text{ and } (y, x) \notin E\}.$$

We use $E^{X \leftarrow}$ as an abbreviation for $E^{V \setminus X \rightarrow}$, and respectively write $E^{x \rightarrow}$, $E^{x \leftarrow}$, $G^{X \rightarrow}$, and $G^{X \leftarrow}$ for $E^{\{x\} \rightarrow}$, $E^{\{x\} \leftarrow}$, $(V, E^{X \rightarrow})$, and $(V, E^{X \leftarrow})$.

Let n be a positive integer. A *partial n -weighted tournament* is a pair $G = (V, w)$ consisting of a finite set of alternatives V and a weight function $w : V \times V \rightarrow \{0, \dots, n\}$ such that for each pair $(x, y) \in V \times V$ with $x \neq y$, $w(x, y) + w(y, x) \leq n$. We say that $T = (V, w)$ is a (*complete*) *n -weighted tournament* if for all $x, y \in V$ with $x \neq y$, $w(x, y) + w(y, x) = n$. A (partial or complete) *weighted tournament* is a (partial or complete) n -weighted tournament for some $n \in \mathbb{N}$. The class of n -weighted tournaments is denoted by \mathcal{T}_n . Observe that with each partial 1-weighted tournament (V, w) we can associate a partial tournament (V, E) by setting $E = \{(x, y) \in V \times V : w(x, y) = 1\}$. Thus, (partial) n -weighted tournaments can be seen to generalize (partial) tournaments, and we may identify \mathcal{T}_1 with \mathcal{T} .

The notations $G \leq G'$ and $[G]$ can be extended naturally to partial n -weighted tournaments $G = (V, w)$ and $G' = (V', w')$ by letting $(V, w) \leq (V', w')$ if $V = V'$ and $w(x, y) \leq w'(x, y)$ for all $x, y \in V$, and $[G] = \{T \in \mathcal{T}_n : G \leq T\}$.

For given $G = (V, w)$ and $X \subseteq V$, we further define $w^{X \rightarrow}$ such that for all $x, y \in V$,

$$w^{X \rightarrow}(x, y) = \begin{cases} n - w(y, x) & \text{if } x \in X \text{ and } y \notin X, \\ w(x, y) & \text{otherwise,} \end{cases}$$

and set $w^{X \leftarrow} = w^{V \setminus X \rightarrow}$. Moreover, $w^{x \rightarrow}$, $w^{x \leftarrow}$, $G^{X \rightarrow}$, and $G^{X \leftarrow}$ are defined in the obvious way.

We use the term *solution concept* for functions S that associate with each (complete) tournament $T = (V, E)$, or with each (complete) weighted tournament $T = (V, w)$, a choice set $S(T) \subseteq V$. A solution concept S is called *resolute* if $|S(T)| = 1$ for each tournament T .

In this paper we will consider the following solution concepts: *Condorcet winners* (*COND*), *Copeland* (*CO*), *top cycle* (*TC*), and *uncovered set* (*UC*) for tournaments, and *maximin* (*MM*), *Borda* (*BO*), and *ranked pairs* (*RP*) for weighted tournaments. Of these only ranked pairs is resolute. Formal definitions will be provided later in the paper.

3 Possible & Necessary Winners

A solution concept selects alternatives from complete tournaments or complete weighted tournaments. A partial (weighted) tournament, on the other hand, can be extended to a number of complete (weighted) tournaments, and a solution concept selects a (potentially different) set of alternatives for each of them.

For a given solution concept S , we can thus define the set of *possible winners* for a partial (weighted) tournament G as the set of alternatives selected by S from *some* completion of G , i.e., as $PW_S(G) = \bigcup_{T \in [G]} S(T)$. Analogously, the set of *necessary winners* of G is the set of alternatives selected by S from *every* completion of G , i.e., $NW_S(G) = \bigcap_{T \in [G]} S(T)$. We can finally write $PWS_S(G) = \{S(T) : T \in [G]\}$ for the set of *sets* of alternatives that S selects for the different completions of G .

