A Multivariate Complexity Analysis of Determining Possible Winners Given Incomplete Votes

Nadja Betzler
Institut für Informatik
Friedrich-Schiller-Universität Jena
D-07743 Jena, Germany
nadja.betzler@uni-jena.de

Susanne Hemmann
Institut für Informatik
FSU Jena
D-07743 Jena, Germany

Rolf Niedermeier
Institut für Informatik
Friedrich-Schiller-Universität Jena
D-07743 Jena, Germany
rolf.niedermeier@uni-jena.de

Abstract

The POSSIBLE WINNER problem asks whether some distinguished candidate may become the winner of an election when the given incomplete votes are extended into complete ones in a favorable way. POSSIBLE WINNER is NP-complete for common voting rules such as Borda, many other positional scoring rules, Bucklin, Copeland etc. We investigate how three different parameterizations influence the computational complexity of POSSIBLE WINNER for a number of voting rules. We show fixed-parameter tractability results with respect to the parameter "number of candidates" but intractability results with respect to the parameter "number of votes". Finally, we derive fixed-parameter tractability results with respect to the parameter "total number of undetermined candidate pairs" and identify an interesting polynomial-time solvable special case for Borda.

1 Introduction

Voting plays a key role for decision making in modern societies and multiagent systems. Classically, a vote one-to-one corresponds to a linear order of the given set of candidates. In many situations, however, only partial information about the voters' preferences is available. For instance, in multiagent systems the agents often must make a joint decision based on their individual preferences concerning a (potentially large) set of alternatives (synonymously, candidates). Sometimes, however, an agent has not enough information or time to provide a linear order on all alternatives and thus only may provide a partial ordering of the alternatives. Then, a vote corresponds to a partial order of the set of candidates. In this context, two central questions arise.

1. Given a set of partial orders, does a distinguished candidate c win for *each* extension of the partial orders into linear ones? This is the NECESSARY WINNER problem.
2. Given a set of partial orders, can a distinguished candidate c win for *at least one* extension of the partial orders into linear ones? This is the POSSIBLE WINNER problem.

In this paper, we focus on the POSSIBLE WINNER problem for several voting systems—in particular, we consider the common voting rules k-approval and, following Xia and Conitzer [2008], Borda, Bucklin, Copeland, maximin, and ranked pairs. Some of our results also hold for broader classes of voting rules.

POSSIBLE WINNER is a well-studied problem [Konczak and Lang, 2005; Lang *et al.*, 2007; Pini *et al.*, 2007; Walsh, 2007; Xia and Conitzer, 2008]. Correcting Konczak and Lang [2005] claiming polynomial-time solvability for scoring rules, it has been recently shown that POSSIBLE WINNER is NP-complete for Borda, many other scoring rules, Bucklin, Copeland, maximin, and ranked pairs [Xia and Conitzer, 2008]. In contrast, NECESSARY WINNER is coNP-complete only for Copeland and ranked pairs but polynomial-time solvable for Bucklin, maximin, and scoring rules [Xia and Conitzer, 2008]. Interestingly, the NP-completeness results for POSSIBLE WINNER even hold when there is only a constant number of candidate pairs per vote for which the relative order is open. In contrast, if the number of candidates is constant, then POSSIBLE WINNER can be solved in polynomial time provided that the voting rule can be executed in polynomial time [Walsh, 2007]. An extensively studied special case of POSSIBLE WINNER is "constructive manipulation" (see, e.g., [Conitzer *et al.*, 2007]). Here, the given set of partial orders consists of two subsets; one subset contains linearly ordered votes and the other one completely unordered votes. Refer to [Walsh, 2007; Xia and Conitzer, 2008] for nice overviews.

The fundamental goal of parameterized algorithmics [Niedermeier, 2006] is to find out whether the seemingly unavoidable combinatorial explosion occurring in algorithms to solve NP-hard problems can be confined to certain problem-specific parameters. The idea is that when such a parameter has only small values in applications, then an algorithm with a running time that is exponential exclusively with respect to the parameter may be efficient and practical. Formally, a given parameterized problem (I, p) with input instance I and parameter p is called *fixed-parameter tractable (FPT)* with respect to the parameter p if it can be solved within running time $f(p) \cdot |I|^{O(1)}$ for some computable function f. In other words, the running time is polynomial for constant parameter values and the degree of the polynomial in the running time is independent of the parameter.

