The Possible Winner Problem with Uncertain Weights[1]

Dorothea Baumeister, Magnus Roos, Jörg Rothe, Lena Schend, and Lirong Xia

Abstract

The original possible winner problem is: Given an unweighted election with partial preferences and a distinguished candidate *c*, can the preferences be extended to total ones such that *c* wins? We introduce a novel variant of this problem in which not some of the voters' preferences are uncertain but some of their weights. Not much has been known previously about the weighted possible winner problem. We present a general framework to study this problem, both for integer and rational weights, with and without upper bounds on the total weight to be distributed, and with and without ranges to choose the weights from. We study the complexity of these problems for important voting systems such as scoring rules, Copeland, ranked pairs, plurality with runoff, and (simplified) Bucklin and fallback voting.

1 INTRODUCTION

Much of the previous work in computational social choice has focused on the complexity of manipulation, control, and bribery problems in voting (see the surveys by Faliszewski et al. [18, 21]). More recently, many papers studied the possible winner problem, which generalizes the (unweighted) coalitional manipulation problem. The original possible winner problem was introduced by Konczak and Lang [24]. The input to this problem is an election with partial (instead of total) preferences and a distinguished candidate, and the question is whether it is possible to extend the partial preferences to total ones such that the distinguished candidate wins. Xia and Conitzer [28] studied this and also the necessary winner problem. Betzler and Dorn [7] and Baumeister and Rothe [5] established a dichotomy result for the possible winner problem, and Betzler et al. [8, 6] investigated the parameterized complexity of this problem.

A number of variants of the possible winner problem have been studied as well. Bachrach, Betzler, and Faliszewski [1] investigated a probabilistic variant thereof. Chevaleyre et al. [10] introduced the *possible winner with respect to the addition of new alternatives* problem, which is related to, yet different from the problem of control via adding candidates[2](see [2, 23]) and is also similar, yet not identical to the cloning problem in elections [16]. Their variant was further studied by, e.g., Xia, Lang, and Monnot [29] and Baumeister, Roos, and Rothe [4]. The latter paper in particular considered a weighted variant of the possible winner problem, and it also introduced and studied this problem under voting rule uncertainty, an approach that was followed up recently by Elkind and Erdélyi [14] who applied it to coalitional manipulation [11]. Baumeister et al. [3] studied variants of the possible winner problem with truncated ballots. Lang et al. [25] and Pini et al. [27] investigated the possible and necessary winner problem for voting trees and multi-round election systems such as STV. Most of the papers listed above consider only *unweighted* elections. We present a general framework to study the *weighted* possible winner problem, and we focus on elections where not some of the voters' preferences, but some of their *weights*, are uncertain. The problems we study in our framework come with integer or rational weights, with or without upper bounds on the total weight to be assigned, and with or without given ranges to choose the weights from. An interesting point in this regard is that while the original possible winner problem generalizes the coalitional manipulation problem [11], certain variants of the possible winner problem with uncertain weights generalize constructive control by adding/deleting voters [2, 23].

[1] To appear in *Proceedings of the 20th European Conference on Artificial Intelligence*, Montpellier, France, August 2012.

[2] We use *candidate* and *alternative* synonymously.

The following situation may motivate why it is interesting to study the possible winner problem with uncertain weights. Imagine a company that is going to decide on its future strategy by voting at the annual general assembly of stockholders. Among the parties involved, everybody's preferences are common knowledge. However, who will succeed with its preferred alternative for the future company strategy depends on the stockholders' weights, i.e., on how many stocks they each own, and there is uncertainty about these weights. Is it possible to assign weights to the parties involved (e.g., by them buying new stocks) such that a given alternative wins? As another example, suppose we want to decide which university is the best in the world based on different criteria (e.g., graduation and retention rates, faculty resources, student selectivity, etc.). Each criterion can be seen as a voter who gives a ranking over all universities (candidates). Suppose the voting rule is fixed (e.g., plurality), but the chair can determine the weights of these criteria. It is interesting to know whether a given university can win if the chair chooses the weights carefully.

2 PRELIMINARIES

An *election* is a pair (C,V) consisting of a finite set C of candidates and a finite list V of voters that are represented by their preferences over the candidates in C and are occasionally denoted by $v_1,\dots,v_{|V|}$. A voting system $\mathscr{E}$ is a set of rules determining the winning candidates according to the preferences in V. The voting systems considered here are all preference-based, that is, the votes are given as linear orders over C. For example, if $C=\{a,b,c,d\}$ then a vote $a>c>b>d$ means that this voter (strictly) prefers a to c, c to b, and b to d. If such an order is not total (e.g., when a voter only specifies $a>c>d$ as her preference over these four candidates), we say it is a partial order. For winner determination in weighted voting systems, a vote v of weight w is considered as if there were w unweighted (i.e., unit-weight) votes v.

For a given election (C,V), the *weighted majority graph (WMG)* is defined as a directed graph whose vertices are the candidates, and we have an edge $c\to d$ of weight $N(c,d)$ between any two vertices c and d, where $N(c,d)$ is the number of voters preferring c to d minus the number of voters preferring d to c. Note that in the WMG of any election, all weights on the edges have the same parity (and whether it is odd or even depends on the parity of the number of votes), and $N(c,d)=-N(d,c)$ (which is why it is enough to give only one of these two edges explicitly).

We will consider the following voting rules.

- **Positional Scoring Rules:** These rules are defined by a *scoring vector* $\vec{\alpha}=(\alpha_1,\alpha_2,\dots,\alpha_m)$, where m is the number of candidates, the α_i are nonnegative integers, and $\alpha_1\geq\alpha_2\geq\cdots\geq\alpha_m$. Let $\rho_i(c)$ denote the position of candidate c in voter v_i's vote. Then c receives $\alpha_{\rho_i(c)}$ points from v_i, and the total score of c is $\sum_{i=1}^{n}\alpha_{\rho_i(c)}$ for n voters. All candidates with the largest score are the $\vec{\alpha}$ winners. In particular, we will consider k-approval elections, $k\leq m$, whose scoring vector has a 1 in the first k positions, and the remaining $m-k$ entries are all 0. The special case of 1-approval is also known as *plurality* and that of $(m-1)$-approval as *veto*. The scoring vector $(m-1,m-2,\dots,2,1,0)$ defines the *Borda* rule.