Note that $NW_S(G)$ may be empty even if S selects a non-empty set of alternatives for each tournament $T \in [G]$, and that $|PWS_S(G)|$ may be exponential in the number of alternatives of G . It is also easily verified that $G \leq G'$ implies $PW_S(G') \subseteq PW_S(G)$ and $NW_S(G) \subseteq NW_S(G')$, and that $PW_S(G) = \bigcup_{G \leq G'} NW_S(G')$ and $NW_S(G) = \bigcap_{G \leq G'} PW_S(G')$.

Deciding membership in the sets $PW_S(G)$, $NW_S(G)$, and $PWS_S(G)$ for a given solution concept S and a partial (weighted) tournament G is a natural computational problem. We will respectively refer to these problems as PW_S , NW_S , and PWS_S , and will study them for the solution concepts mentioned at the end of the previous section.³

For complete tournaments T we have $[T] = \{T\}$ and thus $PW_S(T) = NW_S(T) = S(T)$ and $PWS_S(T) = \{S(T)\}$. As a consequence, for solution concepts S with an NP-hard winner determination problem—like *Banks*, *Slater*, and *TEQ*—the problems PW_S , NW_S , and PWS_S are NP-hard as well. We therefore restrict our attention to solution concepts for which winners can be computed in polynomial time.

For irresolute solution concepts, PWS_S may appear a more complex problem than PW_S . We are, however, not aware of a polynomial-time reduction from PW_S to PWS_S . The relationship between these problems may also be of interest for the “classic” possible winner setting with partial preference profiles.

4 Unweighted tournaments

In this section, we consider the following well-known solution concepts for unweighted tournaments: Condorcet winners, Copeland, top cycle, and uncovered set. Weighted tournaments will then be considered in Section 5.

4.1 Condorcet Winners

Condorcet winners are a very simple solution concept and will provide a nice warm-up. An alternative $x \in V$ is a *Condorcet winner* of a complete tournament $T = (V, E)$ if it dominates all other alternatives, i.e., if $(x, y) \in E$ for all $y \in V \setminus \{x\}$. The set of Condorcet winners of tournament T will be denoted by $COND(T)$; obviously this set is always either a singleton or empty.

³Formally, the input for each of the problems consists of an encoding of the partial (n -weighted) tournament G and, for partial n -weighted tournaments, the number n .

It is readily appreciated that the possible Condorcet winners of a partial tournament $G = (V, E)$ are precisely the undominated alternatives, and that a necessary Condorcet winner of G should already dominate all other alternatives. Both properties can be verified in polynomial time.

Each of the sets in $PWS_{COND}(G)$ is either a singleton or the empty set, and determining membership for a singleton is obviously tractable. Checking whether $\emptyset \in PWS_{COND}(G)$ is not quite that simple. First observe that $\emptyset \in PWS_{COND}(G)$ if and only if there is an extension G' of G in which every alternative is dominated by some other alternative. Given a particular $G = (V, E)$, we can define an extension $G' = (V, E')$ of G by iteratively adding edges from dominated alternatives to undominated ones until this is no longer possible. Formally, let

$$E_0 = E \text{ and } E_{i+1} = E_i \cup \{(x, y) \in X_i \times Y_i : (y, x) \notin E_i\},$$

where X_i and Y_i denote the dominated and undominated alternatives of (V, E_i) , respectively. Finally define $E' = \bigcup_{i=0}^{|V|} E_i$, and observe that this set can be computed in polynomial time.

Now, for every undominated alternative x of G' and every dominated alternative y of G' , we not only have $(x, y) \in E'$, but also $(x, y) \in E$. This is the case because in the inductive definition of E' only edges from dominated to undominated alternatives are added in every step. It is therefore easily verified that $PWS_{COND}(G)$ contains \emptyset if and only if the set of undominated alternatives in G' is either empty or is of size three or more. We have shown the following easy result.

Theorem 1. *PW_{COND} , NW_{COND} , and PWS_{COND} can be solved in polynomial time.*

The results for PW_{COND} and NW_{COND} also follow from Proposition 2 of Lang et al. [16] and Corollary 2 of Konczak and Lang [15]. We further note that Theorem 1 is a corollary of corresponding results for maximin in Section 5.2. The reason is that a Condorcet winner is the maximin winner of a 1-weighted tournament, and a tournament does not admit a Condorcet winner if and only if all alternatives are maximin winners.