The statement that POSSIBLE WINNER can be solved in polynomial time when the number of candidates is bounded by a constant k [Walsh, 2007] is a starting point for our work.

The decisive question is how k influences the degree of the polynomial that upper-bounds the running time. In the stated polynomial-time algorithm, k determines the degree in the way "n^k", where n denotes the input size. Such algorithms become impractical even for modest values of k. We partially answer this question by showing that POSSIBLE WINNER is fixed-parameter tractable with respect to the parameter k (number of candidates) for all of the above mentioned voting rules. Moreover, we initiate a broader algorithmic study of POSSIBLE WINNER by pursuing a multivariate complexity analysis based on further parameterizations. More specifically, we additionally analyze how the parameters "number of votes" and "total number of undetermined pairs" influence the problem complexity. We achieve (parameterized) intractability results in the first case and fixed-parameter tractability results in the second case. More precisely, first we prove that POSSIBLE WINNER for Borda, Bucklin, and k-approval remains NP-complete even in case of a constant number of votes (which excludes any hope for fixed-parameter tractability with respect to this parameter). Second, we prove that with respect to the parameter "total number of undetermined pairs" POSSIBLE WINNER becomes fixed-parameter tractable for all voting rules where in case of linearly ordered votes the winner can be determined in polynomial time. For the special case Borda, we further identify an interesting special case that is solvable in polynomial time and that can be used to improve the running time of the fixed-parameter algorithm.

Due to the lack of space, several details had to be deferred to a full version of the paper.

Preliminaries. Let $C = \{c_1, \ldots, c_m\}$ be the set of *candidates*. Classically, a vote is a linear order (a transitive, antisymmetric, and total relation) on C. An n-voter profile P on C consists of n votes on C. The set of all profiles on C is denoted by $P(C)$.

A *voting rule* r is a function from $P(C)$ to 2^C; the resulting image set is called the set of *co-winners*. Throughout this paper we cast our proofs for the case of uniquely determined co-winners (that is, one-element co-winner sets), just referring to these as winners. All our results hold for the co-winner case as well.

The most frequently used voting rules in this work are Borda and k-approval. Both are *(positional) scoring rules*. Scoring rules are defined by scoring vectors $(\alpha_0, \alpha_1, \ldots, \alpha_{m-1})$ with integers $\alpha_0 \geq \alpha_1 \geq \cdots \geq \alpha_{m-1} \geq 0$. More specifically, a scoring rule r consists of a sequence of scoring vectors $s_1, s_2, \ldots$ such that for any $i \in \mathbb{N}$ there is a scoring vector for i candidates. For a vote v and a candidate $c \in C$, let the score $s(v, c) := \alpha_j$ where j is the number of candidates that are better than c in v. For any profile $P = \{v_1, \ldots, v_n\}$, let $s(P, c) := \sum_{i=1}^{n} s(v_i, c)$. A scoring rule will output c as a co-winner if it maximizes $s(P, c)$. For example, the following three voting rules with their corresponding scoring vectors are common:
- *plurality* with $(1, 0, \ldots)$,
- *Borda* with $(m-1, m-2, \ldots, 0)$, and
- k*-approval* with $(1, 1, 1, \ldots, 0, 0)$ starting with k 1's.

Due to the lack of space, we refer to [Xia and Conitzer, 2008] for the definitions and an overview of the other voting systems studied in this work.

	MANIPULATION	POSSIBLE WINNER
Copeland	NP-c [1]	NP-c
Maximin	NP-c [2]	NP-c
Ranked Pairs	NP-c [2]	NP-c
Bucklin	P [2]	**NP-c**
k-approval	P	**NP-c**
Borda	?	**NP-c**

Table 1: Complexity of MANIPULATION for a constant-size coalition and POSSIBLE WINNER for a constant number of partial votes. Boldface results are new. The polynomial-time solvability of MANIPULATION for k-approval is easy to observe. The other results are from [1][Faliszewski *et al., 2008*] (for specific tie-breaking) and [2][Xia and Conitzer, 2008]. The NP-completeness of POSSIBLE WINNER directly follows from the NP-completeness of MANIPULATION.

