Winner Determination in Voting Trees with Incomplete Preferences and Weighted Votes*

Jérôme Lang
LAMSADE

Maria Silvia Pini
University of Padova

Francesca Rossi
University of Padova

Domenico Salvagnin
University of Padova

Kristen Brent Venable
University of Padova

Toby Walsh
NICTA and UNSW

Abstract

In multiagent settings where agents have different preferences, preference aggregation can be an important issue. Voting is a general method to aggregate preferences. We consider the use of voting tree rules to aggregate agents' preferences. In a voting tree, decisions are taken by performing a sequence of pairwise comparisons in a binary tree where each comparison is a majority vote among the agents. Incompleteness in the agents' preferences is common in many real-life settings due to privacy issues or an ongoing elicitation process. We study how to determine the winners when preferences may be incomplete, not only for voting tree rules (where the tree is assumed to be fixed), but also for the Schwartz rule (in which the winners are the candidates winning for at least one voting tree). In addition, we study how to determine the winners when only balanced trees are allowed. In each setting, we address the complexity of computing necessary (respectively, possible) winners, which are those candidates winning for all completions (respectively, at least one completion) of the incomplete profile. We show that many such winner determination problems are computationally intractable when the votes are weighted. However, in some cases, the exact complexity remains unknown. Since it is generally computationally difficult to find the exact set of winners for voting trees and the Schwartz rule, we propose several heuristics that find in polynomial time a superset of the possible winners and a subset of the necessary winners which are based on the completions of the (incomplete) majority graph built from the incomplete profiles.

1 Introduction

In multiagent settings, agents generally have different preferences, and it can be important to aggregate these preferences, i.e., to select a collectively desirable candidate

*This paper is a revised and extended version of the conference papers [23, 28].

from a set of candidates. Candidates could be, for example, potential presidents, joint plans, allocations of goods or resources, etc. A general method for aggregating preferences in multi-agent systems, in order to take a collective decision, is to run an election among the different options using a voting rule.

Eliciting preferences from agents, to be able to run such an election, is a typical problem in multiagent systems, such as combinatorial auctions and voting systems. However, preference elicitation is often difficult, time-consuming and costly. Agents may be unwilling to reveal all their preferences due to privacy reasons, or due to the set of alternatives being large. Fortunately, we can often determine the outcome before all the preferences have been revealed [10]. For example, it may be that one option has so many votes that it will win whatever happens with the remaining votes.

In this paper, we consider the class of voting tree rules (for short, we shall often say *voting trees*, identifying the tree with the associated rule). These are sometimes called "sequential majority voting" or "Cup." In addition, we consider the Schwartz rule, which can be characterized as a voting tree rule with an undefined tree since the Schwartz winners are exactly those who win for at least one voting tree. We also consider fair winners, i.e., those candidates who win when the voting tree is *balanced*. Fairness comes from the fact that both finalists will have faced the same number of competitions, or the same number plus or minus one.

We consider the computation of the above rules when voters' preferences are incomplete, i.e., represented by partial orders on the set of candidates. We focus on the problem of determining if a given candidate is the *necessary winner* (i.e., the candidate who wins for any completion of these partial preferences) and a *possible winner* (i.e., a candidate who wins for some completion of these partial preferences). Determining if a given candidate is the necessary or a possible winner for voting trees and the Schwartz rule is likely to be computationally hard. We show that they are indeed NP-complete when the input consists of weighted votes, and also when we require that the tree is balanced. We therefore introduce heuristics, that find a superset of the possible winners and a subset of the necessary winners, where the set of completions of the incomplete preferences is replaced by the completions of the *incomplete majority graph* computed from the incomplete preferences. We discuss the links between possible (respectively, necessary) winners from incomplete profiles and those returned by the heuristics, and we show that these heuristics can be computed in polynomial time.

Our results concern several issues in computational social choice: voting trees, voting under incomplete knowledge, and incomplete tournaments. Voting trees have received increasing attention of late in the literature. For example, they have been considered by Trick [34], Procaccia *et al.* [30, 31], Conitzer *et al.* [9, 11], Xia and Conitzer [38], Hazon *et al.* [20], Vu et al. [35], Fischer *et al.* [16, 17], and Vassilevska Williams [37]. The computational aspects of incomplete tournaments have been thoroughly investigated by Brandt *et al.* [6, 7]. Voting under incomplete knowledge has received even more attention [10, 21, 27, 38, 36, 19, 4, 3, 5, 2, 8]. We will discuss some of these papers in Section 7.

2 Voting trees, the Schwartz rule, and fair winners

We define some basic notions related to voting trees. We first focus on unweighted profiles and then extend all the notions to weighted votes.

2.1 Preferences, profiles and majority graphs

We assume that each agent's *preferences* are specified by a strict total order (TO), that is, by an asymmetric, transitive and complete order, on a set of m candidates. The candidates are taken from a set Ω, and they represent the possible options over which agents vote.

Definition 1 (profile) *A* profile P *on* Ω *is a collection of* n *strict total orders over* Ω*, i.e.,* $P = (P_1, \ldots, P_n)$*, where* P_i *is the preference relation (or* vote*) of agent (also called* voter*)* i*.*

For the sake of simplicity we assume throughout the paper that the number n of voters is odd. Our results can be extended to the case where n is even, but at the price of some complications that arise with the handling of ties.

Profiles are denoted using the following notation: (($A > B > C$); ($A > C > B$); ($C > A > B$)) means that voter 1 prefers A to B and B to C, etc.

Definition 2 (voting rule and correspondence) *A (deterministic) voting rule is a mapping from the set of profiles to the set of candidates. A voting correspondence is a mapping from the set of profiles to the set of nonempty sets of candidates.*

A voting rule computes a single winner. On the other hand, a voting correspondence can compute multiple winners. For the sake of simplicity, we will use the word "rule" even for correspondence where it does not cause any ambiguity.

Given a profile P, the induced majority graph $M(P)$ is defined as follows.

Definition 3 (majority graph) *Let* P *be a profile. The majority graph induced by* P*, denoted by* $M(P)$*, is the graph whose set of vertices is the set of candidates* Ω *and where, for all* $A, B \in \Omega$*, there is a directed edge from* A *to* B *(denoted by* $A >_m^P B$*, or by the abbreviated form* $A >_m B$*) if and only if a strict majority of voters prefers* A *to* B*.*

The majority graph $M(P)$ induced from any profile P is asymmetric, but it is not necessarily transitive. Moreover, since the number of voters is odd, $M(P)$ is complete: for each A and $B \neq A$, either $A >_m B$ or $B >_m A$ holds. Therefore, $M(P)$ is a complete and asymmetric graph, also called a *tournament* on Ω [24].

2.2 Voting trees

Given a set of candidates, a voting tree is a binary tree (also called an *agenda*) with one candidate per leaf, such that each candidate appears on at least one leaf. This assumption can be made without loss of generality: if a candidate does not appear

anywhere in the tree then they can be removed from the set of candidates. In addition if every candidate appears on *exactly* one leaf then we say that the voting tree is a *simple voting tree*. Each internal node represents the candidate that wins the pairwise election between the node's children. The winner of each pairwise election is computed by the majority rule, where A beats B if and only if there is a majority of voters stating $A > B$. The candidate at the root of the agenda is the overall winner. Given a binary tree T, we will denote by $V(T)$ (and sometimes with $r_{V(T)}$) the voting tree rule induced by T, and given a profile P, we will denote by $V(P, T)$ the candidate that, given profile P, wins in tree T. Note that, because the number of voters is odd, ties never happen in the pairwise elections, so that $V(T)$ is a (deterministic) voting rule. We denote by V the class of all voting tree rules.

Example 1 *Let $\Omega = \{A, B, C, D\}$. Consider the agenda T shown in Figure 1(a), where A first plays against D, the winner, called w_1, plays against C, and then the winner, called w_2, plays against B. The winner, called w_3, is the overall winner. Consider the profile $P = ((A > B > C > D); (D > A > B > C); (C > D > B > A))$. Its induced majority graph is shown in Figure 1(b). Therefore $V(P, T) = B$.* □

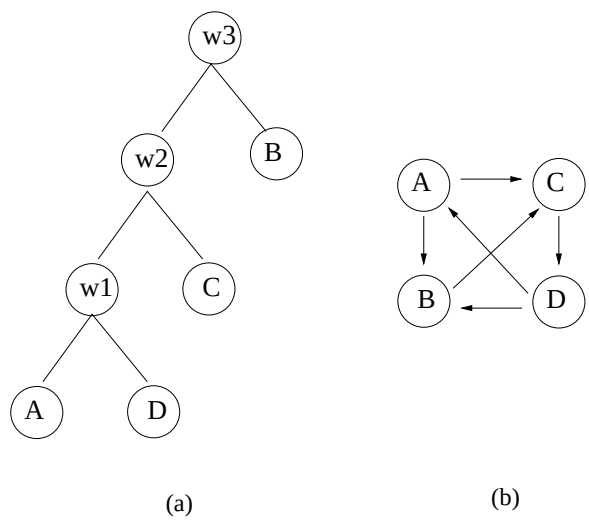

Figure 1: An agenda and a majority graph.

2.3 The Schwartz rule

Given a profile P, a Schwartz winner [32] is a candidate that wins, according to the preferences of P, in some agenda.

Definition 4 (Schwartz winner) *Given a profile P, a candidate A is a Schwartz winner if and only if there exists a binary tree T such that $V(P, T) = A$.*

The Schwartz rule (more precisely, correspondence) denoted by $r_{Schwartz}$ maps any profile P to the set of all Schwartz winners for P.

Example 2 *Consider the profile presented in Example 1, i.e.,* $P = ((A > B > C > D); (D > A > B > C); (C > D > B > A))$. *Then the set of Schwartz winners for* P *is* $\{A, B, C, D\}$. □

The set of Schwartz winners coincides with the *top cycle* of the majority graph [26]. The top cycle of a majority graph G is the set of maximal elements of the reflexive and transitive closure G^* of G. An equivalent characterization of Schwartz winners, that we use later in this paper, is the following:

Theorem 1 (see e.g. [26, 24]) *Given a majority graph* G*, a candidate* A *is a Schwartz winner if and only if, for every other candidate* C*, there exists a path from* A *to* C.

From the characterization above, since finding a path from a single source to all other candidates is polynomial [13], computing the set of the Schwartz winners is polynomial.

Note that voting tree rules, as well as Schwartz, are rules that can be computed from the majority graph. Among the most well-known rules computable from the majority graph, the Schwartz rule is the most liberal. At the other extreme, the Condorcet set contains the Condorcet winner when there exists one, and is empty otherwise; we recall that the Condorcet winner is a candidate who beats every other candidate in the majority graph. In particular, the Schwartz set contains the set of Copeland winners[1] [22], as well as the uncovered set[2], the Banks set [1], and the Slater set[3] [33]. Finally, all these rules are Condorcet-consistent: when a Condorcet winner A exists, all these sets coincide with the singleton $\{A\}$.

2.4 Fair winners

Schwartz winners are candidates who win in at least one binary tree. However, such a tree could be very unbalanced and the winner could compete against only a few candidates. This might be considered "unfair". In the following, we will consider a voting rule *fair* if it has a balanced agenda.