4.2 Copeland

Copeland's solution selects alternatives based on the number of other alternatives they dominate. Define the *Copeland score* of an alternative x in tournament $T = (V, E)$ as $s_{CO}(x, T) = |D_T^+(x)|$. The set $CO(T)$ then consists of all alternatives that have maximal Copeland score. Since Copeland scores coincide with Borda scores in the case of 1-weighted tournaments, the following is a direct corollary of the results in Section 5.1.

Theorem 2. *NW_{CO} , PW_{CO} , and PWS_{CO} can be solved in polynomial time.*

PW_{CO} can alternatively be solved via a polynomial-time reduction to maximum network flow (see, e.g., [8], p. 51).

4.3 Top Cycle

A subset $X \subseteq V$ of alternatives in a (partial or complete) tournament (V, E) is *dominant* if every alternative in X dominates every alternative outside X . The *top cycle* of a tournament $T = (V, E)$, denoted by $TC(T)$, is the unique *minimal* dominant subset of V .

Lang et al. have shown that possible and necessary winners for TC can be computed efficiently by greedy algorithms ([16], Corollaries 1 and 2). For PWS_{TC} , we not only have to check that there exists a completion such that the set in question is dominating, but also that there is no smaller dominating set. It turns out that this can still be done in polynomial time.

Theorem 3. PWS_{TC} can be solved in polynomial time.

Proof sketch. Consider a partial tournament $G = (V, E)$ and a set $X \subseteq V$ of alternatives. If X is a singleton, the problem reduces to checking whether $X \in PWS_{COND}(G)$. If X is of size two or if one of its elements is dominated by an outside alternative, $X \notin PWS_{TC}(G)$. Therefore, we can without loss of generality assume that $|X| \geq 3$ and $(y, x) \notin E$ for all $y \in V \setminus X$ and $x \in X$. The *Smith set* of a partial tournament is defined as the minimal dominant subset of alternatives [22].⁴ It can be shown that there exists a completion $T \in [G]$ with $TC(T) = X$ if and only if the Smith set of the partial tournament $(X, E|_{X \times X})$ equals the whole set X . Since Brandt et al. [4] have shown that the Smith set of a partial tournament can be computed efficiently, the theorem follows. \square

4.4 Uncovered Set

Given a tournament $T = (V, E)$, an alternative $x \in V$ is said to *cover* another alternative $y \in V$ if $D_T^+(y) \subseteq D_T^+(x)$, i.e., if every alternative dominated by y is also dominated by x . The *uncovered set* of T , denoted by $UC(T)$, then is the set of alternatives that are not covered by some other alternative. A useful alternative characterization of the uncovered set is via the *two-step principle*: an alternative is in the uncovered set if and only if it can reach every other alternative in at most two steps.⁵ Formally, $x \in UC(T)$ if and only if for all $y \in V \setminus \{x\}$, either $(x, y) \in E$ or there is some $z \in V$ with $(x, z), (z, y) \in E$. We denote the two-step dominion $D_E^+(D_E^+(x))$ of an alternative x by $D_E^{++}(x)$.

We first consider PW_{UC} , for which we check for each alternative whether it can be reinforced to reach every other alternative in at most two steps.

Theorem 4. PW_{UC} can be solved in polynomial time.

Proof. For a given partial tournament $G = (V, E)$ and an alternative $x \in V$, we check whether x is in $UC(T)$ for some completion $T \in [G]$.

Consider the graph $G' = (V, E')$ where E' is derived from E as follows. First, we let $D^+(x)$ grow as much as possible by letting $E' = E^{x \rightarrow}$. Then, we do the same for its two-step dominion by defining E'' as $E'^{D_{E'}^+(x) \rightarrow}$. Now it can be shown that $x \in PW_{UC}(G)$ if and only if $V = \{x\} \cup D_{E''}^+(x) \cup D_{E''}^{++}(x)$. \square

A similar argument yields the following.

Theorem 5. NW_{UC} can be solved in polynomial time.

Proof. For a given partial tournament $G = (V, E)$ and an alternative $x \in V$, we check whether x is in $UC(T)$ for all completions $T \in [G]$.