A partial order is a transitive, antisymmetric, and reflexive relation. Sometimes, a partial order specifies a whole subset of candidates, e.g., $e \succ D$. This notation means that $e \succ d$ for all $d \in D$ and there is no specified order among the candidates in D. A linear order V *extends* a partial order O if $O \subseteq V$, that is, for any $i, j \leq m$, one has $c_i \succ_O c_j \Rightarrow c_i \succ_V c_j$. For a voting rule r and a profile of partial orders $P_O = (O_1, \ldots, O_n)$ on C, a candidate $c \in C$ is a *possible winner* if there exists a $P = (V_1, \ldots, V_n)$ such that, for each i, V_i extends O_i into a linear order and $c \in r(P)$.

POSSIBLE WINNER
Input: A voting rule r, a set of candidates C, a profile of partial orders $P_O = (O_1, \ldots, O_n)$ on C, and a distinguished candidate $c \in C$.
Question: Is there an extension profile $P = (V_1, \ldots, V_n)$ where V_i extends O_i for all $1 \leq i \leq n$ and $c \in r(P)$?

2 Number of candidates

To assess the parameterized complexity with respect to the parameter "number of candidates", we employ Lenstra's famous algorithm for bounded-variable-cardinality integer linear programming (see [Niedermeier, 2006, Chapter 11]) which has shown usefulness for proving fixed-parameter tractability for control in elections as well [Faliszewski *et al.*, 2007]. Lenstra's result says that it is fixed-parameter tractable with respect to the number of variables to check whether all inequalities of an integer linear program can be fulfilled at the same time (the so-called feasibility problem). By expressing POSSIBLE WINNER as the feasibility problem for an integer linear program with a number of variables bounded by a function solely depending on the parameter "number of candidates", one can obtain the following.

Theorem 1. *For all positional scoring rules, Bucklin, Copeland, ranked pairs, and maximin* POSSIBLE WINNER *is fixed-parameter tractable with respect to the parameter "number of candidates".*

3 Number of votes

A well-studied scenario in voting systems is constructive manipulation with a bounded number of unspecified votes, that is, bounded coalition size [Faliszewski *et al.*, 2008; Xia *et al.*, 2008]. Since constructive manipulation is a special case of POSSIBLE WINNER, this motivates the consideration of the parameter "number of partial votes". An overview of results for this parameterization is given in Table 1. In the following, we show the NP-completeness of POSSIBLE WINNER for a constant *total* number of votes for k-approval, Bucklin, and Borda.

k-approval. The NP-complete INDEPENDENT SET (IS) problem asks, given an undirected graph $G = (V, E)$ and a positive integer t, whether there is a size-t vertex subset $V' \subseteq V$ such that there is no edge between any two vertices of V'.

Theorem 2. *For k-approval,* POSSIBLE WINNER *is NP-complete for a partial profile that consists of two partial orders when k is part of the input.*

Proof. (Sketch) We give a many-one reduction from INDEPENDENT SET to POSSIBLE WINNER for k-approval. Given an IS-instance $((V, E), t)$ with $V = \{v_1, \ldots, v_n\}$, we construct a 2-voter partial profile P_O over a set of candidates C in which the distinguished candidate $c \in C$ is a possible winner according to k-approval with $k := |n + \binom{n}{2} + |E| - tn + \binom{t}{2} + 1|$ iff (G, t) is a yes-instance. The set of candidates is $C := C_V \uplus C_E \uplus \{c\} \uplus D$ with a candidate for every vertex, that is, $C_V := \{c_i \mid v_i \in V\}$, and candidates that are related to pairs of vertices, that is, for $1 \leq i < j \leq n$, if $\{v_i, v_j\} \in E$, then there are two candidates e_{ij} and e'_{ij} in C_E; otherwise, there is one candidate e_{ij} in C_E. Further, D consists of $|d|$ dummy candidates with $d := n - t + \binom{n}{2} + |E| - tn + \binom{t}{2}$. In the following, we describe the construction for the case that $d \geq 0$. The case $d < 0$ can be handled similarly. The two partial orders O_1 and O_2 of P_O are specified as follows:

O_1 : $c \succ D \succ C_V \succ C_E$.
O_2 : $c \succ C_V \cup C_E \succ D$, and for $1 \leq i < j \leq n$,
if $\{v_i, v_j\} \in E$, then $c_i \succ e_{ij}$ and $c_j \succ e'_{ij}$;
otherwise, $c_i \succ e_{ij}$ and $c_j \succ e_{ij}$.