Definition 5 (balanced agenda) *A balanced agenda* T *is a binary tree in which each candidate must appear exactly once in leaves and the difference between the maximum and the minimum depth among the leaves (the level of imbalance, written* $d(T)$*) is less than or equal to 1.*

Assume that we have a profile of strict total orders. We will say that a candidate is a fair winner if he wins under some balanced agenda.

[1] Copeland winners are those candidates who maximize the number of outgoing edges in the majority graph [12].

[2] The uncovered set are those candidates that defeat every other alternative either directly or indirectly at one remove [25].

[3] We do not give formal definitions for Banks and Slater sets since they do not play any further role in the paper.

Definition 6 (fair winner) *Given a profile P, a candidate A is a fair winner if and only if there exists a balanced agenda T, such that $V(P,T) = A$.*

Example 3 *Consider the profile presented in Example 1, i.e., $P = ((A > B > C > D); (D > A > B > C); (C > D > B > A))$. The set of the fair winners for P is $\{A, D\}$. In fact, we have only three balanced trees, up to isomorphism: in the balanced tree where A plays against B, and C plays against D, the winner is A, while in the other two balanced trees the winner is D.* □

Notice that the Condorcet winner is a fair winner since they win in all trees, and thus also in balanced ones. Moreover, a fair winner is a Schwartz winner since there is a tree, in fact a balanced tree in which they win. Moreover, while there could be no Condorcet winner, there is always at least one fair winner, because there always exist at least one balanced tree. Notice also that the candidates that win in every balanced tree are Condorcet winners.

2.5 Weighted votes

Until now we have considered unweighted votes. However, weighted voting systems are used in a number of real-world settings like shareholder meetings and elected assemblies. An agent with integer weight k can be viewed as k agents voting identically. Although human elections are often unweighted, the addition of weights make voting schemes more general. Weights are useful in multiagent systems where we have different types of agents and some agents are more important than others when taking a decision. Moreover, the weighted case informs us about unweighted case when we have probabilistic information about votes [9]. For example, in [9] Conitzer and Sandholm show that, if coalitional manipulation is hard in the weighted case with complete information about the votes, evaluating a candidate's winning probability is hard in the unweighted case when there is uncertainty about the votes provided we permit perfect correlations between the votes (that is, some votes are uncertain but known to be identical).

We now generalize the notions presented above in this section assuming that each agent may have a weight.

Definition 7 (weighted profile) *A weighted profile is a profile where each agent has a given integer weight.*

Note that the unweighted case considered previously is a special case of weighted case where all the agents' weights are equal.

For simplicity, we assume that the sum of the weights of the agents is odd.

Definition 8 (corresponding unweighted profile) *Given a weighted profile P, its corresponding unweighted profile $U(P)$ is the profile obtained from P by replacing every ordering, say O, expressed by an agent with weight k, by k agents with weight 1 all with the same ordering O.*

Example 4 *Consider the profile shown in Example 1, i.e.,* $P = ((A > B > C > D); (D > A > B > C); (C > D > B > A))$. *Assume that the first agent has weight* 1*, and that the second and third agents have weight* 2*. Then the corresponding unweighted profile is* $U(P) = ((A > B > C > D); (D > A > B > C); (D > A > B > C); (C > D > B > A); (C > D > B > A))$. □

Given a weighted profile P, the induced *majority graph* $M(P)$ is the majority graph of the corresponding unweighted profile of P, i.e., $M(P) = M(U(P))$.

Since the sum of weights is odd, $M(P)$ is complete: for each A and $B \neq A$, either $A >_m B$ or $B >_m A$ holds. Therefore, as for unweighted profiles, $M(P)$ is a tournament.

Example 5 *Consider the weighted profile of Example 4. The majority graph induced by such a profile is shown in Figure 1 (b).* □

Voting trees for weighted profiles work as in the case of unweighted profiles, except that the winner of every pairwise competition is computed by the *weighted* majority rule (and not by the classical majority rule), where A beats B if and only if there is a weighted majority of votes stating $A > B$.

3 From incomplete profiles to incomplete majority graphs

When votes are only partially known, applying a voting rule may result in some uncertainty about the winner. We first define what we mean by an incomplete profile; then we will make clear what we mean by applying a voting rule to an incomplete preference profile, via the definition of possible and necessary winners. Finally, in the specific case of voting trees, we define heuristics that find a superset of possible winners and a subset of necessary winners by considering the incomplete majority graph induced from the incomplete profile.

3.1 Incomplete profiles, possible and necessary winners

Previously, we assumed agents gave complete preferences over candidates. We now consider the case where agents' preferences may only be partially known.

Definition 9 (incomplete profile) *Let* Ω *be a set of* candidates *and* $\{1, \ldots, n\}$ *a set of* voters. *An* incomplete preference relation $>$ *on* Ω *is a strict order on* Ω*, that is, a transitive and irreflexive relation on* Ω*. An (*n*-voter)* incomplete profile *on* Ω *is a collection* $P = (P_1, \ldots, P_n)$ *of incomplete preference relations on* Ω*.*

Example 6 *Assume there are three voters and three candidates* A*,* B*, and* C*, then a possible incomplete profile is (*$(A > B > C)$*; (*$B > A$*); (*$A > B$*)).* □

This definition naturally extends to weighted voters: a weighted incomplete profile is a weighted collection of incomplete preference relations.

Definition 10 (complete extension (aka completion)) *Let $P = (P_1, \ldots, P_n)$ be an incomplete profile over a set of candidates Ω. A complete extension (or completion) R of P is a tuple $(R_1, \ldots, R_n)$ such that every R_i is a strict total order on Ω containing P_i. We will denote by Ext(P) the set of all complete extensions of P.*

Definition 11 (possible and necessary winners) *Let r be a voting rule, P be an incomplete profile and A a candidate. A is a possible winner for P and r, denoted by $A \in$ Poss(P, r), if there exists a completion Q of P such that $r(Q) = A$. A is the necessary winner for P and r, denoted by $A \in$ Nec(P, r), if for any completion Q of P we have $r(Q) = A$.*

Possible and necessary winners were introduced in [21]. Xia and Conitzer showed that, for unweighted votes and an unbounded number of candidates, for most commonly used voting rules computing possible winners is NP-hard, whilst computing necessary winners is coNP-hard for some voting rules but polynomial for others [38]. Also, Betzler and Dorn [3, 5], and Baumeister and Rothe [2], showed that, for an unbounded number of candidates and unweighted votes, computing possible winners is NP-hard for some classes of scoring rules and polynomial for others.

These definitions apply to any voting rule, and thus to V(T) and Schwartz. For instance, A is a *possible Schwartz winner* for P if it is a possible winner with respect to the incomplete profile P and the Schwartz rule, i.e., $A \in$ Poss$(P, r_{Schwartz})$.

We also define possible and necessary Condorcet winners [21], and possible and necessary fair winners:

Definition 12 (possible and necessary Condorcet winners) *Let P be an incomplete profile and A a candidate, A is a possible Condorcet winner for P if there exists a completion Q of P such that A is the Condorcet winner for Q, and the necessary Condorcet winner for P if for every completion Q of P, A is the Condorcet winner for Q.*

Definition 13 (possible and necessary fair winners) *Let P be an incomplete profile and A a candidate, A is a* possible fair winner *for P if there exists a completion Q of P such that A is a fair winner for Q, and the* necessary fair winner *for P if for every completion Q of P, A is the fair winner for Q.*

3.2 Winners defined from the majority graph

We now consider an abstraction of the incomplete profile: the incomplete majority graph.

Definition 14 (incomplete majority graph) *Given an incomplete (weighted) profile P with n voters, the incomplete majority graph $M(P)$ induced by P is the graph whose set of vertices is Ω and where there is an edge from A to B, denoted by $A >_m^P B$ (or by the abbreviated form $A >_m B$), if the number of voters (or the sum of the weights of the voters) who prefer A to B is greater than $n/2$ (half of the voters' total weight).*

$M(P)$ is an incomplete asymmetric graph (or *weak tournament*) over Ω.

Example 7 *Consider the incomplete profile shown in Example 6, i.e., (($A > B > C$); ($B > A$); ($A > B$)), and assume that the weight of the first agent is* 1 *and the weight of every other agent is* 2*. The majority graph induced by this incomplete weighted profile, that we call* Q*, is the graph with three nodes* A*,* B*, and* C *and only one edge* $A >_m^Q B$*.* □

A *completion* of $M(P)$ is a (complete) tournament containing $M(P)$. Note that the set of all completions of $M(P)$ is a *superset* of the set of majority graphs induced by all possible completions of P. Formally, let *Ext*$(M(P))$ be the set of all complete extensions of $M(P)$ and $M(\mathit{Ext}(P)) = \{M(R) \mid R \in \mathit{Ext}(P)\}$ the set of all majority graphs induced from the extensions of P. Then, the following result holds.

Proposition 1 $M(\mathit{Ext}(P)) \subseteq \mathit{Ext}(M(P))$.

This inclusion can be strict, which means that moving from P to $M(P)$ implies a loss of information. An incomplete majority graph throws away information about how individual agents have voted. For example, we lose information about correlations between votes. Such correlations may prevent a candidate from being able to win.

Example 8 *Consider an incomplete profile* P *with just one agent and* $\Omega = \{A, B, C\}$*, where the only vote is* $A > B$ *(the relations between* A *and* C*, and between* B *and* C *have not been specified). The induced majority graph* $M(P)$ *has only one edge from* A *to* B*. In this situation, there is a completion of* $M(P)$ *(with edges from* B *to* C *and from* C *to* A*) and an agenda where* B *wins (we first compare* A *with* C*,* C *wins, and then* C *with* B*, and* B *wins). However, there is no way to complete profile* P *and set up the agenda so that* B *wins. In fact, the possible completions of* P *are* $A > B > C$*,* $A > C > B$*, and* $C > A > B$*, and in all these cases* B *is always beaten at least by* A*. Hence, there is no agenda where* B *wins. Note that the completion of the majority graph which allows us to conclude that* B *can win, cannot be obtained from any possible completion of the agent's preferences of* P *since it violates transitivity.* □

Clearly, the more "complete" P is, the more complete *Ext*(P) is ($M(P)$ can even be complete without P being so) and the better $M(P)$ approximates $M(\mathit{Ext}(P))$.

For any rule r based on the majority graph, including V(T) and Schwartz, we can define notions of the possible (respectively, necessary) winner from the incomplete majority graph.

Definition 15 (possible/necessary winners from the incomplete majority graph) *Let* Q *be a complete profile,* A *a candidate, and* r *a voting rule based on the majority graph, that is, there is a function* f_r *such that* $r(Q) = f_r(M(Q))$*. Let* P *be an incomplete profile,* A *is a possible winner for* $M(P)$ *and* r *if there exists a completion* M' *of* $M(P)$ *such that* $A = f_r(M')$*.* A *is the necessary winner for* $M(P)$ *and* r *if for every completion* M' *of* $M(P)$ *we have* $A = f_r(M')$*. We denote by Poss*$(M(P), r)$ *and Nec*$(M(P), r)$[4] *respectively the set of possible and necessary winners for* $M(P)$ *and* r*.*

[4]Note that, given a voting rule r, *Poss*(x, r) and *Nec*(x, r) represent two different functions, depending on whether the first argument x is a profile or a majority graph.