Consider the graph $G' = (V, E')$ with E' defined as follows. First, let $E' = E^{x \leftarrow}$. Then, expand it to $E'' = E'^{D_{E'}^-(x) \rightarrow}$. Intuitively, this makes it as hard as possible for x to beat alternatives outside of its dominion in two steps. Then it can be shown that $x \in NW_{UC}(G)$ if and only if $V = \{x\} \cup D_{E''}^+(x) \cup D_{E''}^{++}(x)$. \square

For all solution concepts considered so far—Condorcet winners, Copeland, and top cycle— PW and PWS have the same complexity. One might wonder whether a result like this holds more generally, and whether there could be a polynomial-time reduction from PWS to PW . The following result shows that this is not the case, unless $P=NP$.

Theorem 6. PWS_{UC} is NP-complete.

This can be shown by a rather intricate reduction from SAT. We have to omit the construction due to space constraints but a sketch is presented in the appendix.

⁴For complete tournaments, the Smith set coincides with the top cycle.

⁵In graph theory, vertices satisfying this property are often called *kings*.

5 Weighted Tournaments

We now turn to weighted tournaments, and in particular consider the solution concepts Borda, maximin, and ranked pairs.

5.1 Borda

The Borda solution (BO) is typically used in a voting context, where it is construed as based on voters' rankings of the alternatives: each alternative receives $|V| - 1$ points for each time it is ranked first, $|V| - 2$ points for each time it is ranked second, and so forth; the solution concept then chooses the alternatives with the highest total number of points. In the more general setting of weighted tournaments, the *Borda score* of alternative $x \in V$ in $G = (V, w)$ is defined as $s_{BO}(x, G) = \sum_{y \in V \setminus \{x\}} w(x, y)$ and the *Borda winners* are the alternatives with the highest Borda score. If $w(x, y)$ represents the number of voters that rank x higher than y , the two definitions are equivalent.

Before we proceed further, we define the notion of a b -matching, which will be used in the proofs of two of our results. Let $H = (V_H, E_H)$ be an undirected graph with vertex capacities $b : V_H \rightarrow \mathbb{N}_0$. Then, a b -matching of H is a function $m : E_H \rightarrow \mathbb{N}_0$ such that for all $v \in V_H$, $\sum_{e \in \{e' \in E_H : v \in e'\}} m(e) \leq b(v)$. The *size* of b -matching m is defined as $\sum_{e \in E_H} m(e)$. It is easy to see that if $b(v) = 1$ for all $v \in V_H$, then a maximum size b -matching is equivalent to a maximum cardinality matching. In a b -matching problem with upper *and* lower bounds, there further is a function $a : V_H \rightarrow \mathbb{N}_0$. A feasible b -matching then is a function $m : E_H \rightarrow \mathbb{N}_0$ such that $a(v) \leq \sum_{e \in \{e' \in E_H : v \in e'\}} m(e) \leq b(v)$.

If H is bipartite, then the problem of computing a maximum size feasible b -matching with lower and upper bounds can be solved in strongly polynomial time ([21], Chapter 21). We will use this fact to show that PW_{BO} and PWS_{BO} can both be solved in polynomial time. While the following result for PW_{BO} can be shown using Theorem 6.1 of [14], we give a direct proof that can then be extended to PWS_{BO} .

Theorem 7. *PW_{BO} can be solved in polynomial time.*

Proof sketch. Let $G = (V, w)$ be a partial n -weighted tournament, $x \in V$. We give a polynomial-time algorithm for checking whether $x \in PW_{BO}(G)$, via a reduction to the problem of computing a maximum size b -matching of a bipartite graph.

Let $G^{x \rightarrow} = (V, w^{x \rightarrow})$ denote the graph obtained from G by maximally reinforcing x , and $s^* = s_{BO}(x, G^{x \rightarrow})$ the Borda score of x in $G^{x \rightarrow}$. From $G^{x \rightarrow}$, we then construct a bipartite graph $H = (V_H, E_H)$ with vertices $V_H = V \setminus \{x\} \cup E^{<n}$, where $E^{<n} = \{\{i, j\} \subseteq V \setminus \{x\} : w(i, j) + w(j, i) < n\}$,⁶ and edges $E_H = \{\{v, e\} : v \in V \setminus \{x\} \text{ and } v \in e \in E^{<n}\}$. We further define vertex capacities $b : V_H \rightarrow \mathbb{N}_0$ such that $b(\{i, j\}) = n - w(i, j) - w(j, i)$ for $\{i, j\} \in E^{<n}$ and $b(v) = s^* - s_{BO}(v, G^{x \rightarrow})$ for $v \in V \setminus \{x\}$.