- **Copeland$^\alpha$ (for each rational number α, $0\leq\alpha\leq 1$):**[3] For any two alternatives c and c', we can simulate a *pairwise election* between them, by seeing how many voters prefer c to c', and how many prefer c' to c; the winner of the pairwise election is the one preferred more often. Then, an alternative receives one point for each win in a pairwise election, α points for each tie, and zero points for each loss. This is the Copeland score of the alternative. A Copeland winner maximizes the Copeland score.

- **Ranked pairs:** This rule first creates an entire ranking of all the candidates. In each step, we consider a pair of candidates c,c' that we have not previously considered; specifically, we

[3]The original Copeland system [12] is defined for the specific value of $\alpha=1/2$; the generalization to other α values is due to Faliszewski et al. [20].

choose among the remaining pairs one with the highest $N(c,c')$ value (in case of ties, we use some tie-breaking mechanism) and then fix the order $c > c'$, unless this contradicts previous orders already fixed (i.e., unless this order violates transitivity). We continue until we have considered all pairs of candidates (and so we have a full ranking). A candidate at the top of the ranking for some tie-breaking mechanism is a winner.

- **Plurality with runoff:** This rule proceeds in two rounds. First, all alternatives except those two with the highest plurality score are eliminated; in the second round (the runoff), the plurality rule is used to select a winner among these two. Some tie-breaking rule is applied in both rounds if needed.
- **Bucklin and fallback voting (both simplified):** In a Bucklin election, the voters' preferences are linear orders and the level ℓ score of a candidate c is the number of voters ranking c among their top ℓ positions. The Bucklin score of a candidate c is the smallest number t such that more than half of the voters rank c somewhere in their top t positions. A Bucklin winner minimizes the Bucklin score.[4] In (simplified) fallback elections, on the other hand, nontotal (more specifically, "top-truncated" as defined in [3]) preference orders are allowed. Every Bucklin winner is also a fallback winner, but if no Bucklin winner exists (which may happen due to the voters' partial orders) and ℓ is the length of a longest preference order among the votes, all candidates with the greatest level ℓ score are the fallback winners. Throughout this paper we will refer to "simplified Bucklin" and "simplified fallback" simply as Bucklin and fallback voting.

We will use the following notation. If the set of candidates is, say, $C = B \cup D \cup \{c\}$, then we mean by $c > \overrightarrow{D} > \cdots$ that c is preferred to all candidates, where $\overrightarrow{D}$ is an arbitrarily fixed ordering of the candidates occurring in D, and "$\cdots$" indicates that the remaining candidates (those from B in this example) can be ranked in an arbitrary order afterwards.

Some proofs in this paper use *McGarvey's trick* [26] (applied to WMGs), which constructs a list of votes whose WMG is the same as some targeted weighted directed graph. This will be helpful because when we present our proofs, we only need to specify the WMG instead of the whole list of votes, and then by using McGarvey's trick for WMGs, a votes list can be constructed in polynomial time. More specifically, McGarvey showed that for every unweighted majority graph, there is a particular list of preferences that results in this majority graph. Extending this to WMGs, the trick works as follows. For any pair of candidates, (c,d), if we add two votes, $c > d > c_3 > \cdots > c_m$ and $c_m > c_{m-1} > \cdots > c_3 > c > d$, to a vote list, then in the WMG, the weight on the edge $c \to d$ is increased by 2 and the weight on the edge $d \to c$ is decreased by 2, while the weights on all other edges remain unchanged.

3 PROBLEM DEFINITIONS AND DISCUSSION

We now define our variants of the possible winner problem with uncertain weights. Let $\mathscr{E}$ be a given voting system and $\mathbb{F} \in \{\mathbb{Q}^+, \mathbb{N}\}$.

$\mathscr{E}$-Possible-Winner-with-Uncertain-Weights-$\mathbb{F}$ ($\mathscr{E}$-PWUW-$\mathbb{F}$)	
Given:	An $\mathscr{E}$ election $(C, V_0 \cup V_1)$, $V_0 \cap V_1 = \emptyset$, where the weights of the voters in V_0 are not specified yet and weight zero is allowed for them, yet all voters in V_1 have weight one, and a designated candidate $c \in C$.
Question:	Is there an assignment of weights $w_i \in \mathbb{F}$ to the votes v_i in V_0 such that c is an $\mathscr{E}$ winner of election $(C, V_0 \cup V_1)$ when v_i's weight is w_i for $1 \leq i \leq \lvert V_0 \rvert$?

[4] We consider only this simplified version of Bucklin voting. In the full version (see, e.g., [17]), among all candidates with smallest Bucklin score, say t, for c to win it is also required that c's level t score is largest.

We distinguish between allowing nonnegative rational weights (i.e., weights in $\mathbb{Q}^+$) and nonnegative integer weights (i.e., weights in $\mathbb{N}$). In particular, we allow weight-zero voters in V_0. Note that for inputs with $V_0 = \emptyset$ (which is not excluded in the problem definition), we obtain the ordinary unweighted (i.e., unit-weight) winner problem for $\mathscr{E}$. Allowing weight zero for voters in V_0 in some sense corresponds to control by deleting voters (see [2, 23]); later in this section we also briefly discuss the relationship with control by adding voters. The reason why we distinguish between votes with uncertain weights and unit-weight votes in our problem instances is that we want to capture these problems in their full generality; just as votes with total preferences are allowed to occur in the instances of the original possible winner problem. The requirement of normalizing the weights in V_1 to unit-weight, on the other hand, *is* a restriction (that doesn't hurt) and is chosen at will. This will somewhat simplify our proofs.