Note that these notions of possible and necessary winners only apply to rules that are based on the majority graph. Note also that *Poss*$(M(P), r) \supseteq$ *Nec*$(M(P), r)$, and that when the majority graph is complete, *Poss*$(M(P), r) =$ *Nec*$(M(P), r) = r(P)$.

We also define the possible and necessary Condorcet winners from a majority graph in a similar way:

Definition 16 *Let P be a profile and A a candidate, A is a possible Condorcet winner for $M(P)$ if there exists a completion M' of $M(P)$ such that A beats every other candidate in M'. A is the necessary Condorcet winner for $M(P)$ if for every completion M' of $M(P)$, A beats every other candidate in M'.*

We will denote by *PossCond*(P) and *NecCond*(P) (respectively, *PossCond*$(M(P))$ and *NecCond*$(M(P))$) the possible and necessary Condorcet winners from the profile P (respectively, from the majority graph $M(P)$).

Example 9 *Consider the profile P given in Example 7. We have that NecCond$(P) =$ Nec$(P, r_{Schwartz}) = \emptyset$, PossCond$(P) = \{A, C\}$, and Poss$(P, r_{Schwartz}) = \{A, B, C\}$. Notice that in this example the possible and necessary Schwartz and Condorcet winners obtained by considering the completions of P coincide with those obtained by considering the completions of the majority graph induced by P. However, as shown in Example 8, this is not true in general.*

Given an incomplete unweighted profile P and the incomplete majority graph G induced by P, that is, $G = M(P)$, we already observed that the completions of G are a (possibly proper) superset of the set of complete majority graphs induced by all possible completions of P. This observation leads to the following results.

Proposition 2 *Let P be an incomplete weighted profile and r be a rule based on the majority graph.*

1. *Poss*$(M(P), r) \supseteq$ *Poss*(P, r).
2. *Nec*$(M(P), r) \subseteq$ *Nec*(P, r).
3. *PossCond*$(P) =$ *PossCond*$(M(P))$.
4. *NecCond*$(P) =$ *NecCond*$(M(P))$

Proof: We first give the proof when P is unweighted. Points 1 and 2 are straightforward consequences of Proposition 1.

For point 3, if a candidate A is a possible Condorcet winner for P then, by Proposition 1, it follows that it is also a possible Condorcet winner for $M(P)$. Also the converse holds. If a candidate A is a possible Condorcet winner for $M(P)$, then there must be one or more completions of the majority graph where A is the Condorcet winner, i.e., where there is an outgoing edge from A to all the other candidates. We now show that there is completion of P in which A is the Condorcet winner. For this we complete P by first adding, to every agent preference relation $>_i$ and every candidate C such that neither $A >_i C$ nor $C >_i A$, the relation $A >_i C$. Let $>'_i$ the relation

obtained. First, we claim that $>_i'$ has no cycles. Assume that $>_i'$ has a cycle, then there exists an agent i and a cycle in $>_i'$ involving A (otherwise the cycle would already be in $>_i$, which is not possible because $>_i$ is a strict partial order). Without loss of generality, this cycle is $A >_i' C_1 >_i' C_2 >_i' \ldots >_i' C_m >_i' A$, for some $\{C_1, \ldots, C_m\} \subseteq \Omega \setminus \{A\}$, where $m \geq 1$. Because every C_i is different from A, $C_i > C_{i+1}$ was already in $>_i$, as well as $C_m > A$, that is, $C_1 >_i C_2 >_i \ldots >_i C_m >_i A$. Since $>_i$ is transitive, this implies that $C_1 >_i A$. This contradicts the fact that $A >_i C_1$ was added to $>_i$, because the relation between A and C_1 was *not* unspecified in $>_i$. Now, for every agent i, let $>_i^*$ be the transitive closure of $>_i'$. Because $>_i'$ has no cycles, $>_i^*$ has no cycles and is a partial order. Now, we complete $>_i^*$ by instantiating all remaining unspecified pairs in an arbitrary way that preserves transitivity. Let us denote by P' such a completion of P. According to Proposition 1 we have, $M(P') \in \mathit{Ext}(M(P))$, therefore in $M(P')$ there is an outgoing edge from A to all the other candidates and thus A is the Condorcet winner for P', and therefore a possible Condorcet winner for P.

Now for point 4, the fact that the necessary Condorcet winner for a majority graph $M(P)$ is the necessary Condorcet winner for the profile P is an obvious consequence of Proposition 1. For the converse inclusion, let us suppose A is not the necessary Condorcet winner for $M(P)$. Then there exists B such that $(A, B) \notin M(P)$, that is, in the profile we have $\#\{voter\, i \mid A >_i B\} < \frac{n}{2}$. Let us complete the profile by adding $B >_j A$ for all voters j for which the relation between A and B is unspecified: we get an extension of P in which a majority of voters prefers B to A, therefore A is not the necessary Condorcet winner for P.

The proof can be easily extended to weighted profiles. Let P be an incomplete weighted profile, $\mathit{Poss}(M(P),\ r) \supseteq \mathit{Poss}(P, r)$, $\mathit{Nec}(M(P), r) \subseteq \mathit{Nec}(P, r)$, $\mathit{NecCond}(M(P)) = \mathit{NecCond}(P)$, and $\mathit{PossCond}(M(P)) = \mathit{PossCond}(P)$. In fact, obviously, $M(P) = M(U(P))$, and it is easy to show that $\mathit{NecCond}(P) = \mathit{NecCond}(U(P))$ and $\mathit{PossCond}(P) = \mathit{PossCond}(U(P))$. We recall that $U(P)$ is the corresponding unweighted profile of P (see Definition 8). □

Notice that there are cases in which the subset relation $\mathit{Poss}(M(P), r_{Schwartz}) \supseteq \mathit{Poss}(P, r_{Schwartz})$ is strict. In fact, a candidate can be a Schwartz winner for a completion of $M(P)$ which is not induced by any completion of P, as shown previously in Example 8.

Note that if r is a Condorcet-consistent rule, then $\mathit{NecCond}(P) \subseteq \mathit{Nec}(P, r)$ and $\mathit{PossCond}(P) \subseteq \mathit{Poss}(P, r)$.

Notice also that the set of winners from an incomplete profile and from its incomplete majority graph coincide when we allow also for completions of the preferences that may contain cycles, i.e., when we allow that the voters are irrational. Such preferences do occur in practice and have been explored in the context of computational social choice [14, 15].

4 Computing winners with incomplete preferences in voting trees

We now analyze the computational complexity of computing the various kinds of winners in voting trees when preferences are incomplete[5] . We will consider first winners from an incomplete profile and then winners from an incomplete majority graph.

4.1 Incomplete profiles

Assume agents' preferences have been specified by an incomplete weighted profile P. It is possible to show that, when we have 3 or more candidates, it is difficult to determine if a candidate is a possible winner for $V(T)$, even if T is a balanced tree.

Theorem 2 *Given a fixed number m of candidates with $m \geq 3$, an incomplete weighted profile P and any agenda T on these candidates, deciding if a candidate is in Poss$(P, r_{V(T)})$ is NP-complete, even if every candidate appears exactly once in T, and even if T is a balanced tree.*

Proof: We give a reduction from the number partitioning problem. Consider the tree T where A plays against B and the winner thus plays against C. We have a bag of integers, k_i with sum $2k$ and we wish to decide if they can be partitioned into two bags, each with sum k. We will show that we can build an election where we can complete the incomplete weighted profile so that C wins (i.e., C is a possible winner) if and only if such a partition exists. We suppose the following votes are given: 1 vote for $C > B > A$ of weight 1, 1 vote $C > A > B$ of weight $2k-1$, and 1 vote $B > C > A$ of weight $2k-1$. Hence, C is ahead of A by a weight of $4k-1$, C is ahead of B by a weight of 1 and B is ahead of A by a weight of 1. For each k_i in the bag of integers, we have an incomplete vote of weight $2k_i$ in which $A > C$ is fixed, but the rest of the vote is incomplete. We are sure A beats C in the final result by a weight of 1 whatever completion takes place. We now show that the incomplete weighted profile can be completed so that B beats A and C beats B (and thus C is a possible winner) if and only if there is a partition of size k.

Assume that such a partition exists, and that votes in one partition have $A > C > B$ and the votes in the other have $B > A > C$. Thus, B beats A overall and C beats B. Thus C is the winner. On the other hand, suppose there is a way to complete the preferences so that C wins. This can only happen if B beats A and C then beats B. In fact, if A beats B in the first round, A will beat C in the second round, since we have shown before that A beats C overall, and then A will be the final winner. For C to beat B, at least half the weight of incomplete votes must rank C above B. Similarly, for B to beat A, at least half the weight of incomplete votes must rank B above A. Since all votes rank A above C, B cannot be both above A and below C. Thus precisely half the weight of incomplete votes ranks B above A and half ranks C above B. Hence, we have a partition of equal weight. Therefore, we can complete the incomplete profile so that C wins if and only if there is a partition of size k. The result can also be extended

[5] For a complete treatment of the basic notions of complexity theory, see [18].

to the trees with more than 3 candidates which extend T or a tree isomorphic to it, by placing any additional candidate at the bottom of every voters' preference ordering (it does not matter how).

We can use the same construction above when T is a balanced tree. Given the profile constructed here, the only possible balanced trees in which C wins are those in which A plays against B, and the winner then plays against C. All the additional candidates will be defeated by A, B and C, so that they can be placed anywhere in the balanced tree. □

In [29] it is shown that, when we have 3 candidates and the agenda is a simple voting tree, the necessary winner from the incomplete profile and from the incomplete majority graph coincide. Therefore, since by Theorem 4, it is polynomial to compute the necessary winner from the incomplete majority graph, it is also polynomial to compute the necessary winner from the incomplete profile. On the other hand, we now show that, when we have 4 or more candidates, it is a coNP-complete problem to determine if a candidate is the necessary winner.

Theorem 3 *Given a fixed number m of candidates with $m \geq 4$, an incomplete weighted profile P and any agenda T on these candidates, deciding if a candidate is in Nec$(P, r_{V(T)})$ is coNP-complete, even if every candidate appears exactly once in T.*

Proof: To show that with at least 4 candidates it is coNP-complete to decide if a candidate wins for every completion of the preferences, we will consider the tree T where A plays against B, the winner then plays against C, and the winner of this match goes forward to the final match against D. We will reduce number partitioning to deciding if, given a particular incomplete weighted profile, we can complete such a profile to make a given candidate win.

We have a bag of integers, k_i with sum $2k$ and we wish to decide if they can be partitioned into two bags, each with sum k. We construct an incomplete profile where the following votes are given: 1 vote for $C > D > B > A$ of weight 1, 1 vote $C > D > A > B$ of weight $2k - 1$, and 1 vote $D > B > C > A$ of weight $2k - 1$. For the first number, k_1 in the bag of integers, we have a vote for $D > B > A > C$ of weight $2k_1$. For each other number, k_i where $i > 1$, we have an incomplete vote of weight $2k_i$ in which $A > C$ is fixed, but the rest of the vote is incomplete. We are sure that A beats C in the final result by 1 vote whatever completion takes place. Similarly, we are also sure that D beats A, and D beats B.