The distinguished candidate c has a total score of two. If there is an extension in which all other candidates have a total score of at most one, then c is a possible winner. This is equivalent to the demand that every other candidate must assume at least one "zero-position". Since we have $1 + |C_V| + |C_E| + |D| = 1 + n + \binom{n}{2} + |E| + n - t + \binom{n}{2} + |E| - tn + \binom{t}{2}$ candidates and consider k-approval with $k = n + \binom{n}{2} + |E| - tn + \binom{t}{2} + 1$, there are $\binom{n}{2} + |E| + n - t$ zero-positions per vote.

In O_1, the candidates of C_E are beaten by all other candidates, and thus assume zero-positions in every extension. Hence, $|C_E| = \binom{n}{2} + |E|$ zero-positions are already occupied. The remaining $n - t$ zero-positions can only be assigned to candidates of C_V. In an extension in which c wins, those t candidates from C_V with a one-position in O_1 must assume a zero-position in O_2. According to the definition of O_2, every candidate from C_V is placed before $n - 1$ candidates from C_E. Hence, placing a candidate from C_V at a zero-position in O_2 implies that $n - 1$ candidates from C_E must also assume zero-positions. Since every candidate of C_E has a zero-position in O_1, assigning zero-positions to candidates from C_E in an extension of O_2 is not necessary to make c the possible winner. The basic idea of the construction is that two candidates $c_i, c_j \in C_V$ that assume zero-positions in O_2 can "share" a candidate from C_E if $\{v_i, v_j\} \notin E$. In this case, both candidates enforce a zero-position for the same candidate e_{ij} whereas otherwise the two candidates e_{ij} and e'_{ij} must assume zero-positions. As we argue in the following, the construction ensures that for no pair of "selected" candidates from C_V there are two corresponding candidates from C_E that are forced to a zero-position. Otherwise, the number of remaining zero-positions would not be sufficient for the selected candidates from C_V and c would not win.

Since all candidates from D assume zero-positions in O_2, it is easy to verify that there are exactly $tn - \binom{t}{2}$ zero-positions left over (and c beats all candidates of D). Now, suppose that the candidates from C_V that must assume a zero-position in O_2 correspond to an independent set. Then, the number of candidates of C_E that are forced to assume a zero-position is $t \cdot (n - 1)$ minus the number of candidates that are counted twice. For all pairs of candidates that correspond to non-adjacent vertices a candidate of C_E is counted twice. Since the t candidates correspond to an independent set, there are $\binom{t}{2}$ such candidates and the number of zero-positions is sufficient. Hence, extending O_2 by placing the C_V-candidates that correspond to an independent set (and all "enforced" candidates) to zero-positions and extending O_1 by placing the remaining candidates of C_V to zero-positions yields an extension in which c is the winner.

We omit the argument for the reverse direction. □

Bucklin. The *Bucklin score* of a candidate is the smallest position s such that the candidate takes position s or smaller in more than half of the votes. The candidate with the smallest Bucklin score wins. The reduction from IS used in Theorem 2 for k-approval can be adapted to work for Bucklin. We sketch the basic idea: Additionally to the two partial orders, the modified profile contains three linear orders ensuring the following. The Bucklin score of the distinguished candidate is set to a certain value s. All other candidates have Bucklin score at most s if they do not assume a position greater than s in one of the partial votes, and have Bucklin score higher than s, otherwise. Now, a position higher than s is equivalent to a "zero-position", that is a position greater than k, in the proof of Theorem 2. Hence, one can argue in analogy to there. Regarding the construction of the linear votes, the Bucklin score of the distinguished candidate can be fixed by setting it to position s in all three linear votes. All other candidates are set to a position smaller than s in one of the linear orders and to a position higher than s in the two other linear orders. This idea can be used to show the following.