We also consider the following restrictions of $\mathscr{E}$-PWUW-$\mathbb{F}$:

- In $\mathscr{E}$-PWUW-RW-$\mathbb{F}$, an $\mathscr{E}$-PWUW-$\mathbb{F}$ instance and regions (i.e., intervals) $R_i \subseteq \mathbb{F}$, $1 \leq i \leq |V_0|$, are given, and the question is the same as in $\mathscr{E}$-PWUW-$\mathbb{F}$, except that each weight w_i must be chosen from R_i in addition.
- In $\mathscr{E}$-PWUW-BW-$\mathbb{F}$, an $\mathscr{E}$-PWUW-$\mathbb{F}$ instance and a positive bound $B \in \mathbb{F}$ is given, and the question is the same as in $\mathscr{E}$-PWUW-$\mathbb{F}$, except that $\sum_{i=1}^{|V_0|} w_i \leq B$ must hold in addition (i.e., the total weight that can be assigned must be bounded by B).
- In $\mathscr{E}$-PWUW-BW-RW-$\mathbb{F}$, an $\mathscr{E}$-PWUW-BW-$\mathbb{F}$ instance and regions (i.e., intervals) $R_i \subseteq \mathbb{F}$, $1 \leq i \leq |V_0|$, are given, and the question is the same as in $\mathscr{E}$-PWUW-BW-$\mathbb{F}$, except that each weight w_i must be chosen from R_i in addition.

One could also define other variants of $\mathscr{E}$-PWUW-$\mathbb{F}$ (e.g., the *destructive* variant where the question is whether c's victory can be prevented by some weight assignment) or other variants of $\mathscr{E}$-PWUW-BW-RW-$\mathbb{F}$ and $\mathscr{E}$-PWUW-RW-$\mathbb{F}$ (e.g., by allowing *sets of intervals* for each weight), but here we focus on the eight problems defined above. We focus on the *winner* model (aka. the *co-winner* or the *nonunique-winner* model) where the question is whether c can be made *a* winner by assigning appropriate weights. By minor proof adjustments, most of our results can be shown to also hold in the *unique-winner* model where we ask whether c can be made the only winner.

We assume that the reader is familiar with common complexity-theoretic notions, such as the complexity classes P and NP, and the notions of hardness and completeness with respect to the polynomial-time many-one reducibility, which we denote by $\leq_m^p$.

The following reductions hold trivially among our problems, by setting the bound on the total weight allowed to the sum of the highest possible weights for the first two reductions and by setting the intervals to $[0, B]$ (where B is the bound on the total weight) for the last two reductions:

$$
\begin{aligned}
\text{PWUW-RW-}\mathbb{Q}^+ &\leq_m^p \text{PWUW-BW-RW-}\mathbb{Q}^+ && (1)\\
\text{PWUW-RW-}\mathbb{N} &\leq_m^p \text{PWUW-BW-RW-}\mathbb{N} && (2)\\
\text{PWUW-BW-}\mathbb{Q}^+ &\leq_m^p \text{PWUW-BW-RW-}\mathbb{Q}^+ && (3)\\
\text{PWUW-BW-}\mathbb{N} &\leq_m^p \text{PWUW-BW-RW-}\mathbb{N}. && (4)
\end{aligned}
$$

Related to our variants of the PWUW problem is the problem of constructive control by adding voters (see [2]), CCAV for short. Here, a set C of candidates with a distinguished candidate $c \in C$, a list V of registered voters, an additional list V' of as yet unregistered voters, and a positive integer k are given. The question is whether it is possible to make c win the election by adding at most k voters from V' to the election.

Obviously, there is a direct polynomial-time many-one reduction from CCAV to PWUW-BW-RW-$\mathbb{N}$. The voters in V_1 are the registered voters from V and the voters in V_0 are those from V', where the weights can be chosen from $\{0, 1\}$ for all votes in V_0, and the total bound on the weight B is set

PWUW-	Scoring Rules, Plurality with runoff	Plurality, 2-AV, Veto	3-AV	k-AV, $k \geq 4$	Bucklin, Fallback	Copeland, Ranked Pairs
$\mathbb{Q}^+$	P	P	P	P	P	?
$\mathbb{N}$	?	P	P	P	NP-c.	NP-c.
BW-RW-$\mathbb{Q}^+$	P	P	P	P	P	?
BW-RW-$\mathbb{N}$	?	P	?	NP-c.	NP-c.	NP-c.
BW-$\mathbb{Q}^+$	P	P	P	P	P	?
BW-$\mathbb{N}$	?	P	?	NP-c.	NP-c.	NP-c.
RW-$\mathbb{Q}^+$	P	P	P	P	P	?
RW-$\mathbb{N}$	?	P	P	P	NP-c.	NP-c.

Table 1: Overview of results. "NP-c." stands for NP-complete.

to k. If succinct representation is assumed,[5] there is also a polynomial-time many-one reduction in the other direction. The registered voters are those from V_1, and the unregistered voters are those from V_0, where each vote is added according to its maximal weight in the PWUW instance. The number k of voters who may be added equals the bound B on the total weight.

Since there are reductions in both directions, complexity results carry over from CCAV to PWUW-BW-RW-$\mathbb{N}$ when we assume succinct representation. For the voting systems considered in this paper, this implies that PWUW-BW-RW-$\mathbb{N}$ is NP-complete for Copeland^0 and Copeland^1, and is solvable in polynomial time for plurality (see [20, 2]). (Note that the NP-hardness results on CCAV for Bucklin and fallback voting from [17] concern the full, not the simplified versions of these voting rules.) These already known cases are nevertheless covered by our proofs in the next section, since they handle several restrictions of the PWUW problems at the same time. Conversely, the results from the next section for PWUW-BW-RW-$\mathbb{N}$ all carry over to CCAV if we assume succinct representation.