If we complete preferences so that in all incomplete votes D beats C in their pairwise election, then D will win overall. We now show that there is a completion that makes C win if and only if there is a partition of equal weight.

Suppose there is such a partition and that the complete votes in one partition have $B > A > C$ and the complete votes in the other have $A > C > B$. Thus, B beats A overall, C beats B. We put D in the bottom position in all such completions as far as possible, and so C beats D. Thus C is the overall winner.

On the other hand, suppose there is a way to complete the preferences so that C wins. This can only happen if B beats A, C then beats B and C finally beats D. If A beats B in the first round, A will beat C in the second round and then A will be beaten

by D. For C to beat B, at least half the weight of incomplete votes must rank C above B. Similarly, for B to beat A, at least half the weight of incomplete votes must rank B above A. Since all votes rank A above C, B cannot be both above A and below C. Thus precisely half the weight of incomplete votes ranks B above A and half ranks C above B, and we have a partition of equal weight. Hence, if there is a partition of equal weight then both C and D are possible winners; if not, then only D is a possible winner, therefore the necessary winner. Therefore, D is the necessary winner if and only if there is no partition in the initial problem. The proof can be extended to all the trees with more than 4 candidates, when they extend T or a tree isomorphic to it, by placing any additional candidate at the bottom of every voter's preference ordering (it does not matter how). □

These results can be compared with Theorem 7 in [9] which states that weighted constructive manipulation for $V(T)$ is polynomial. A candidate can be made a winner by a coalition of weighted votes if it is a possible winner for a specific incomplete profile consisting of full votes and empty votes. By allowing any kind of incomplete profile, we increase the computational complexity.

The possible winner (respectively, necessary winner) problem for $V(T)$ was proved to be NP-complete (respectively, coNP-complete) for unweighted votes and an unbounded number of candidates even when T is balanced [39] and easy to compute for unweighted votes and a bounded number of candidates [36].

4.2 Incomplete majority graphs

In this Section (unlike the rest of the paper) we focus on *simple voting trees*, that is, voting trees where every candidate appears exactly once. We present an algorithm, called *Win*, for determining, given a simple voting tree T and an incomplete majority graph G, the set of possible winners (i.e., *Poss*$(G, r_{V(T)})$). We will represent a simple voting tree T with a binary tree with a root, that we call *root*(T), a left subtree, called *left*(T), and a right subtree, called *right*(T).

Algorithm *Win* recursively takes in input a simple voting tree T and an incomplete majority graph G, and it returns a set of candidates W, which is the set of possible winners. If *root*(T) is not empty, and both *left*(T) and *right*(T) are empty, then the algorithm returns *root*(T). Otherwise, the set of winners at the root of T is the set of all candidates who are possible winners in the left (respectively, right) branch of T and who beat at least one candidate who is a possible winner in the right (respectively, left) branch of T.

To check whether there exists a necessary winner for an incomplete majority graph G, it suffices to run the algorithm *Win* on P; if it outputs a single candidate, then it is the necessary winner, otherwise there is no necessary winner. We will call such a procedure *NecessaryWin*.

Example 10 *We now show how to determine possible and necessary winners given a fixed simple voting tree by applying Win. Let* $\Omega = \{A, B, C, D, E, F, H, I\}$ *and consider the simple voting tree* T *over* Ω *with left*$(root(T)) =$ *Win*$($*Win*$(\{C\}, \{D\})$, *Win*$(\{E\}, \{F\}))$ *and right*$(root(T)) =$ *Win*$($*Win*$(\{A\}, \{B\})$, *Win*$(\{I\}, \{H\}))$.

Algorithm 1: *Win*

```
Input: T: a simple voting tree, G: an incomplete majority graph;
Output: W: set of candidates;
if T contains only one node then
    W ← label(root(T))
else
    W_1 ← Win(left(T), G);
    W_2 ← Win(right(T), G);
    W ← ∅;
    foreach (s, t) ∈ W_1 × W_2 do
        if s >_m t then
            W ← W ∪ {s}
        else
            if t >_m s then
                W ← W ∪ {t}
            else
                W ← W ∪ {s, t}
return W
```

Consider also the incomplete majority graph G with edges $A >_m B$, $A >_m C$, $A >_m D$, $A >_m E$, $A >_m I$, $E >_m F$, and $I >_m H$.

The application of Win to the subtree of T that contains only C and D gives Win$(\{C\}, \{D\}) = \{C, D\}$, because G contains no edge between C and D. Its application to the subtree that contains only E and F gives Win$(\{E\}, \{F\}) = \{E\}$, because G contains an edge from E to F. Next, the application of Win to the left subtree of T gives Win(Win$(\{C\}, \{D\})$, Win$(\{E\}, \{F\})) =$ Win$(\{C, D\}, \{E\}) = \{C, D, E\}$, because E is not beaten by both C and D, C is not beaten by E and D is not beaten by E. As for the right subtree of T, we get Win(Win$(\{A\}, \{B\})$, Win$(\{I\}, \{H\})) =$ Win$(\{A\}, \{I\}) = \{A\}$. Finally, the application of Win to the root of T gives Win$(\{C, D, E\}, \{A\}) = \{A\}$, because A beats all of C, D and E in the graph G. Thus, Win returns $\{A\}$. Since the set returned by Win contains only one candidate, A is also the necessary winner. □

Theorem 4 *Consider a set of candidates Ω, an agenda T over Ω, and a possibly incomplete majority graph G over Ω. Then,*

- *Algorithm Win(T, G) returns Poss$(G, r_{V(T)})$.*
- *Algorithm NecessaryWin(T, G) returns Nec$(G, r_{V(T)})$.*
- *Both algorithms run in polynomial time in the number of candidates.*

Proof: Let us first consider Algorithm *Win*. We prove by induction that the following induction hypothesis holds for any tree T:

$$H(T)\text{: } \mathit{Poss}(G, r_{V(T)}) = \mathit{Win}(T, G)$$

We omit G in $H(T)$ because the induction is over trees, and not on G.

If T is a tree with a single node, then $H(T)$ holds trivially, because T contains only one candidate, namely *label*(*root*(T)), and we have *Win*(T, G) = *Poss*$(G, r_{V(T)})$ = {*label*(*root*(T))}. Now we need to show that if T is a tree whose left subtree is T_1 and right subtree is T_2, then $H(T_1)$ and $H(T_2)$ imply $H(T)$. Assume now that $H(T_1)$ and $H(T_2)$ hold.

Assume that $A \in$ *Win*(T, G). Then either $A \in$ *Win*(T_1, G) or $A \in$ *Win*(T_2, G). Without loss of generality, assume $A \in$ *Win*(T_1, G). Then there exists a B in *Win*(T_2, G) such that G either contains an edge from A to B or no edge between A and B.

By the induction hypothesis, $A \in$ *Win*(T_1, G) and $B \in$ *Win*(T_2, G) imply $A \in$ *Poss*$(G, r_{V(T_1)})$ and $B \in$ *Poss*$(G, r_{V(T_2)})$, which in turn imply that there exist two complete tournaments G_1 and G_2 extending G such that $A = r_{V(T_1)}(G_1)$ and $B = r_{V(T_2)}(G_2)$. Now, let G_3 be the tournament obtained from G_1 and G_2 in the following way: for any two candidates X and Y, (a) if both X and Y appear in T_1 then G_3 contains the edge $X >_m Y$ if and only if G_1 contains it; (b) if both X and Y appear in T_2 then G_3 contains the edge $X >_m Y$ if and only if G_2 contains it; (c) G_3 contains the edge $A >_m B$; (d) in all other cases, if G contains an edge $X >_m Y$ then G_3 contains it as well, and if G does not contain any edge between X and Y then we can decide arbitrarily whether there is an edge from X to Y or an edge from Y to X. By construction, G_3 is an extension of G, because both G_1 and G_2 are and because G either contains or does not contain the edge $B >_m A$.

Now, because no candidate appears in more than one leaf of T, the sets of candidates appearing in T_1 and T_2 are disjoint. Therefore no pair of candidates can be concerned by more than one of the cases (a), (b), and (c), and the definition of G_3 is well-founded. Now, because the way $r_{V(T_1)}(G)$ (respectively, $r_{V(T_2)}(G)$) is determined depends only on the restriction of the graph on the candidates appearing in T_1 (respectively, T_2), we have that $r_{V(T_1)}(G_3) = r_{V(T_1)}(G_1) = A$ and $r_{V(T_2)}(G_3) = r_{V(T_2)}(G_2) = B$. From this and the fact that G_3 contains the edge $A >_m B$, we get that $r_{V(T)}(G_3) = A$. Thus we have found a complete extension of G in which the winner of $r_{V(T)}$ is A, which shows that $A \in$ *Poss*$(G, r_{V(T)})$.

Conversely, assume that $A \in$ *Poss*$(G, r_{V(T)})$. Then there is a completion G' of G such that $r_{V(T)}(G') = A$. Obviously, A appears in T. Without loss of generality, assume it appears in T_1. Because T_1 is a subtree of T, $r_{V(T)}(G') = A$ implies $r_{V(T_1)}(G') = A$; therefore, since G' is an extension of G, we have $A \in$ *Poss*$(G, r_{V(T_1)})$, and by the induction hypothesis, we have $A \in$ *Win*(T_1, G). Now, $r_{V(T)}(G') = A$ implies that A beats $r_{V(T_2)}(G')$ in G'. Denote $r_{V(T_2)}(G')$ by B. Again by the induction hypothesis, we have $B \in$ *Poss*$(G, r_{V(T_2)})$. Finally, G does not contain the edge $B >_m A$, otherwise G' (which contains the edge $A >_m B$) would not be an extension of G. This implies that *Win*(T, G) contains A. We now have proven that $H(T)$ holds, and therefore we have shown that the algorithm *Win* is sound and complete, *i.e.*, returns the set of possible winners for $r_{V(T)}$.

Finally, if we have m candidates, Algorithm *Win* performs at most $O(m)$ steps, where at each step it may consider $O(m^2)$ pairs of candidates. For each pair it performs a constant amount of work. Thus, the overall complexity is $O(m^3)$. The same

complexity holds for Algorithm *NecessaryWin*. □

Since *Poss*$(M(P),\ r_{V(T)}) \supseteq$ *Poss*$(P,\ r_{V(T)})$, Algorithm *Win* is a heuristic that finds a superset of *Poss*$(P, r_{V(T)})$.

Notice that, in general, there is no guarantee on the quality of this heuristic, as shown by the following example.

Example 11 *Assume we have an incomplete profile, say P, over $n+2$ candidates, say A, $B_1, \ldots, B_n$, and C, where there is only one voter that expresses his preferences as follows: $A > B_i$, $\forall i$. Assume also that the agenda T has A first play against C. For example, T may be the agenda where first A plays against C, the winner plays against B_1, and each new winner at step i plays against B_{i+1}. Then, Poss$(M(P), r_{V(T)}) = \{A, B_1, \ldots, B_n, C\}$, while Poss$(P, r_{V(T)}) = \{A, C\}$. This means that all candidates are possible winners for the incomplete majority graph $M(P)$, while only two of them, i.e., A and C, are possible winners for the incomplete profile P. In fact, a candidate B_i may win only if C beats A and B_i beats C. To achieve this, there must be a completion with a cycle $A > B_i > C > A$, which is allowed in the majority graph, but it is forbidden by transitivity in the profile.* □

Since *Poss*$(P, r_{V(T)}) \subseteq$ *Poss*$(M(P), r_{V(T)})$, the quality of the heuristic can be measured by the ratio between the number of winners returned by *Poss*$(M(P), r_{V(T)})$ and *Poss*$(P, r_{V(T)})$. Example 11 shows that this ratio can be as bad as $\frac{m+2}{2}$.