Theorem 3. *For Bucklin,* POSSIBLE WINNER *is NP-complete for a partial profile consisting of two partial and three linear orders.*

Borda. The 3-PARTITION problem is defined as follows. Given a multi-set $A = \{a_1, \ldots, a_n\}$ of positive integers and $B := (3/n) \cdot \sum_{a_i \in A} a_i$, it asks whether there is a partition of A into size-3 subsets $A_1, \ldots, A_{n/3}$ such that $\sum_{a_i \in A_j} a_i = B$ for all $j \in \{1, \ldots, n/3\}$. The 3-PARTITION problem is strongly NP-complete [Garey and Johnson, 1979].

Informally, this means that the NP-hardness still holds when the integers of A have values polynomially bounded in n. We denote the special case that each integer $a_i \in A$ must be a multiple of n as 3-n-PARTITION. It is not hard to verify that the 3-n-PARTITION problem is strongly NP-hard since every 3-PARTITION instance can be reduced to 3-n-PARTITION by multiplying all input integers with n.

Theorem 4. *For Borda,* POSSIBLE WINNER *is NP-complete for a partial profile consisting of three partial and three linear orders.*

Proof. Let $A = \{a_1, \ldots, a_n\}$ denote a 3-n-PARTITION instance with $B := (3/n) \cdot \sum_{a_i \in A} a_i$. To ease the presentation, we assume that $a_i < a_{i+1}$ for $i = 1, \ldots, n-1$ and $a_1 = n$. It is not hard to modify the following many-one reduction to work for general instances. We construct a partial profile P_O over a set C of candidates in which the distinguished candidate $c \in C$ can become a winner iff A is a yes-instance for 3-n-PARTITION. The set of candidates is $C := \{c\} \uplus E \uplus T \uplus D$, with one candidate for every member of A, that is, $E := \{e_i \mid a_i \in A\}$, candidates representing the subsets resulting from the partition into 3-sets, that is $T := \{t_1, \ldots, t_{n/3}\}$, and a set of dummy candidates $D := \biguplus_{i=1}^{n} D_i$ only needed to "fill" positions (specified later). The partial profile P_O consists of three linear orders and three identical partial orders. Every partial order $O_q, q \in \{1, 2, 3\}$, of P_O is given by $c \succ T$ and

$$c \succ D_1 \succ e_1 \succ D_2 \succ \cdots \succ D_i \succ e_i \succ \cdots \succ D_n \succ e_n$$

with $|D_1| = a_1 - 1$ and $|D_i| = a_i - a_{i-1} - 1$ for $i \in \{2, \ldots, n\}$. This definition fixes the number of dummy candidates; more precisely, $|D| = \sum_{D_i \in D} |D_i| = a_1 - 1 + \sum_{i=2}^{n}(a_i - a_{i-1} - 1) = a_n - n$. Thus, the total number of candidates is $m = 1+|E|+|T|+|D| = 1+n+n/3+a_n-n = a_n+n/3+1$. Since 3-$n$-PARTITION is *strongly* NP-complete, we can assume that a_n and, thus, m is polynomial in n. This also allows that the integers from A are presented by the candidates as follows. For every candidate $e_i \in E$, there are exactly $a_i - 1$ candidates $s \in D \cup E$ with $c \succ s$ and $s \succ e_i$ in $O_q, q \in \{1, 2, 3\}$. Further, note that the position and thus the total score of the distinguished candidate c is already fixed. In contrast, every subset candidate $t_j \in T$ can be "inserted" at any position behind c in the three partial votes. The basic idea of this construction is that the "choice" of the positions for t_j in the three partial orders corresponds to the choice of three numbers from A into the corresponding subset A_j. For example, inserting t_j directly before the candidate e_i in one of the partial votes means that $a_i \in A_j$. More specifically, we would like to ensure the following two points for every possible extension in which c wins:

1. Every number of A is *selected* exactly once, that is, for every candidate $e_i \in E \backslash \{e_1\}$ there is exactly one candidate $t_j \in T$ with $e_{i-1} \succ t_j$ and $t_j \succ e_i$ in one of the three partial votes, and one candidate $t_j \in T$ with $t_j \succ e_1$.
2. For all $t_j \in T$, the sum corresponding to the three "number candidates" from E selected by t_j is B.