4 RESULTS AND SELECTED PROOFS

Table 1 gives an overview of our results. In the next section, we will provide or sketch some of the proofs for these results. Due to space constraints, not all proofs can be presented in full detail.

4.1 Integer Weights

We begin with the results for the integer cases.

Proposition 1 *1. Each of the four variants of* plurality-PWUW-$\mathbb{N}$, veto-PWUW-$\mathbb{N}$, *and* 2-approval-PWUW-$\mathbb{N}$ *studied in this paper is in* P.

2. *For each* $k \geq 1$, k-approval-PWUW-$\mathbb{N}$ *and* k-approval-PWUW-RW-$\mathbb{N}$ *are in* P.

PROOF. For the first statement, we present the proof details for 2-approval-PWUW-BW-RW-$\mathbb{N}$, where for each vote in V_0 the range of allowed weights is $\{0,1\}$. The proof can be adjusted to also work when other ranges are given.

Given a 2-approval-PWUW-BW-RW-$\mathbb{N}$ instance as above, we construct the following max-flow instance. Let V_0' denote the list of votes in V_0 where c is ranked among the top two positions. We may assume, without loss of generality, that the given bound B on the total weight satisfies $B \leq |V_0'|$.[6] The vertices are $\{s,s',t\} \cup V_0' \cup (C \setminus \{c\})$ with the following edges:

[5] This means that when there are several identical votes, we don't list them all but rather store a number in binary saying how often this vote occurs.

[6] Otherwise, the optimal strategy is to let the weights of the votes in V_0' be 1 and to let the weights of all other votes be 0.

- There is an edge $s \rightarrow s'$ with capacity B and an edge from s' to each node in V_0' with capacity 1.
- There is an edge from a node L in V_0' to a node d in $C \setminus \{c\}$ with capacity 1 if and only if d is ranked besides c among the top two positions in L.
- There is an edge from each node $d \in C \setminus \{c\}$ to t with capacity $B + score(c,V_1) - score(d,V_1)$, where $score(e,V_1)$ is the 2-approval score of any $e \in C$ in vote list V_1.[7]

In the max-flow problem, we are asked whether there exists a flow whose value is B. We note that in the PWUW instance, it is always optimal to choose B votes in V_0' and to let their weights be 1. The bound on $d \rightarrow t$ for $d \in C \setminus \{c\}$ ensures that the 2-approval score of d is no more than the 2-approval score of c.

The claims for 2-approval-PWUW-RW-$\mathbb{N}$ and 2-approval-PWUW-BW-$\mathbb{N}$ follow from (2) and (4).

For the second statement, it suffices to maximize the weights of the votes in V_0' that rank c among their top k positions, and to minimize the weights of the other votes.

❑

In particular, it is open whether 3-approval-PWUW-BW-RW-$\mathbb{N}$ and 3-approval-PWUW-BW-$\mathbb{N}$ are also in P. For $k \geq 4$, however, we can show that these problems are NP-complete.

Theorem 2 *For each $k \geq 4$, k-approval-PWUW-BW-RW-$\mathbb{N}$ and k-approval-PWUW-BW-$\mathbb{N}$ are NP-complete.*

PROOF. It is easy to see that both problems belong to NP. For proving NP-hardness, we give a proof for 4-approval-PWUW-BW-$\mathbb{N}$ by a reduction from the NP-complete problem EXACT COVER BY 3-SETS (X3C): Given a set $\mathscr{B} = \{b_1, \ldots, b_{3q}\}$ and a collection $\mathscr{S} = \{S_1, \ldots, S_n\}$ with $|S_i| = 3$ and $S_i \subseteq \mathscr{B}$, $1 \leq i \leq n$, does $\mathscr{S}$ contain an exact cover for $\mathscr{B}$, i.e., a subcollection $\mathscr{S}' \subseteq \mathscr{S}$ such that every element of $\mathscr{B}$ occurs in exactly one member of $\mathscr{S}'$?

Construct an instance of k-approval-PWUW-BW-$\mathbb{N}$ with the set

$$C = \{c, b_1, \ldots, b_{3q}, b_1^1, \ldots, b_{3q}^1, b_1^2, \ldots, b_{3q}^2, b_1^3, \ldots, b_{3q}^3\}$$

of candidates, where c is the designated candidate, and with the set V_0 of n votes of the form $c > \overrightarrow{S_i} > \cdots$, the set V_1 of $q-1$ votes of the form $b_j > b_j^1 > b_j^2 > b_j^3 > \cdots$ for each j, $1 \leq j \leq 3q$, and the bound $B = q$ on the total weight of the votes in V_0. Recall that the votes in V_1 all have fixed weight one, and those of the votes in V_0 are from $\mathbb{N}$. We show that $\mathscr{S}$ has an exact cover for $\mathscr{B}$ if and only if we can set the weights of the voters in this election such that c is a winner.

Assume that there is an exact cover $\mathscr{S}' \subseteq \mathscr{S}$ for $\mathscr{B}$. By setting the weights of the votes $c > \overrightarrow{S_i} > \cdots$ to one for those q subsets S_i contained in $\mathscr{S}'$, and to zero for all other votes in V_0, c is a winner of the election, as c and all b_j, $1 \leq j \leq 3q$, receive exactly q points, whereas b_j^1, b_j^2, and b_j^3, $1 \leq j \leq 3q$, receive $q-1$ points each.

Conversely, assume that c can be made a winner of the election by choosing the weights of the votes in V_0 appropriately. Note that the bound on the total weight for the votes in V_0 is $B = q$. Every b_i gets $q-1$ points from the votes in V_1, and c gets points only from the votes in V_0. Since there are always some b_j getting points if a vote from V_0 has weight one, there are at least three b_j having q points if a vote from V_0 has weight one. Hence c must get q points from the votes in V_0 by setting the weight of q votes to one. Furthermore, every b_j can occur only once in the votes having weight one in V_0, as otherwise c would not win. Thus, the S_i corresponding to the votes of weight one in V_0 must form an exact cover for $\mathscr{B}$.