Now observe that in any completion $T = (V, w') \in [G^{x \rightarrow}]$, $w'(i, j) + w'(j, i) = n$ for all $i, j \in V$ with $i \neq j$. The sum of the Borda scores in T is therefore $n|V|(|V| - 1)/2$. Some of the weight has already been used up in $G^{x \rightarrow}$; the weight which has not yet been used up is equal to $\alpha = n|V|(|V| - 1)/2 - \sum_{v \in V} s_{BO}(v, G^{x \rightarrow})$. We claim that $x \in PW_{BO}(G)$ if and only if H has a b -matching of size at least α . □

This idea can be extended to a polynomial-time algorithm for PWS_{BO} where we use a similar construction for a given $G = (V, w)$, a candidate set $X \subset V$ and a target Borda score s^* . Binary search can be used to efficiently search the interval $I = [\max_{x \in X} s_{BO}(x, G), n(|V| - 1)]$ of possible target scores. The full proof is omitted.

⁶Note that $w(i, j) = w^{x \rightarrow}(i, j)$ for alternatives $i, j \in V \setminus \{x\}$.

Theorem 8. PWS_{BO} can be solved in polynomial time.

We conclude this section by showing that NW_{BO} can be solved in polynomial time as well.

Theorem 9. NW_{BO} can be solved in polynomial time.

Proof. Let $G = (V, w)$ be a partial weighted tournament, $x \in V$. We give a polynomial-time algorithm for checking whether $x \in NW_{BO}(G)$.

Let $G' = G^{x\leftarrow}$. We want to check whether some other alternative $y \in V \setminus \{x\}$ can achieve a Borda score of more than $s^* = s_{BO}(x, G')$. This can be done separately for each $y \in V \setminus \{x\}$ by reinforcing it as much as possible in G' . If for some y , $s_{BO}(y, G'^{y\rightarrow}) > s^*$, then $x \notin NW_{BO}(G)$. If, on the other hand, $s_{BO}(y, G'^{y\rightarrow}) \leq s^*$ for all $y \in V \setminus \{x\}$, then $x \in NW_{BO}(G)$. \square

Since the Borda and Copeland solutions coincide in unweighted tournaments, the above results imply that PW_{CO} and NW_{CO} can be solved in polynomial time. The same is true for PWS_{CO} , because the Copeland score is bounded by $|V| - 1$.

5.2 Maximin

The *maximin score* $s_{MM}(x, T)$ of an alternative x in a weighted tournament $T = (V, w)$, is given by its worst pairwise comparison, i.e., $s_{MM}(x, T) = \min_{y \in V \setminus \{x\}} w(x, y)$. The *maximin solution*, also known as *Simpson's method* and denoted by MM , returns the set of all alternatives with the highest maximin score.

We first show that PW_{MM} is polynomial-time solvable by reducing it to the problem of finding a maximum cardinality matching of a graph.

Theorem 10. PW_{MM} can be solved in polynomial time.

Proof sketch. We show how to check whether $x \in PW_{MM}(G)$ for a partial n -weighted tournament $G = (V, w)$. Consider the graph $G^{x\rightarrow} = (V, w^{x\rightarrow})$. Then, $s_{MM}(x, G^{x\rightarrow})$ is the best possible maximin score x can get among all completions of G . If $s_{MM}(x, G^{x\rightarrow}) \geq \frac{n}{2}$, then we have $s_{MM}(y, T) \leq w^{x\rightarrow}(y, x) \leq \frac{n}{2}$ for every $y \in V \setminus \{x\}$ and every completion $T \in [G^{x\rightarrow}]$ and therefore $x \in PW_{MM}(G)$. Now consider $s_{MM}(x, G^{x\rightarrow}) < \frac{n}{2}$. We will reduce the problem of checking whether $x \in PW_{MM}(G)$ to that of finding a maximum cardinality matching, which is known to be solvable in polynomial time [11]. We want to find a completion $T \in [G^{x\rightarrow}]$ such that $s_{MM}(x, T) \geq s_{MM}(y, T)$ for all $y \in V \setminus \{x\}$. If there exists a $y \in V \setminus \{x\}$ such that $s_{MM}(x, G^{x\rightarrow}) < s_{MM}(y, G^{x\rightarrow})$, then we already know that $x \notin PW_{MM}(G)$. Otherwise, each $y \in V \setminus \{x\}$ derives its maximin score from at least one particular edge (y, z) where $z \in V \setminus \{x, y\}$ and $w(y, z) \leq s_{MM}(x, G^{x\rightarrow})$. Moreover, it is clear that in any completion, y and z cannot both achieve a maximin score of less than $s_{MM}(x, G^{x\rightarrow})$ from edges (y, z) and (z, y) at the same time.