However, we are able to give an upper bound on the number of possible winners from the incomplete majority graph, which depends on the number of the missing arcs in the incomplete majority graph.

Proposition 3 *Let P be an incomplete profile over m candidates, $M(P)$ its corresponding incomplete majority graph, and k the number of the missing arcs in $M(P)$. |Poss$(M(P), r_{V(T)})| \leq 2^k$.*

Proof: If there are k missing arcs in $M(P)$, then there are 2^k possible completions of $M(P)$. Therefore, there are at most 2^k results and so there are at most 2^k possible winners. □

Parameterizations based on the total number of undetermined candidate pairs have also been considered in [4] to analyze the complexity of the possible winner problem for some voting rules. However, these voting rules do not include the voting tree rule.

4.3 Experimental evaluation of the heuristic

We tested computationally the quality of the heuristic given by Algorithm *Win* in computing an upper bound on the set of possible winners.

4.3.1 Experimental setting

To do this, we generated profiles according to the following parameters:

- number of candidates $m \in \{4, 8, 16, 32\}$;
- number of agents $n \in \{3, 5, 7, 9, 11\}$;
- type of agenda:
  - completely balanced: a balanced simple voting tree,
  - completely unbalanced: a simple voting tree where each internal node has at least one child which is a leaf.
- level of completeness of the preference of each agent. We considered three levels of completeness: *High*, also denoted by *H* (80% of the relations of a complete linear order), *Medium*, also denoted by *M* (50% of the relations of a complete linear order) and *Low*, also denoted by *L* (20% of the relations of a complete linear order).

For each combination of the above parameters, we randomly generated 100 instances according to the impartial culture assumption, that is, using a uniform distribution over profiles. For each incomplete profile P and agenda T, we computed the set of possible winners *Poss*$(M(P), T)$ as by Algorithm *Win* and then we checked for each candidate in such a set if it is indeed a possible winner from the profile, that is, if it is in *Poss*(P, T).

We then measured the following:

- the error rate: the percentage of instances where *Poss*$(M(P), T)$ is larger than *Poss*(P, T);
- the error mean: the ratio $\frac{|\mathit{Poss}(M(P),T)|-|\mathit{Poss}(P,T)|}{|\mathit{Poss}(P,T)|\cdot(m-1)}$ averaged over the instances in which there is an error. The error is $|Poss(M(P), T)| - |Poss(P, T)|$ divided by $|Poss(P, T)|$. We have then multiplied it by $1/(m - 1)$ in the definition of the error mean, since we wanted to have a value which is in $[0, 1]$ independently of the number of candidates.

All the experiments were done on an Intel Xeon 3.2GHz machine with 16GB of RAM.

4.3.2 IP model for incomplete profiles

To generate the possible winners from an incomplete profile, due to the combinatorial nature of this problem, we adopt a Integer Programming model (IP)[6], defined as follows, and solved with IBM ILOG Cplex 12.2.

We are given m candidates, n agents, and a fixed agenda T. We want a model that proves that a given candidate is a possible winner, by looking for a profile completion where such a candidate wins, or proving that no such profile exists.

Let us introduce the following variables:

$$x_{ijk} = \begin{cases} 1 & \text{if candidate } i \text{ is preferred to candidate } j \text{ according to agent } k \\ 0 & \text{otherwise} \end{cases}$$

[6]Notice that we are not trying to optimize a specific function. We will consider a feasibility program.

$$y_{ij} = \begin{cases} 1 & \text{if candidate } i \text{ is preferred to candidate } j \text{ according to the majority graph} \\ 0 & \text{otherwise} \end{cases}$$

For each agent k, variables x_{ijk} define the ordering of the candidates according to agent k: such an order can be seen as a permutation of m elements and so it can be easily modeled with linear constraints (as done, for example, in the linear ordering problem). In particular, the following constraints suffice:

$$x_{ijk} + x_{jik} \leq 1 \quad \forall (i,j)$$

$$x_{ijk} + x_{jlk} + x_{lik} \leq 2 \quad \forall (i,j,l)$$

Next we need to link variables y to variables x with the logical constraint

$$y_{ij} = 1 \leftrightarrow \sum_k x_{ijk} \geq \lceil n/2 \rceil \quad \forall (i,j)$$

and this also can be modeled with the following linear constraints:

$$\lceil n/2 \rceil y_{ij} \leq \sum_k x_{ijk} \leq \lfloor n/2 \rfloor + \lceil n/2 \rceil y_{ij} \quad \forall (i,j)$$

Finally, we must compute the winner for each node a of the tree. To this end, we introduce variables

$$w_{ai} = \begin{cases} 1 & \text{if candidate } i \text{ is the winner at node } a \\ 0 & \text{otherwise} \end{cases}$$

$$z_{aij} = \begin{cases} 1 & \text{if } (i,j) \text{ is the pair of possible winners at node } a \\ 0 & \text{otherwise} \end{cases}$$

Some constraints are needed to ensure the correct computation of the winner at node a starting from the winners at its child nodes l and r

$$w_{ai} \rightarrow w_{li} \vee w_{ri} \quad \forall (a,i)$$

$$z_{aij} \rightarrow w_{ai} \geq y_{ij} \quad \forall (a,i,j)$$

$$z_{aij} = (w_{li} \vee w_{ri}) \wedge (w_{li} \vee w_{ri}) \quad \forall (a,i,j)$$

The first set of constraints says that a candidate i can be a winner at node a only if it is a winner in one of its child nodes. The second set of constraints uses the majority graph to compute the winner among the two possible candidates. The third set of constraints just links variables z and w. All of them can be stated as linear constraints

$$w_{ai} \leq w_{li} + w_{ri} \quad \forall (a,i)$$

$$w_{ai} \geq y_{ij} + z_{nij} - 1 \quad \forall (a,i,j)$$

$$z_{aij} \leq w_{li} + w_{ri} \quad \forall(a,i,j)$$
$$z_{aij} \leq w_{lj} + w_{rj} \quad \forall(a,i,j)$$
$$z_{aij} \geq w_{li} + w_{ri} + w_{lj} + w_{rj} - 1 \quad \forall(a,i,j)$$

Notice that the model has $O(m^2n + m^3)$ variables and $O(m^3n)$ constraints. However, for small numbers of candidates, the size of the model is reasonably compact. Also, several variables can be fixed right from the beginning, the more so if we run algorithm *Win* and use the information collected by it. This variable fixing, together with the logical constraints, allows for a very effective preprocessing and constraint propagation on the model. To test if a candidate i is a possible winner we just need to solve the model by setting $w_{root,i} = 1$.

Notice also that, given m candidates and n agents, we have $(m!)^n$ possible complete profiles. In fact, each agent gives a strict total order over the candidates. Such an ordering is a permutation of the candidates, and thus in the worst case (when there is full incompleteness) we have $m!$ possible orderings for every agent. For example, in an instance with 8 candidates and 11 agents, we have $(8!)^{11}$, that is, more than 10^{50}, possible completions, which cannot be enumerated via a brute-force algorithm.

4.3.3 Experimental results

Figures 2 - 5 show the error rate and the error mean for the cases with 4, 8, 16, and 32 candidates.

It is easy to notice that the shape of the agenda (completely balanced vs. completely unbalanced) does not seem to affect the quality of the heuristic. On the other hand, the level of completeness of the profile does influence the quality of the heuristic, since both the error rate and the error mean decrease as the level of completeness decreases. Indeed, when the profile is rather incomplete, the set of possible winners is larger and thus closer to what the heuristic can return. However, the error mean is always very small. Moreover, the error rate increases with the number of candidates and it decreases as the number of agents increases. Notice that, with 4 candidates and a high level of completeness, the error is always 0, since in this case the profile is almost complete.

The largest instances we have considered (32 candidates and 11 agents) have been solved in about 30 seconds (we need to solve an IP model for each of the candidates returned by the heuristic). Thus, we could certainly tackle much larger instances, but we have preferred to focus on smaller instances and consider 100 of them to get a reasonable average. This is the time needed by the exact algorithm. Instead, the time needed by the heuristics is less than 1 millisecond in all instances. This is expected, since it is just a visit of the nodes of the agenda based on the majority graph.

5 Computing winners with incomplete preferences in the Schwartz rule

We now analyze the complexity of determining Schwartz winners when we have uncertainty in the preferences.

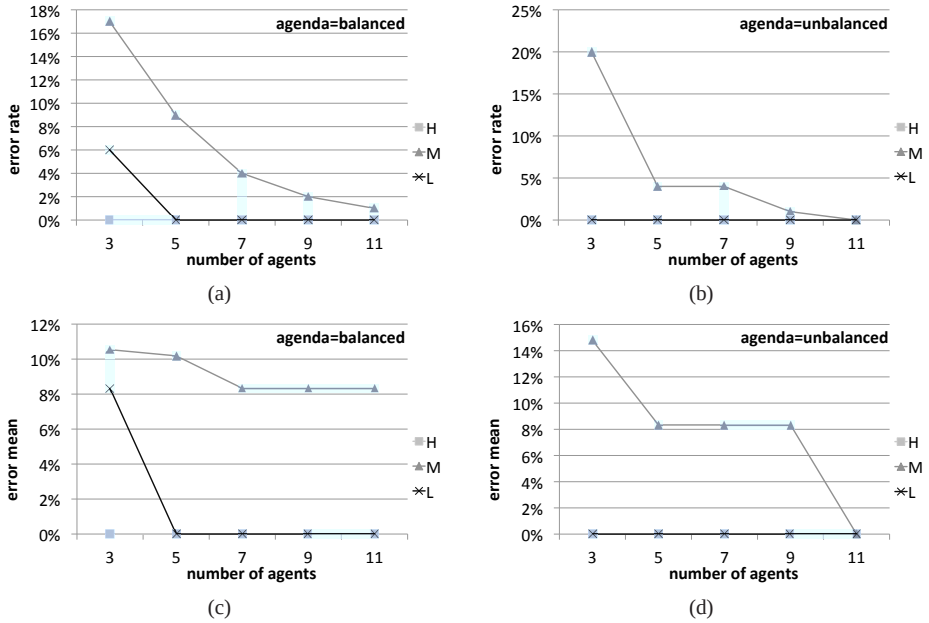

Figure 2: Error rate and error mean for agendas with 4 candidates.