[7] Note that if this capacity is negative, the given 2-approval-PWUW-BW-RW-$\mathbb{N}$ instance is trivially a no-instance, since c can never be made a winner.

By adding dummy candidates to fill the positions receiving points, we can adapt this proof for k-approval for any fixed $k > 4$. NP-hardness for k-approval-PWUW-BW-RW-$\mathbb{N}$, $k \geq 4$, then follows from the trivial reduction (4) stated in Section 3. ❑

We now show that all variants of PWUW with integer weights are NP-complete for Copeland$^{\alpha}$, ranked pairs, Bucklin, and fallback elections.

Theorem 3 *For each rational number α, $0 \leq \alpha \leq 1$, every variant of* Copeland$^{\alpha}$-PWUW-$\mathbb{N}$ *studied in this paper is* NP-*complete.*

PROOF. NP membership is easy to see for all problem variants. We first prove NP-hardness for Copeland$^{\alpha}$-PWUW-$\mathbb{N}$, and then show how to modify the proof for the variants of the problem. Given an X3C instance $(\mathscr{B}, \mathscr{S})$ with $\mathscr{B} = \{b_1, \dots, b_{3q}\}$ and $\mathscr{S} = \{S_1, \dots, S_n\}$, we construct the following PWUW instance for Copeland$^{\alpha}$, where the set of candidates is $\mathscr{B} \cup \{c, d, e\}$. Without loss of generality we assume that $q \geq 4$ and we are asked whether c can be made a winner.

The votes on C are defined as follows. V_0 will encode the X3C instance and V_1 will be used to implement McGarvey's trick. V_0 consists of the following n votes: For each j, $1 \leq j \leq n$, there is a vote $d > e > \overrightarrow{S_j} > c > \cdots$. V_1 is the vote list whose WMG has the following edges:

- $c \rightarrow d$ with weight $q+1$, $d \rightarrow e$ with weight $q+1$, and $e \rightarrow c$ with weight $q+1$.
- For every i, $1 \leq i \leq 3q$, $d \rightarrow b_i$ and $e \rightarrow b_i$ each with weight $q+1$, and $b_i \rightarrow c$ with weight $q-3$.
- The weight on any other edge not defined above is no more than 1.

It follows that no matter what the weights of the votes in V_0 are, d beats e and e beats c in pairwise elections, and both d and e beat all candidates in $\mathscr{B}$ in pairwise elections. For c to be a winner, c must beat d in their pairwise election, which means that the total weight of the votes in V_0 is no more than q. On the other hand, c must beat all candidates in $\mathscr{B}$. This happens if and only if the votes in V_0 that have positive weights correspond to an exact cover of $\mathscr{B}$, and all of these votes must have weight one. This means that Copeland$^{\alpha}$-PWUW-$\mathbb{N}$ is NP-hard.

For the BW and BW-RW variants, we let $B = q$; for the RW and BW-RW variants, we let the range of each vote in V_0 be $\{0, 1\}$. ❑

Theorem 4 *All variants of* ranked-pairs-PWUW-$\mathbb{N}$ *studied in this paper are* NP-*complete.*

PROOF. The proof is similar to the proof of Theorem 3. That the problems are in NP is easy to see. For the hardness proof, given an X3C instance $(\mathscr{B}, \mathscr{S})$ with $\mathscr{B} = \{b_1, \dots, b_{3q}\}$ and $\mathscr{S} = \{S_1, \dots, S_n\}$, we construct the following ranked-pairs-PWUW-$\mathbb{N}$ instance, where the set of candidates is $\mathscr{B} \cup \{c, d\}$. We are asked whether c can be made a winner. V_0 consists of the following n votes: For each $j, 1 \leq j \leq n$, there is a vote $e > \overrightarrow{S_j} > c > d > \cdots$. V_1 is the vote list whose WMG has the following edges, and is constructed by applying McGarvey's trick:

- $c \rightarrow d$ with weight $2q+1$, $d \rightarrow e$ with weight $4q+1$, and $e \rightarrow c$ with weight $2q+1$.
- For every i, $1 \leq i \leq 3q$, $d \rightarrow b_i$ and $e \rightarrow b_i$ each with weight $2q+1$, and $b_i \rightarrow c$ with weight $4q-1$.
- The weight on any other edge not defined above is 1.

If the total weight of votes in V_0 is larger than q, then the weight on $e \rightarrow c$ and $e \rightarrow b_i$ in the WMG is at least $3q+2$, and the weight on $d \rightarrow e$ is no more than $3q$, which means that c is not a winner for ranked pairs. Moreover, if c is a winner, then the weight on any $b_i \rightarrow c$ should not be strictly higher than the weight on $c \rightarrow d$, otherwise $b_i \rightarrow c$ will be fixed in the final ranking. It follows that if c is a winner, then the votes in V_0 that have positive weights correspond to an exact cover of $\mathscr{B}$, and all of these votes must have weight one. This means that ranked-pairs-PWUW-$\mathbb{N}$ is NP-hard.