Construct the following undirected and unweighted graph $H = (V_H, E_H)$ where $V_H = V \setminus \{x\} \cup \{\{i, j\} \subseteq V : i \neq j\}$. Build up E_H such that: $\{i, \{i, j\}\} \in E_H$ if and only if $i \neq j$ and $w^{x\rightarrow}(i, j) \leq s_{MM}(x, G^{x\rightarrow})$. In this way, if i is matched to $\{i, j\}$ in H , then i derives a maximin score of less than or equal to $s_{MM}(x, G^{x\rightarrow})$ from his comparison with j . Clearly, H is polynomial in the size of G . Then, the claim is that $x \in PW_{MM}(G)$ if and only if there exists a matching of cardinality $|V| - 1$ in H . \square

For NW_{MM} we apply a similar technique as for NW_{BO} : to see whether $x \in NW_{MM}(G)$, we start from the graph $G^{x\leftarrow}$ and check whether some other alternative can achieve a higher maximin score than x in a completion of $G^{x\leftarrow}$.

Theorem 11. NW_{MM} can be solved in polynomial time.

We conclude the section by showing that PWS_{MM} can be solved in polynomial time. The proof proceeds by identifying the maximin values that could potentially be achieved simultaneously by all elements of the set in question, and solving the problem for each of these values using similar techniques as in the proof of Theorem 10. Only a polynomially bounded number of problems need to be considered.

Theorem 12. PWS_{MM} can be solved in polynomial time.

5.3 Ranked Pairs

The method of *ranked pairs* (RP) is the only resolute solution concept considered in this paper. Given a weighted tournament $T = (V, w)$, it returns the unique undominated alternative of a transitive tournament T' on V constructed in the following manner. First order the (directed) edges of T in decreasing order of weight, breaking ties according to some exogenously given tie-breaking rule. Then consider the edges one by one according to this ordering. If the current edge can be added to T' without creating a cycle, then do so; otherwise discard the edge.⁷

It is readily appreciated that this procedure, and thus the winner determination problem for RP , is computationally tractable. The possible winner problem, on the other hand, turns out to be NP-hard. This also shows that tractability of the winner determination problem, while necessary for tractability of PW , is not generally sufficient.

Theorem 13. PW_{RP} is NP-complete.

Proof sketch. Membership in NP is obvious, as for a given completion and a given tie-breaking rule, the ranked pairs winner can be found efficiently.

NP-hardness can be shown by a reduction from SAT. For a Boolean formula φ in conjunctive normal-form with a set C of clauses and set P of propositional variables, we construct a partial 8-weighted tournament $G_\varphi = (V_\varphi, w_\varphi)$ as follows. For each variable $p \in P$, V_φ contains two *literal alternatives* p and \bar{p} and two *auxiliary alternatives* p' and \bar{p}' . For each clause $c \in C$, there is an alternative c . Finally, there is an alternative d for which membership in $PW_{RP}(G_\varphi)$ is to be decided.

In order to conveniently describe the weight function w_φ , let us introduce the following terminology. For two alternatives $x, y \in V_\varphi$, say that there is a *heavy* edge from x to y if $w_\varphi(x, y) = 8$ (and therefore $w_\varphi(y, x) = 0$). A *medium* edge from x to y means $w_\varphi(x, y) = 6$ and $w_\varphi(y, x) = 2$, and a *light* edge from x to y means $w_\varphi(x, y) = 5$ and $w_\varphi(y, x) = 3$. Finally, a *partial* edge between x and y means $w_\varphi(x, y) = w_\varphi(y, x) = 1$.