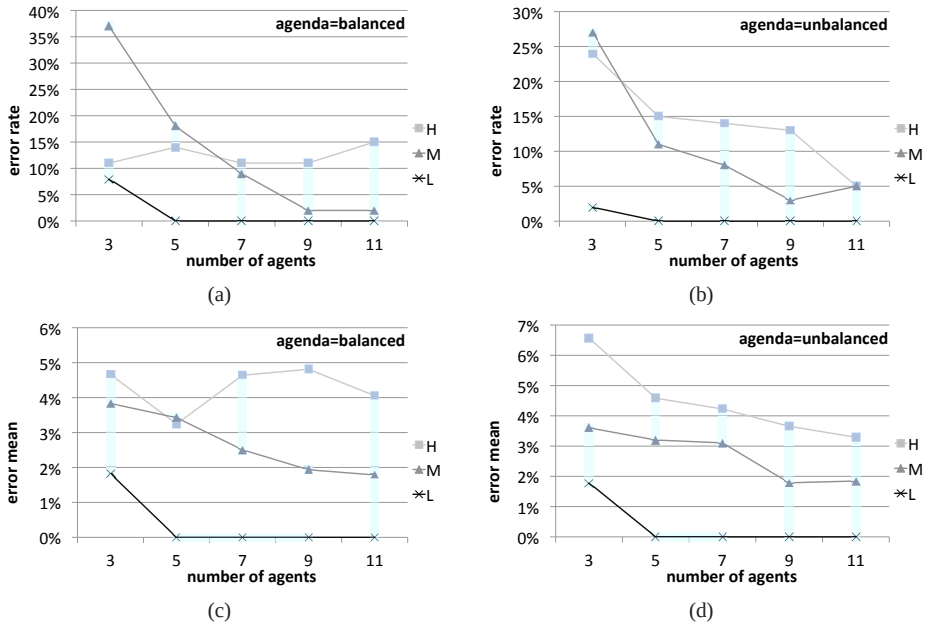

Figure 3: Error rate and error mean for agendas with 8 candidates.

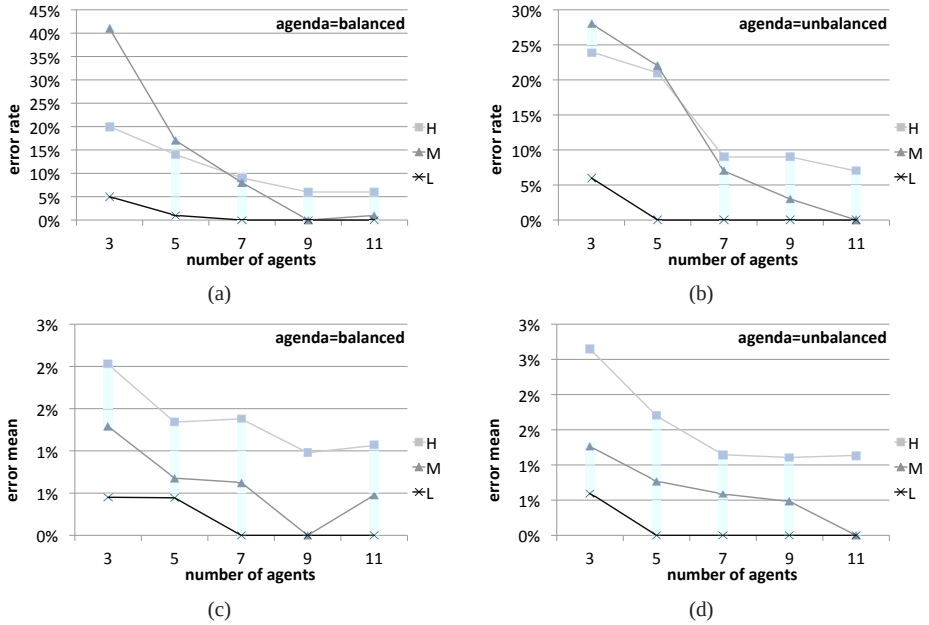

Figure 4: Error rate and error mean for agendas with 16 candidates.

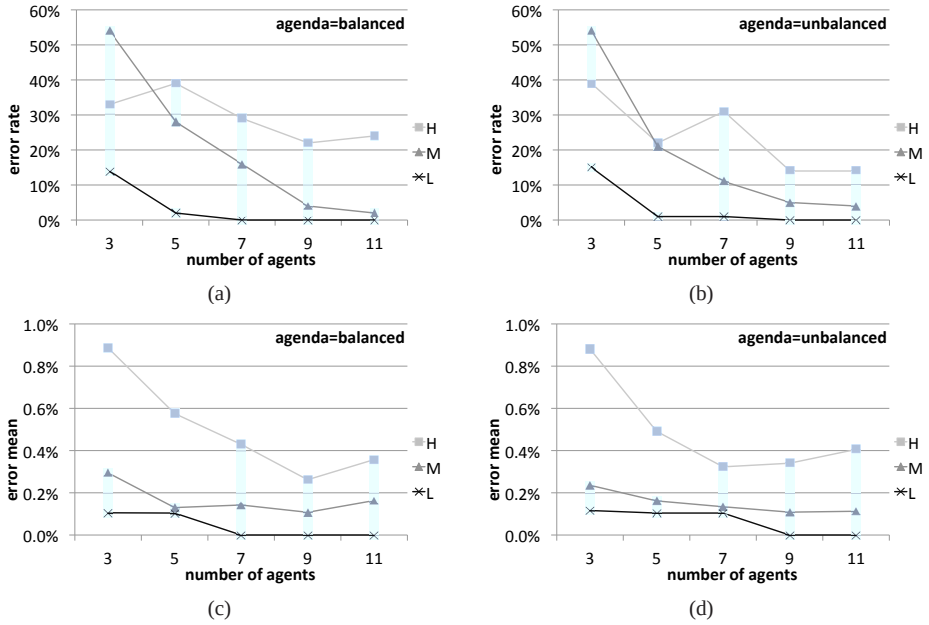

Figure 5: Error rate and error mean for agendas with 32 candidates.

We show that deciding whether a given candidate is a Schwartz winner from incomplete weighted profiles is intractable in general, but polynomial from the induced majority graphs. Hence, computing the winners for incomplete majority graphs can be a computationally efficient heuristic that finds a superset of the winners for incomplete weighted profiles. We conjecture that computing the set of all the possible Schwartz winners for incomplete unweighted profiles is NP-hard as well, but so far have not been able to show this.

5.1 Incomplete weighted profiles

We prove that determining whether a given candidate is a possible Schwartz winner for weighted profiles is NP-hard for incomplete weighted profiles where there are 3 or more candidates.

Theorem 5 *Given a fixed number m of candidates with $m \geq 3$ and an incomplete weighted profile P on these candidates, deciding if a candidate is in Poss$(P, r_{Schwartz})$ is NP-complete.*

Proof: Clearly the problem is in NP as a polynomial witness is a completion and an agenda in which the candidate wins. To show it is NP-complete, we give a reduction from the number partitioning problem. The reduction is based on constructing a Condorcet cycle and is similar to those used in [9] to show that manipulation is computationally hard even with a small number of candidates when votes are weighted.

We have a bag of integers, k_i with sum $2k$ and we wish to decide if they can be partitioned into two bags, each with sum k. We want to show that a candidate B is a possible Schwartz winner if and only if such a partition exists. We construct an incomplete profile over three candidates (A, B, and C) as follows. We have 1 vote for $B > C > A$ of weight 1, 1 vote $B > A > C$ of weight $2k-1$, and 1 vote $C > B > A$ of weight $2k-1$. At this point, the total weight of votes with $B > A$ exceeds that of $A > B$ by $4k-1$, the total weight of votes with $B > C$ exceeds that of $C > B$ by 1, and the total weight of votes with $C > A$ exceeds that of $A > C$ by 1.

We also have, for each k_i, a partially specified vote of weight $2k_i$ in which we know just that $A > B$. As the total weight of these partially specified votes is $4k$, we are sure A beats B in the final result by 1 vote. The only agenda in which B can win is the one in which A plays against C and the winner then plays against B. In addition, for B to win, the partially specified votes need to be completed so that B beats C, and C beats A in the final result. We show that this is possible if and only if there is a partition of size k. Suppose there is such a partition. Then let the votes in one bag be $A > B > C$ and the votes in the other be $C > A > B$. Then, A beats B, B beats C and C beats A, all by 1 vote in the final result. On the other hand, suppose there is a way to cast the votes to give the result A beats B, B beats C, and C beats A. All the uncast votes rank A above B. In addition, at least half the weight of votes must rank B above C, and at least half the weight of votes must rank C above A. Since A is above B, C cannot be both above A and below B. Thus precisely half the weight of votes ranks C above A and half ranks B above C. Hence we have a partition of equal weight. We therefore conclude that B can win if and only if there is a partition of size k. That is, deciding if

B is a possible Schwartz winner is NP-complete. We can extend the reduction to more than 3 candidates by placing any additional candidate at the bottom of every voters' preference ordering in whatever way we wish. □

Computing possible Schwartz winners for weighted profiles seems to be intractable, however we will show that is easy to find a superset of these winners considering the majority graph.

5.2 Incomplete majority graphs

We now give graph characterizations of possible and necessary Schwartz winners for incomplete majority graphs, which allow us to conclude that they can be computed in polynomial time.

Theorem 6 *Given an incomplete majority graph G and a candidate A, $A \in$ Nec$(G, r_{Schwartz})$ if and only if, for every other candidate B, there is a path from A to B in G.*

Proof: ($\Leftarrow$) Suppose that for each $B \neq A$ there is a path from A to B in G. Then these paths remain in every completion of G. Therefore, using Theorem 1, A is a Schwartz winner in every completion of G, i.e., it is the necessary Schwartz winner. ($\Rightarrow$) Suppose there is no path from A to B in G. Let us define the following three subsets of the set of candidates Ω: $R(A)$ is the set of candidates reachable from A in G (including A); $R^{-1}(B)$ is the the set of candidates from which B is reachable in G (including B); and *Others*$= \Omega \setminus (R(A) \cup R^{-1}(B))$. Because there is no path from A to B in G, we have that $R(A) \cap R^{-1}(B) = \emptyset$ and therefore $\{R(A), R^{-1}(B), \textit{Others}\}$ is a partition of Ω. Now, let us build the complete majority graph $\hat{G}$ as follows:

1. $\hat{G} := G$;
2. $\forall\, x \in R(A)$, $y \in R^{-1}(B)$, add (y, x) to $\hat{G}$;
3. $\forall\, x \in R(A)$, $y \in$ *Others* (i.e., $y \in \Omega \setminus (R(A) \cup R^{-1}(B))$), add (y, x) to $\hat{G}$;
4. $\forall\, x \in$ *Others* (i.e., $x \in \Omega \setminus (R(A) \cup R^{-1}(B))$), $y \in R^{-1}(B)$, add (y, x) to $\hat{G}$;
5. $\forall\, x, y$ belonging to the same element of the partition: if neither (x, y) nor (y, x) is in G then add one of them (arbitrarily) in $\hat{G}$.