For the BW and BW-RW variants, we let $B = q$; for the RW and BW-RW variants, we let the range of each vote in V_0 be $\{0,1\}$. ❑

Theorem 5 *All variants of* Bucklin-PWUW-$\mathbb{N}$ *studied in this paper are* NP-*complete.*

PROOF. NP membership is easy to see for all problem variants. We first prove NP-hardness for Bucklin-PWUW-$\mathbb{N}$, and then show how to modify the proof for the variants of the problem. Given an X3C instance $(\mathscr{B},\mathscr{S})$ with $\mathscr{B} = \{b_1,\dots,b_{3q}\}$ and $\mathscr{S} = \{S_1,\dots,S_n\}$, we construct the following Bucklin-PWUW-$\mathbb{N}$ instance. The set of candidates is $\mathscr{B} \cup \{c,d\} \cup D \cup D'$, where $D = \{d_1,\dots,d_{3q}\}$ and $D' = \{d'_1,\dots,d'_{3q}\}$ are sets of auxiliary candidates. We are asked whether c can be made a winner. V_0 consists of the following n votes: For each j, $1 \le j \le n$, there is a vote $d > \overrightarrow{S_j} > c > \overrightarrow{D} > \overrightarrow{D'} > \cdots$. V_1 consists of $q-1$ copies of $\overrightarrow{\mathscr{B}} > c > \overrightarrow{D'} > \overrightarrow{D} > d$ and one copy of $\overrightarrow{D'} > c > \overrightarrow{\mathscr{B}} > d > \overrightarrow{D}$.

If the total weight of votes in V_0 is larger than q, then d is the unique candidate that is ranked in top positions for more than half of the votes, which means that c is not a winner. Suppose the total weight of the votes in V_0 is at most q. Then, the Bucklin score of c is $3q+1$ and the Bucklin score of any candidate in D and D' is larger than $3q+1$. Therefore, c is a Bucklin winner if and only if the Bucklin score of any candidate in $\mathscr{B}$ is at least $3q+1$. This happens if and only if the votes in V_0 that have positive weights correspond to an exact cover of $\mathscr{B}$, and all of these votes must have weight one. This means that Bucklin-PWUW-$\mathbb{N}$ is NP-hard.

For the BW and BW-RW variants, we let $B = q$; for the RW and BW-RW variants, we let the range of each vote in V_0 be $\{0,1\}$. ❑

Bucklin voting can be seen as the special case of fallback voting where all voters give complete linear orders over all candidates. So the NP-hardness results for Bucklin voting transfer to fallback voting, while the upper NP bounds are still easy to see.

Corollary 6 *All variants of* fallback-PWUW-$\mathbb{N}$ *studied in this paper are* NP-*complete.*

4.2 Rational Weights and Voting Systems that Can Be Represented by Linear Inequalities

Chamberlin and Cohen [9] observed that various voting rules can be represented by systems of linear inequalitites, see also [19]. We use this property to formulate linear programs, thus being able to solve the PWUW problem variants with rational weights for these voting rules efficiently, provided that the size of the systems describing the voting rules is polynomially bounded. Note that an LP with rational instead of integer values can be solved in polynomial time [22].

What voting rules does this technique apply to? The crucial requirement a voting rule needs to satisfy is that the scoring function used for winner determination can be described by linear inequalities and that this description is in a certain sense independent of the voters' weights. By "independent of the voters' weights" we mean that the points a candidate gains from a vote are determined essentially in the same way in both a weighted and an unweighted electorate, but in the former we have a weighted sum of these points that gives the candidate's score, whereas in the latter we have a plain sum. Scoring functions satisfying this condition are said to be *weight-independent*. This requirement is fulfilled by, e.g., the scoring functions of all scoring rules, Bucklin, and fallback voting. Copeland's scoring function, on the other hand, does not satisfy it. In a Copeland election, every candidate gets one point for each other candidate she beats in a pairwise contest. Who of the two candidates wins a pairwise contest and thus gains a Copeland point depends directly on the voters' weights. Thus, the Copeland score in a weighted election is not a weighted sum of the Copeland scores in the corresponding unweighted election in the above sense.

In what follows, we have elections where the voter list consists of the two sublists V_0 and V_1. We have to assign weights $x_1, \dots, x_{|V_0|}$ to the voters in V_0. We don't exclude the case where weight zero can be assigned, but we will seek to find solutions where all weights are strictly positive, since assigning weight zero to a voter is equivalent to excluding this voter entirely from the election. For $c \in C$, let $\rho_i^0(c)$ denote the position of c in the preference of the ith voter in V_0, $1 \leq i \leq |V_0|$, and let $\rho_j^1(c)$ denote the position of c in the preference of the jth voter in V_1, $1 \leq j \leq |V_1|$.

Lemma 7 *Let $\mathscr{E}$ be a voting rule with a weight-independent scoring function that can be described by a system A of polynomially many linear inequalities. Then* $\mathscr{E}$-PWUW-$\mathbb{Q}^+$, $\mathscr{E}$-PWUW-BW-$\mathbb{Q}^+$, $\mathscr{E}$-PWUW-RW-$\mathbb{Q}^+$, *and* $\mathscr{E}$-PWUW-BW-RW-$\mathbb{Q}^+$ *are each in* P.

PROOF. Let $x_1, x_2, \dots, x_n$ be the variables of the system A that describes $\mathscr{E}$ for an $\mathscr{E}$ election with n voters. The following linear program can be used to solve $\mathscr{E}$-PWUW-BW-RW-$\mathbb{Q}^+$. Let an instance of this problem be given: an election $(C, V_0 \cup V_1)$ with as yet unspecified weights in V_0, a designated candidate $c \in C$, a bound $B \in \mathbb{Q}^+$, and regions $R_i \subseteq \mathbb{Q}^+$, $1 \leq i \leq |V_0|$. The *vector of variables* of our linear program is $\vec{x} = (x_1, x_2, \dots, x_{|V_0|}, \chi) \in \mathbb{R}^{|V_0|+1}$ and we maximize the *objective function* $\vec{c} \cdot \vec{x}^T$ with $\vec{c} = (0, 0, \dots, 0, 1)$ and the following constraints:

$$A \tag{5}$$
$$x_i - \chi \geq 0 \qquad \text{for } 1 \leq i \leq |V_0| \tag{6}$$
$$\chi \geq 0 \tag{7}$$
$$\sum_{i=1}^{|V_0|} x_i \leq B \tag{8}$$
$$x_i \leq r_i \qquad \text{for } 1 \leq i \leq |V_0| \tag{9}$$
$$-x_i \leq -\ell_i \quad \text{for } 1 \leq i \leq |V_0| \tag{10}$$

Constraint (5) gives the linear inequalitites that have to be fulfilled for the designated candidate c to win under $\mathscr{E}$. By maximizing the additional variable χ in the objective function we try to find solutions where the weights are positive, this is accomplished by constraint (6). Constraint (8) implements our given upper bound B for the total weight to be assigned and constraints (9) and (10) implement our given ranges $R_i = [\ell_i, r_i] \subseteq \mathbb{Q}$ for each weight.