We are now ready to define w_φ . For each variable $p \in P$, we have heavy edges from p to \bar{p}' and from \bar{p} to p' , and partial edges between p and p' and between \bar{p} and \bar{p}' . For each clause $c \in C$, we have a medium edge from c to d and a heavy edge from the literal alternative $\ell_i \in \{p, \bar{p}\}$ to c if the corresponding literal ℓ_i appears in the clause c . Finally, we have heavy edges from d to all auxiliary alternatives and light edges from d to all literal alternatives. For all pairs x, y for which no edge has been specified, we define $w_\varphi(x, y) = w_\varphi(y, x) = 4$.

Observe that the only pairs of alternatives for which w_φ is not fully specified are those pairs that are connected by a partial edge. It can be shown that alternative d is a possible

⁷The variant of ranked pairs originally proposed by Tideman [23], which was also used by Xia and Conitzer [25], instead chooses a set of alternatives, containing any alternative that is selected by the above procedure for *some* way of breaking ties among edges with equal weight. We do not consider this irresolute version of ranked pairs because it was recently shown that winner determination for this variant is NP-hard [5]. As mentioned in Section 3, this immediately implies that all problems concerning possible or necessary winners are NP-hard as well.

ranked pairs winner in G_φ if and only if φ is satisfiable. Intuitively, choosing a completion w' of w_φ such that $w'(p', p)$ is large and $w'(\bar{p}', \bar{p})$ is small corresponds to setting the variable p to “true.” \square

Since the ranked pairs method is resolute, hardness of PWS_{RP} follows immediately.

Corollary 1. *PWS_{RP} is NP-complete.*

Computing necessary ranked pairs winners turns out to be coNP-complete. This is again somewhat surprising, as computing necessary winners is often considerably easier than computing possible winners, both for partial tournaments and partial preference profiles [25].

Theorem 14. *NW_{RP} is coNP-complete.*

Proof sketch. Membership in coNP is again obvious. For hardness, we give a reduction from UNSAT that is a slight variation of the reduction in the proof of Theorem 13. We introduce a new alternative d^* , which has heavy edges to all alternatives in V_φ except d . Furthermore, there is a light edge from d to d^* . It can be shown that d^* is a necessary ranked pairs winner in this partial 8-weighted tournament if and only if φ is unsatisfiable. \square

6 Discussion

The problem of computing possible and necessary winners for partial preference profiles has recently received a lot of attention. In this paper, we have investigated this problem in a setting where partially specified (weighted or unweighted) *tournaments* instead of profiles are given as input. We have summarized our findings in Table 1.

A key conclusion is that computational problems for partial tournaments can be significantly easier than their counterparts for partial profiles. For example, possible Borda or maximin winners can be found efficiently for partial tournaments, whereas the corresponding problems for partial profiles are NP-complete [25].

While tractability of the winner determination problem is necessary for tractability of the possible or necessary winners problems, the results for ranked pairs in Section 5.3 show that it is not sufficient. We further considered the problem of deciding whether a given subset of alternatives equals the winner set for some completion of the partial tournament. The results for the uncovered set in Section 4.4 imply that this problem cannot be reduced to the computation of possible or necessary winners, but whether a reduction exists in the opposite direction remains an open problem.

Partial tournaments have also been studied in their own right, independent of their possible completions. For instance, Peris and Subiza [18] and Dutta and Laslier [10] have generalized several tournament solutions to incomplete tournaments by directly adapting their definitions. In this context, the notion of possible winners suggests a canonical way to generalize a tournament solution to incomplete tournaments. The positive computational results in this paper are an indication that this may be a promising approach.

Furthermore, we have not examined the complexity of computing possible and necessary winners for some attractive tournament solutions such as the minimal covering set, the bipartisan set [17] and weighted versions of the top cycle and the uncovered set [9].

An interesting related question that goes beyond the computation of possible and necessary winners is the following: when the winners are not yet fully determined, which unknown comparisons need to be learned, or which matches should be played? The construction of a policy tree defining an optimal protocol minimizing the number of questions to be asked or the number of matches to be played, in the worst case or on average, is an even more challenging issue that we leave for further research.