Let us first show that $\hat{G}$ is a complete majority graph. If $x \in R(A)$ and $y \in R^{-1}(B)$, then $(x, y) \notin G$ (otherwise there would be a path from A to B in G). If $x \in R(A)$ and $y \in$ *Others*, then $(x, y) \notin G$, otherwise y would be in $R(A)$. If $x \in$ *Others* and $y \in R^{-1}(B)$, then $(x, y) \notin G$, otherwise x would be in $R^{-1}(B)$. Therefore, whenever x and y belong to two distinct elements of the partition, $\hat{G}$ contains (y, x) or (x, y), but not both. Now, if x and y belong to the same element of the partition, by Step 5, $\hat{G}$ contains exactly one edge in $\{(x, y), (y, x)\}$. Therefore, $\hat{G}$ is a complete majority graph. Let us show now that there is no path from A to B in $\hat{G}$. Suppose there is one, that is, there exist $z_0 = A, z_1, \ldots, z_{m-1}, z_m = B$ such that

$\{(z_0, z_1), (z_1, z_2), \ldots, (z_{m-1}, z_m)\} \subseteq \hat{G}$. Now, for all $x \in R(A)$ and all y such that $(x, y) \in \hat{G}$, by construction of $\hat{G}$, we necessarily have $y \in R(A)$. Therefore, for all $i < m$, if $z_i \in R(A)$ then $z_{i+1} \in R(A)$. Now, since $z_0 = A \in R(A)$, by induction we have $z_i \in R(A)$ for all i, thus $B \in R(A)$, which is impossible. Therefore, there is no path from A to B in $\hat{G}$. Thus, $\hat{G}$ is a complete majority graph with no path from A to B, which implies that A is not a Schwartz winner with respect to $\hat{G}$. Lastly, by construction, $\hat{G}$ contains G. So $\hat{G}$ is a complete extension of G for which A is not a Schwartz winner. This shows that A is not the necessary Schwartz winner for G. □

A procedure based on the previous theorem gives us a polynomial time algorithm to find necessary Schwartz winners for incomplete majority graphs. Hence, we have a heuristic which computes in polynomial time a subset of necessary Schwartz winners for incomplete profiles.

Corollary 1 *Given an incomplete majority graph G, Nec$(G, r_{Schwartz})$ can be computed in polynomial time.*

Proof: By Theorem 6, we know that $A \in$ *Nec*$(G, r_{Schwartz})$ if and only if, for every other candidate B, there is a path from A to B in G. Since finding a path from a single source to all other candidates is polynomial [13], then also computing *Nec*$(G, r_{Schwartz})$ is polynomial. □

We now consider computing possible Schwartz winners. Let G be an asymmetric graph, Ω the set of candidates, and $A \in \Omega$. Let us consider Algorithm 2. It tries to complete G in order to have a path from A to every other candidate Z by putting edges directed to Z, where it is possible, from the candidates that are reachable from A.

Algorithm 2: PossibleSchwartzWinner

```
Input: G: an asymmetric graph, A: a candidate,
Output: Σ: a set of candidates;
Σ ← {A} ∪ {X | there is a path from A to X in G};
G' ← G;
repeat
    foreach (Y, Z) ∈ Σ × (Ω \ Σ) do
        if (Z → Y) ∉ G' then
            add (Y → Z) to G';
    foreach Z ∈ Ω \ Σ do
        if there is a path from A to Z in G' then
            add Z to Σ;
until Σ = Ω or there is no (Y, Z) ∈ Σ × (Ω \ Σ) s. t. (Z → Y) ∉ G' ;
return Σ
```

Theorem 7 *Given an incomplete majority graph G and a candidate A, $A \in$ Poss$(G, r_{Schwartz})$ if and only if* PossibleSchwartzWinner$(G, A) = \Omega$.

Proof: We start with the following observation: the graph G' obtained at the end of the algorithm is asymmetric and extends G. It is asymmetric because it is so at the start of the algorithm (since G is asymmetric) and it is still asymmetric if an edge $Y \to Z$ is added to G' when $Z \to Y$ is not already in G'.

Now, assume PossibleSchwartzWinner$(G, A) = \Omega$. Let G'' be a majority graph extending G' (and, a fortiori, G). Such a G'' exists (because G' is asymmetric). By construction of G', there is a path in G' from A to every node of PossibleSchwartzWinner$(G,\ A)\backslash\{A\}$, hence to every node of $\Omega \setminus \{A\}$; since G'' extends G', this holds a fortiori for G'', hence A is a Schwartz winner in G'' and therefore a possible Schwartz winner for G.

Conversely, assume PossibleSchwartzWinner$(G, A) = \Sigma \neq \Omega$. Denote $\Theta = \Omega \backslash \Sigma$. Then, for all $(Y, Z) \in \Sigma \times \Theta$ we have $Z \to Y \in G'$. Now, $Z \in \Theta$ means that no edge $Z \to Y$ (for $Z \in \Theta$ and $Y \in \Sigma$) was added to G'; hence, for every $Y \in \Sigma$ and $Z \in \Theta$, we have that $Z \to Y \in G'$ if and only if $Z \to Y \in G$. This implies that for all $(Y, Z) \in \Sigma \times \Theta$ we have $Z \to Y \in G$, therefore, in every majority graph G'' extending G, every candidate of Θ beats every candidate of Σ, and in particular A. Therefore, there cannot be a path in G'' from A to a candidate in Z, which implies that A is not a Schwartz winner in G''. Since the latter holds for every majority graph G'' extending G, A is not a possible Schwartz winner for G. □

Algorithm 2 is in fact a polynomial-time algorithm for computing a superset of possible Schwartz winners for incomplete profiles.

Corollary 2 *Given an incomplete majority graph G, the set Poss$(G,\ r_{Schwartz})$ can be computed in polynomial time.*

Proof: It follows from Theorem 7 together with the fact that Algorithm PossibleSchwartzWinner (Algorithm 2) runs in polynomial time, since it performs at most m iterations, each iteration considering at most m^2 pairs of candidates. □

By Lemma 2 of [7], the set of the Smith winners[7] for an incomplete tournament G coincides with the set of possible Schwartz winners for G. Therefore, the above polynomiality result is also subsumed by Theorem 5 of [7].

An equivalent characterization of possible Schwartz winners says that $Poss(G, r_{Schwartz})$ is the smallest subset Θ of G satisfying this condition: for every $Z \in \Theta$ and every $X \in \Omega \setminus \Theta$, $(Z, X) \in G$. The proof of this result can be found in [23].

To conclude, we have shown that it is easy to compute a superset of the possible Schwartz winners from the incomplete profile by considering the majority graph. Regarding the quality of this heuristic, similar results can be derived as in Section 4.2 since Example 11 can be easily generalized to work for every agenda.

[7] The Smith winners are the candidates that belong to the Smith set, which is the set of candidates that dominate any candidate not in the set [7].

6 Computing winners over balanced agendas

We now turn our attention to fair winners, that is, candidates that win in some balanced agendas. We study the computational complexity of determining whether there is completion of the incomplete preferences such that a given candidate can win in a balanced tree, i.e., whether he is a possible fair winner.

We recall that the candidates that win in every balanced agenda coincide with Condorcet winners, thus the candidates that win in every balanced agenda for some completion (respectively, all completions) of the preferences coincide with possible (respectively, necessary) Condorcet winners, and thus they are polynomial to find both for weighted and unweighted profiles, and for majority graphs.

Given an incomplete weighted profile P, since every balanced agenda is also an agenda, we have that every possible fair winner is also a possible Schwartz winner. We already know from Theorem 5 that determining whether a given candidate is in *Poss*$(P, r_{Schwartz})$ is intractable. We will now show that this remains so also when we require that the agenda is balanced.

Theorem 8 *Given a fixed number m of candidates with $m \geq 3$ and an incomplete weighted profile P, deciding if a candidate is a possible fair winner for P is NP-complete.*

Proof: We use the same construction as in the proof of Theorem 5. Given the profile constructed there, the only possible fair agendas in which B wins are those in which A plays against C, and (at some later point) B then plays against the winner. All the additional candidates will be defeated by A, B, and C so can be placed anywhere in the fair agenda. □

The computational complexity of determining whether a given candidate is the necessary fair winner remains an open question, just as the complexity of computing necessary Schwartz winners. However, we know that, since every balanced agenda is also an agenda, every necessary fair winner is also a necessary Schwartz winner.

7 Related work

Since we have studied a number of different issues (fair winners with complete profiles, possible and necessary winners for voting trees and Schwartz, from incomplete profiles and from incomplete majority graphs), the related work spans several areas of research. We therefore structure this section according to the kind of topic studied. There are three related streams of work: voting rules with incomplete preferences, computational aspects of voting trees, and computing winners from incomplete tournaments.

7.1 Computing voting rules with incomplete preferences

We have given several results about voting trees and the Schwartz rule when voters' preferences are incomplete. The notions of possible and necessary winners for voting rules, as well as the notions of possible and necessary Condorcet winners, were

introduced by Konczak and Lang [21]. The computational complexity of computing possible and necessary winners for many common voting rules was thoroughly studied by Xia and Conitzer [38], by Betzler and Dorn [5], and by Baumeister and Rothe [2]. We complete the picture by adding results concerning voting trees and Schwartz. Moreover, we consider both the standard case of incomplete profiles and the case of incomplete majority graphs, which leads to a superset of the sets of possible winners which is easier to compute. The complexity of determining whether a given candidate is a possible or necessary winner has also been considered in Pini *et al.* [27] for the STV rule. Moreover, they give a preference elicitation procedure which focuses just on the set of possible winners. The problem of deciding if preference elicitation is over (that is, the problem of determining if the remaining votes can be cast so that a given candidate does not win) has been previously investigated also by Conitzer *et al.* [10, 11, 36] .

7.2 Voting trees

We have considered the computational complexity of determining possible winners for voting trees. A special case of the possible winner problem is when a subset of voters provide complete votes and the other voters provide empty votes. In such situations, determining whether a given candidate is a possible winner is equivalent to deciding whether there is a constructive coalitional manipulation for this candidate. Conitzer *et al.* [11] have investigated the complexity of weighted coalitional manipulation for $V(T)$ and found out that the latter problem is polynomial. We prove that the more general possible winner problem for $V(T)$ is NP-complete for weighted voters and 3 candidates or more.

Vu *et al.* [35] consider the following problem. The input consists of a set of candidates and a collection of pairwise winning probabilities (i, j) representing the probability that candidate i beats candidates j in a pairwise election; the problem (called "tournament schedule control") consists of designing a voting tree satisfying some given conditions (such as balancedness) that maximizes the winning probability of a target candidate. It is shown ([35], Theorems 1 and 3) that finding a balanced tree maximizing the probability that a given candidate wins is NP-hard, even if the pairwise winning probabilities are all in $\{0, 0.5, 1\}$. Our fair winners correspond to winners for which the maximum winning probability over all balanced trees is 1, when the probability matrix is composed only of 0's and 1's (for any two candidates, we know with certainty which of them beats the other in a pairwise comparison). Whether the NP-hardness result of [35] still holds when winning probabilities are $0/1$, that is, whether computing fair winners is NP-hard, is still an open problem. However, Vassilevska Williams [37] has shown that, if we allow only some of the matches between players, then the problem of determining if a candidate is a fair winner is NP-complete. Moreover, she has shown that if a candidate satisfies certain conditions, they are a fair winner and it is possible to construct in polynomial time a balanced tree in which this candidate wins. For example, if a candidate A has a number of outgoing edges in the majority graph G that is greater than or equal to the number of outgoing edges of every candidate B that beats A in G, then A is a fair winner and it is possible to compute in polynomial time a binomial spanning arborescence of G rooted in A [37].