Omit (8) for $\mathscr{E}$-PWUW-RW-$\mathbb{Q}^+$, omit (9) and (10) for $\mathscr{E}$-PWUW-BW-$\mathbb{Q}^+$, and omit (8), (9), and (10) for $\mathscr{E}$-PWUW-$\mathbb{Q}^+$.

A solution in $\mathbb{Q}$ for a linear program with polynomially bounded constraints can be found in polynomial time. ❑

In the following theorems we present the specific systems of linear inequalities describing scoring rules in general, and the voting systems Bucklin, fallback, and plurality with runoff. These can be used to formally specify the complete linear program stated in the proof of Lemma 7.

Theorem 8 *For each scoring rule* $\vec{\alpha}$, $\vec{\alpha}$-PWUW-$\mathbb{Q}^+$, $\vec{\alpha}$-PWUW-BW-$\mathbb{Q}^+$, $\vec{\alpha}$-PWUW-RW-$\mathbb{Q}^+$, *and* $\vec{\alpha}$-PWUW-BW-RW-$\mathbb{Q}^+$ *are in* P.

PROOF. We are given an election with m different candidates in C, where $c \in C$ is the distinguished candidate. Recall that $\rho_i^0(c)$ denotes c's position in the preference of voter $v_i \in V_0$, and that $\alpha_{\rho_i^0(c)}$ denotes the number of points c gets for this position according to the scoring vector $\vec{\alpha}$. Let $S_{V_1}(c)$ denote the number of points candidate c gains from the voters in V_1 (recall that those have all weight one). Then the distinguished candidate c is a winner if and only if for all candidates $c' \in C$ with

$c' \neq c$, we have $\left(\left(\alpha_{\rho_j^0(c)} - \alpha_{\rho_j^0(c')}\right)_{1\leq j\leq |V_0|}\right)\vec{x}^T \geq S_{V_1}(c') - S_{V_1}(c)$, where $\vec{x} = (x_1, x_2, \ldots, x_{|V_0|}) \in \mathbb{R}^{|V_0|}$ are the weights that will be assigned to the voters in V_0. The linear program for scoring rule $\vec{\alpha}$ is of the following form. As in the proof of Lemma 7, we have the vector of variables $\vec{x} = (x_1, x_2, \ldots, x_{|V_0|}, \chi) \in \mathbb{R}^{|V_0|+1}$ and we maximize the objective function $\vec{c} \cdot \vec{x}^T$ with $\vec{c} = (0, 0, \ldots, 0, 1)$ and the following constraints:

$$-\sum_{i=1}^{|V_0|} \left(\alpha_{\rho_i^0(c)} - \alpha_{\rho_i^0(c')}\right) x_i \leq S_{V_1}(c) - S_{V_1}(c') \;\forall\, c' \neq c \tag{11}$$

$$x_i - \chi \geq 0 \qquad \text{for } 1 \leq i \leq |V_0| \tag{12}$$

$$\chi \geq 0 \tag{13}$$

$$\sum_{i=1}^{|V_0|} x_i \leq B \tag{14}$$

$$x_i \leq r_i \qquad \text{for } 1 \leq i \leq |V_0| \tag{15}$$

$$-x_i \leq -\ell_i \quad \text{for } 1 \leq i \leq |V_0| \tag{16}$$

Here again, constraints (14) to (16) are needed only for the restricted variants.

Since we have at most $(m-1)|V_0| + 3|V_0| + 2 = (m+2)|V_0| + 2$ constraints, this linear program can be solved in polynomial time. ❑

Note that by adding χ to the left-hand side of (11), a solution where χ is positive is an assignment of weights making the distinguished candidate a unique winner.

Being level-based voting rules, for Bucklin and fallback voting we have to slightly expand the presented approach. Due to space constraints, we omit the proof of Theorem 9 and only briefly sketch the idea. Intuitively, it is clear that we first try to make the distinguished candidate a level 1 winner; if this attempt fails, we try the second level; and so on. So the linear program in the proof of Theorem 9 has to be solved for each level beginning with the first until a solution has been found. For Bucklin voting, the representation by linear inequalities is due to Dorn and Schlotter [13], and we adapt it for the simplified version of Bucklin and fallback voting. For the latter, we add appropriate constraints if the approval stage is reached.

Theorem 9 *Let $\mathscr{E}$ be either Bucklin or fallback voting.* $\mathscr{E}$-PWUW-$\mathbb{Q}^+$, $\mathscr{E}$-PWUW-BW-$\mathbb{Q}^+$, $\mathscr{E}$-PWUW-RW-$\mathbb{Q}^+$, *and* $\mathscr{E}$-PWUW-BW-RW-$\mathbb{Q}^+$ *are each in* P.

Note that the proof of Theorem 9 does not work in the unique-winner case.

For plurality with runoff we can take a similar approach: For each candidate d different from c, we use a set of linear inequalities to figure out whether there exists a set of weights such that (1) c and d enter the runoff (i.e., the plurality scores of c and d are at least the plurality score of any other candidate), and (2) c beats d in their pairwise election. Therefore, we have the following corollary whose proof does not work in the unique-winner case.

Theorem 10 *Let PR be the plurality with runoff rule.* PR-PWUW-$\mathbb{Q}^+$, PR-PWUW-BW-$\mathbb{Q}^+$, PR-PWUW-RW-$\mathbb{Q}^+$, *and* PR-PWUW-BW-RW-$\mathbb{Q}^+$ *are each in* P.