Acknowledgments We thank Gerhard Woeginger for hints towards improving our previous pseudo-polynomial time algorithms for PWS_{BO} and PWS_{MM} . This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/9-1, BR 2312/10-1, and FI 1664/1-1. Jérôme Lang has been supported by the ANR project ComSoc (ANR-09-BLAN-0305).

References

1. D. Baumeister and J. Rothe. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. *Information Processing Letters*, 112(5):186–190, 2012.
2. N. Betzler and B. Dorn. Towards a dichotomy for the possible winner problem in elections based on scoring rules. *Journal of Computer and System Sciences*, 76(8):812–836, 2010.
3. N. Betzler, S. Hemmann, and R. Niedermeier. A multivariate complexity analysis of determining possible winners given incomplete votes. In *Proc. of 21st IJCAI*, pages 53–58, 2009.
4. F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. *Mathematical Logic Quarterly*, 55(4):444–459, 2009.
5. M. Brill and F. Fischer. The price of neutrality for the ranked pairs method. In J. Hoffmann and B. Selman, editors, *Proc. of 26th AAAI Conference*. AAAI Press, 2012. Forthcoming.
6. I. Charon and O. Hudry. A survey on the linear ordering problem for weighted or unweighted tournaments. *4OR*, 5(1):5–60, 2007.
7. Y. Chevaleyre, J. Lang, N. Maudet, J. Monnot, and L. Xia. New candidates welcome! Possible winners with respect to the addition of new candidates. *Mathematical Social Sciences*, 64(1):74–88, 2012.
8. W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization*. Wiley and Sons, 1998.
9. P. De Donder, M. Le Breton, and M. Truchon. Choosing from a weighted tournament. *Mathematical Social Sciences*, 40(1):85–109, 2000.
10. B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
11. J. Edmonds. Paths, trees and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
12. N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.
13. M. Kalech, S. Kraus, G. A. Kaminka, and C. V. Goldman. Practical voting rules with partial information. *Journal of Autonomous Agents and Multi-Agent Systems*, 22(1):151–182, 2011.
14. W. Kern and D. Paulusma. The computational complexity of the elimination problem in generalized sports competitions. *Discrete Optimization*, 1(2):205–214, 2004.

15. K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proc. of Multidisciplinary Workshop on Advances in Preference Handling*, pages 124–129. 2005.
16. J. Lang, M. S. Pini, F. Rossi, D. Salvagnin, K. B. Venable, and T. Walsh. Winner determination in voting trees with incomplete preferences and weighted votes. *Journal of Autonomous Agents and Multi-Agent Systems*, 25(1):130–157, 2012.
17. J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer, 1997.
18. J. E. Peris and B. Subiza. Condorcet choice correspondences for weak tournaments. *Social Choice and Welfare*, 16(2):217–231, 1999.
19. M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Dealing with incomplete agents' preferences and an uncertain agenda in group decision making via sequential majority voting. In *Proc. of 11th KR Conference*, pages 571–578. AAAI Press, 2008.
20. M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Possible and necessary winners in voting trees: Majority graphs vs. profiles. In *Proc. of 10th AAMAS Conference*, pages 311–318, 2011.
21. A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.
22. J. H. Smith. Aggregation of preferences with variable electorate. *Econometrica*, 41(6): 1027–1041, 1973.
23. T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
24. T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proc. of 22nd AAAI Conference*, pages 3–8. AAAI Press, 2007.
25. L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.
26. L. Xia, J. Lang, and J. Monnot. Possible winners when new alternatives join: New results coming up! In *Proc. of 10th AAMAS Conference*, pages 829–836, 2011.

Haris Aziz, Markus Brill, Paul Harrenstein, Hans Georg Seedig
Institut für Informatik
Technische Universität München
Garching bei München, Germany
Email: {aziz,brill,harrenst,seedigh}@in.tum.de

Felix Fischer
Statistical Laboratory
University of Cambridge
Cambridge, UK
Email: fischerf@statslab.cam.ac.uk

Jérôme Lang
LAMSADE
Université Paris-Dauphine
Paris, France
Email: lang@lamsade.dauphine.fr