We have given some slightly different complexity results to [35] and [39] for certain problems about voting trees with incomplete preferences. The differences concern mainly the nature of the incomplete preferences. In [35], the input is a probability matrix specifying, for every pair of candidates x, y, how likely it is that x beats y. In Sections 4.2 and 5.2, we start from an incomplete majority graph, therefore the 'qualitative' counterpart of [35]. In [39] and in Sections 4.1 and 5.1, the input is a collection of partial orders over the set of candidates.

Hazon *et al.* [20] also address the maximization over all voting trees of the probability that a given candidate wins the election, and show that a modified version of the problem of designing a tree such that the probability that a given candidate wins be greater than or equal to a certain value is NP-hard.

7.3 Incomplete tournaments

Brandt *et al.* [7] generalize several tournament solutions to incomplete tournaments, and study their computational complexity. Among the concepts they study there are the Smith set [7] and the Schwartz set (which are known to coincide for complete tournaments). Therefore, their input is an incomplete tournament, just as we study in Section 3.2. By Lemma 2 of [7] and Definition 15, the set of the Smith winners for an incomplete tournament G defined as in [7] coincides with the set of our possible Schwartz winners for G. Therefore, our polynomiality result (Corollary 2) is subsumed by their Theorem 5. Note that these two works have been developed independently (both conference papers were published in 2007) with very different starting points and motivations.

Since the Schwartz set is contained in the Smith set even for incomplete tournaments (Theorem 1 in [7]), we may wonder whether it coincides with the set of necessary Schwartz winners. However, this does not hold; for example, consider the graph G with four candidates A, B, C, D, and the edges $A >_m B$, $B >_m C$, and $C >_m A$: A is in the Smith set for G, but from our Theorem 6, A is not a necessary Schwartz winner.

8 Conclusions and future work

We have considered multi-agent settings where agents' preferences may be incomplete and are aggregated using voting trees. We have addressed various computational issues of determining winners for complete and incomplete profiles, as well as majority graphs.

All our complexity results about determining the various types of winners are summarized in Table 1: in the first four rows we consider the voting tree rule associated with a fixed tree ($V(T)$) with incomplete preferences, as well as the case of balanced trees. The following three rows concern the Schwartz rule and the fair winner determination when the agents' preferences are incomplete.

We first considered voting trees when the tree is fixed, i.e., $V(T)$, when the agents' preferences are incomplete (rows 1, 2, 3, and 4 of Table 1). In this context, we have shown that, if we assume we know the agents' preferences in the form of incomplete

	incomplete weighted profiles (number of candidates)	incomplete majority graphs
Possible winner for $V(T)$	NP-complete (≥ 3) (Th. 2)	P (Th. 4)
Necessary winner for $V(T)$	coNP-complete (≥ 4) (Th. 3)	P (Th. 4)
Possible winner for $V(T)$, T balanced	NP-complete (≥ 3) (Th. 2)	P(Th. 4)
Necessary winner for $V(T)$, T balanced	?	P(Th. 4)
Necessary Schwartz winner	?	P (Cor. 1)
Possible Schwartz winner	NP-complete (≥ 3) (Th. 5)	P (Cor. 2)
Possible fair winner	NP-complete (≥ 3) (Th. 8)	?

Table 1: Complexity of winner determination. All the polynomial results for $V(T)$ hold for simple voting trees.

majority graphs, it is polynomial to compute both the necessary and the possible winners. However, if we assume that agents' preferences are expressed by an incomplete weighted profile, then possible and necessary winners are computationally intractable to compute. We have also shown that possible winners are intractable to compute even when the agenda is balanced. The same results have been shown also in [39] in the case of incomplete unweighted profiles with an unbounded number of candidates.

We also evaluated experimentally the quality of our heuristic to compute possible winners, showing that it performs very well in every case.

We then analyzed the Schwartz rule in the case of incomplete preferences (rows 5 and 6 of Table 1). When the preferences are expressed via an incomplete majority graph, it is easy to find possible and necessary Schwartz winners, i.e., those candidates that are Schwartz winners in some completion or in all the completions of the preferences, respectively. On the other hard, when the agents' preferences are expressed via an incomplete weighted profile, it is intractable to determine if a given candidate is a possible Schwartz winner. These two results imply that every candidate has an incentive to participate in the election which is desirable. If it were easy to find the set of possible Schwartz winners, then a candidate that was not in this set, i.e., that is a loser, might choose not to participate.

We also have defined fair winners, i.e., candidates that win in at least one balanced agenda. The problem of determining whether a given candidate is a fair winner for a given complete profile is an open problem that we intend to investigate in the future. We have also analyzed the complexity of determining whether a given candidate is a possible fair winner when the agents' votes are incomplete and weighted, and we have shown that this problem is intractable (row 7 of Table 1).

The computational complexity of determining necessary Schwartz winners is related to the problem of testing whether constructive control is possible. This problem is polynomial-time solvable for some classes of incomplete profiles [29]. Nevertheless, it remains open in general and needs to be studied further when we consider incomplete profiles. However, we have shown that it is polynomial-time solvable if we consider incomplete majority graphs.

An interesting direction for future work is deciding which candidates are most likely to win, which is related to probabilistic approaches to voting theory. We also

plan to study other forms of uncertainty in the application of the voting rule, such as uncertain weights in a scoring rule and uncertainty about the voting rule.

Acknowledgements

We would like to thank the referees for their helpful comments. This work has been partially supported by the MIUR PRIN 20089M932N project "Innovative and multidisciplinary approaches for constraint and preference reasoning". Jérôme Lang has been supported by the ANR project ComSoc (ANR-09-BLAN-0305). Toby Walsh is supported by the Australian Government's Department of Broadband, Communications and the Digital Economy and the Australian Research Council.

References

[1] J. S. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 1(4):295–306, 1985.

[2] D. Baumeister and J. Rothe. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. In *Proceedings of ECAI'10*, pages 1019–1020, Lisbon, Portugal, 2010.

[3] N. Betzler and B. Dorn. Towards a dichotomy of finding possible winners in elections Based on Scoring Rules. In *Proceedings of MFCS'09*, volume 5734 of *LNCS*, pages 124–136, Novy Smokovec, High Tatras, Slovakia, 2009.

[4] N. Betzler, S. Hemmann, and R. Niedermeier. A multivariate complexity analysis of determining possible winners given incomplete votes. In *Proceedings of IJCAI'09*, pages 53–58, Pasadena, California, USA, 2009.

[5] Nadja Betzler and Britta 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.

[6] F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. In *Proceedings of TARK'07*, pages 82–91, Brussels, Belgium, 2007.

[7] F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. *Mathematical Logic Quarterly*, 55(4):444–459, 2009.

[8] Y. Chevaleyre, J. Lang, N. Maudet, and J. Monnot. Possible winners when new candidates are added: the case of scoring rules. In *Proceedings of AAAI'10*, Atlanta, Georgia, USA, 2010.

[9] V. Conitzer and T. Sandholm. Complexity of manipulating an election with few candidates. In *Proceedings of AAAI'02*, pages 314–319, Edmonton, Alberta, Canada, 2002.

[10] V. Conitzer and T. Sandholm. Vote elicitation: complexity and strategy-proofness. In *Proceedings of AAAI'02*, pages 392–397, Edmonton, Alberta, Canada, 2002.

[11] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate. *Journal of the ACM*, 54(3), 2007.

[12] A. H. Copeland. A reasonable social welfare function. *University of Michigan Seminar on Applications of Mathematics to the Social Sciences*, 1951.

[13] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2002.

[14] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Llull and Copeland voting broadly resist bribery and control. In *Proceedings of AAAI'07*, pages 724–730, Vancouver, Canada, 2007.

[15] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research (JAIR)*, 35:275–341, 2009.

[16] F. A. Fischer, A. D. Procaccia, and A. Samorodnitsky. A new perspective on implementation by voting trees. In *Proceedings of EC'09*, pages 31–40, Stanford, California, 2009.

[17] F. A. Fischer, A. D. Procaccia, and A. Samorodnitsky. A new perspective on implementation by voting trees. *Random Structures and Algorithms*, forthcoming.

[18] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W.H. Freeman, 1979.

[19] N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. Evaluation of election outcomes under uncertainty. In *Proceedings of AAMAS'08, Volume 2*, pages 959–966, Estoril, Portugal, 2008.

[20] N. Hazon, P. E. Dunne, S. Kraus, and M. Wooldridge. How to rig elections and competitions. In *Proceedings of COMSOC'08*, Liverpool, UK, 2008.

[21] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedings of IJCAI'05 Multidisciplinary Workshop on Advances in Preference Handling*, Edinburgh, Scotland, UK, 2005.

[22] G. Laffond, J.-F. Laslier, and M. Le Breton. Condorcet choice correspondences: A set-theoretical comparison. *Mathematical Social Sciences*, pages 23–35, 1995.

[23] J. Lang, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Winner determination in sequential majority voting. In *Proceedings of IJCAI'07*, pages 1372–1377, Hyderabad, India, 2007.

[24] J.-F. Laslier. *Tournament solutions and majority voting*. Springer-Verlag, 1997.

[25] N. Miller. A new solution set for tournaments and majority voting: further graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, 24:68–9, 1980.

[26] H. Moulin. *Axioms of Cooperative Decision Making*. Cambridge University Press, 1988.

[27] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Incompleteness and incomparability in preference aggregation. In *Proceedings of IJCAI'07*, pages 1464–1469, Hyderabad, India, 2007.

[28] 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 *Proceedings of KR'08*, pages 571–578, Sydney, Australia, 2008.

[29] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Possible and necessary winners in voting trees: majority graphs vs. profiles. In *Proceedings of AAMAS'11*, Taipei, Taiwan, 2011.

[30] A. D. Procaccia, A. Zohar, Y. Peleg, and J. S. Rosenschein. Learning voting trees. In *Proceedings of AAAI'07*, pages 110–115, Vancouver, British Columbia, 2007.

[31] A. D. Procaccia, A. Zohar, Y. Peleg, and J. S. Rosenschein. The learnability of voting rules. *Artificial Intelligence*, 173(12-13):1133–1149, 2009.

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

[33] P. Slater. Inconsistencies in a schedule of paired comparisons. *Biometrika*, 48(3-4):303–312, 1961.

[34] M. Trick. Small binary voting trees. In *Proceedings of COMSOC'06*, pages 500–511, Amsterdam, Netherlands, 2006.

[35] T. Vu, A. Altman, and Y. Shoham. On the complexity of schedule control problems for knockout tournaments. In *Proceedings of AAMAS'09, Volume 1*, pages 225–232, Budapest, Hungary, 2009.

[36] T. Walsh. Complexity of terminating preference elicitation. In *Proceedings of AAMAS'08*, pages 967–974, Estoril, Portugal, 2008.

[37] V. Vassilevska Williams. Fixing a tournament. In *Proceedings of AAAI'10*, Atlanta, Georgia, USA, 2010.

[38] L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proceedings of AAAI'08*, pages 196–201, Chicago, Illinois, USA, 2008.

[39] L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. A longer unpublished version of [38], available at *http://www.cs.duke.edu/∼lxia*, 2010.