PROOF. For each candidate d different from c, there exists a set of linear inequalities that are similar to those in the proof of Theorem 8 such that c and d enter the runoff if and only if these inequalities can be satisfied. We also add the following inequality: $\sum_{\{i \mid c >_{\rho_{0,i}} d\}} x_i + |\{k \mid c >_{\rho_{1,k}} d\}| \geq \sum_{\{i \mid d >_{\rho_{0,i}} c\}} x_i + |\{k \mid d >_{\rho_{1,k}} c\}|$, where $\{i \mid c >_{\rho_{j,i}} d\}$ denotes those voters $v_i \in V_j$ for $j \in \{0,1\}$ that prefer c to d. Then, for each candidate d different from c we construct an LP that is similar to the LP in the proof of Theorem 8. It follows that c is a possible winner if and only if at least one of these LPs has a feasible solution. ❑

5 CONCLUSIONS AND OPEN QUESTIONS

We introduced the possible winner problem with uncertain weights, where not the preferences but the weights of the voters are uncertain, and we studied this problem and its variants in a general framework. We showed that some of these problem variants are easy to solve and some are hard to solve for some of the most important voting rules. Interestingly, while the original possible winner problem (in which there is uncertainty about the voters' preferences) generalizes the coalitional manipulation problem and is a special case of swap bribery [15], the possible winner problem with uncertain weights generalizes the problem of constructive control by adding or deleting voters.

Some interesting issues remain open, as indicated in Table 1, e.g., regarding 3-approval, Copeland voting, positional scoring rules, and plurality with runoff. Also, it would be interesting to study an even more general variant: the weighted possible winner problem with uncertainty about both the voters' preferences and their weights.

Acknowledgments: This work was supported in part by DFG grant RO-1202/15-1, SFF grant "Cooperative Normsetting" of HHU Düsseldorf, ARC-DP110101792, and a DAAD grant for a PPP-PROCOPE project. Lirong Xia is supported by NSF under Grant #1136996 to the Computing Research Association for the CIFellows Project.

References

[1] Y. Bachrach, N. Betzler, and P. Faliszewski. Probabilistic possible winner determination. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 697–702. AAAI Press, July 2010.

[2] J. Bartholdi III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modelling*, 16(8/9):27–40, 1992.

[3] D. Baumeister, P. Faliszewski, J. Lang, and J. Rothe. Campaigns for lazy voters: Truncated ballots. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 577–584. IFAAMAS, June 2012.

[4] D. Baumeister, M. Roos, and J. Rothe. Computational complexity of two variants of the possible winner problem. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 853–860. IFAAMAS, May 2011.

[5] 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.

[6] N. Betzler. On problem kernels for possible winner determination under the k-approval protocol. In *Proceedings of the 35th International Symposium on Mathematical Foundations of Computer Science*, pages 114–125. Springer-Verlag *Lecture Notes in Computer Science #6281*, August 2010.

[7] 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.

[8] N. Betzler, S. Hemmann, and R. Niedermeier. A multivariate complexity analysis of determining possible winners given incomplete votes. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 53–58. IJCAI, July 2009.

[9] J. Chamberlin and M. Cohen. A linear inequality method of establishing certain social choice conjectures. *Public Choice*, 33(2):5–16, 1978.

[10] 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.

[11] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 2007.

[12] A. Copeland. A "reasonable" social welfare function. Mimeographed notes from a Seminar on Applications of Mathematics to the Social Sciences, University of Michigan, 1951.

[13] B. Dorn and I. Schlotter. Multivariate complexity analysis of swap bribery. *Algorithmica*, 64(1):126–151, 2012.

[14] E. Elkind and G. Erdélyi. Manipulation under voting rule uncertainty. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 627–634. IFAAMAS, June 2012.

[15] E. Elkind, P. Faliszewski, and A. Slinko. Swap bribery. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory*, pages 299–310. Springer-Verlag *Lecture Notes in Computer Science #5814*, October 2009.

[16] E. Elkind, P. Faliszewski, and A. Slinko. Cloning in elections. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, pages 768–773. AAAI Press, July 2010.

[17] G. Erdélyi, L. Piras, and J. Rothe. The complexity of voter partition in Bucklin and fallback voting: Solving three open problems. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 837–844. IFAAMAS, May 2011.

[18] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*, 53(11):74–82, 2010.

[19] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Multimode control attacks on elections. *Journal of Artificial Intelligence Research*, 40:305–351, 2011.

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

[21] P. Faliszewski and A. Procaccia. AI's war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.

[22] L. Hačijan. A polynomial algorithm in linear programming. *Soviet Math. Dokl.*, 20:191–194, 1979.

[23] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6):255–285, 2007.

[24] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, pages 124–129, July/August 2005.

[25] J. Lang, M. Pini, F. Rossi, D. Salvagnin, K. Venable, and T. Walsh. Winner determination in voting trees with incomplete preferences and weighted votes. *Journal of Autonomous Agents and Multi-Agent Systems*, pages 130–157, 2012.

[26] D. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.

[27] M. Pini, F. Rossi, K. Venable, and T. Walsh. Incompleteness and incomparability in preference aggregation: Complexity results. *Artificial Intelligence*, 175(7):1272–1289, 2011.

[28] L. Xia and V. Conitzer. Determining possible and necessary winners given partial orders. *Journal of Artificial Intelligence Research*, 41:25–67, 2011.

[29] L. Xia, J. Lang, and J. Monnot. Possible winners when new alternatives join: New results coming up! In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 829–836. IFAAMAS, May 2011.

Dorothea Baumeister, Magnus Roos, Jörg Rothe, and Lena Schend
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany
Email: `{baumeister,roos,rothe,schend}@cs.uni-duesseldorf.de`

Lirong Xia
Center for Research on Computation and Society (CRCS)
Harvard University
Cambridge, MA 02138, USA
Email: `lxia@seas.harvard.edu`