The two points can be realized by setting the linear orders of P_O appropriately. To this end, for $c' \in C \backslash \{c\}$ let the *maximum partial score* $s_p^{\max}(c')$ denote the maximum score that c' may get within the three partial votes without beating c. Then, it is not hard to construct three linear orders of P_O such that the following conditions hold. (We defer the description of the linear orders to the full version of this work.)

- $s_p^{\max}(e_i) = 3(m-1) - 3a_i - i$ for all $e_i \in E$,
- $s_p^{\max}(t_j) = 3(m-1) - B$ for all $t_j \in T$, and
- $s_p^{\max}(d) \geq 3m$ for all $d \in D$.

By construction, c will make $3(m-1)$ points in any extension. Hence, the last condition implies that a candidate $d \in D$ can never beat c. Further, every $e_i \in E$ must "loose" at least $3a_i + i$ points against c and every $t_j \in T$ must loose at least B points against c.

Now, we show that there is a solution for 3-n-PARTITION iff there is an extension of P_O such that c wins.

"$\Rightarrow$" Let $\{A_1, \ldots, A_{n/3}\}$ with $A_j = \{a_{j_1}, a_{j_2}, a_{j_3}\}$ denote a solution of 3-n-PARTITION for A. Then, extend $O_q, q \in \{1, 2, 3\}$, such that $t_j \succ e_{j_q}$ and $D_{j_q} \succ t_j$. This extension is unambiguous since all $a_{j_q} \in A$ are pairwise distinct. As explained before, in every partial vote, for every $e_i \in E$, there are exactly $a_i - 1$ candidates $s \in D \cup E$ with $c \succ s$ and $s \succ e_i$. Thus, without inserting any $t_j \in T$ before e_i, e_i "looses" $3a_i$ points against c. For $q \in \{1, 2, 3\}$, let $\tau_{i,q}$ denote the number of candidates from T that are inserted before e_i to extend O_q. Then, for every e_i, we have $\tau_{i,1} + \tau_{i,2} + \tau_{i,3} = i$ since for all $z \leq i$ a candidate from T is inserted directly before e_z in one of the three partial votes. Thus, e_i looses $3a_i + i$ points in this extension and c beats e_i. It remains to show that c beats t_j. Since $a_{j_1} + a_{j_2} + a_{j_3} = B$, due to the construction the number of candidates that are "better" than t_j in the three partial votes is at least $|\{s \in (D \cup E) : s \succ_{O_1} e_{j_1}\}| + |\{s \in (D \cup E) : s \succ_{O_2} e_{j_2}\}| + |\{s \in (D \cup E) : s \succ_{O_3} e_{j_3}\}| = B - 3$. Thus, $s_p^{\max}(t_j) \leq 3(m-1) - B$ and c will beat t_j.

"$\Leftarrow$" Let V_1, V_2, V_3 denote an extension of P_O in which c is a winner. As explained before, without inserting any $t_i \in T$ before e_i, e_i "looses" $3a_i$ points against c. Hence, e_i must loose further i points. This can only be achieved by inserting candidates of T. Hence, at least i times a candidate of T must be inserted before e_i, that is, $t_{i_1} + t_{i_2} + t_{i_3} \geq i$. We denote this as property (I). In the following, we show first that for a candidate $t_j \in T$ that selects $e_{j_1}, e_{j_2}, e_{j_3}$, one must have $\sum_{q=1}^{3} a_{j_q} = B$ and, second that every $e_i \in E$, that is, every number of A, is selected exactly once.

First, we show by contradiction that for a candidate $t_j \in T$ that selects $e_{j_1}, e_{j_2}, e_{j_3}$, one can neither have $\sum_{q=1}^{3} a_{j_q} < B$ nor $\sum_{q=1}^{3} a_{j_q} > B$. Assume that there is a t_j with $\sum_{q=1}^{3} a_{j_q} < B$. Then, the minimum number of points that t_j will make in any extension is as follows: In $V_q, q \in \{1, 2, 3\}$, there are at most $a_{j_q} - 1$ candidates of $D \cup E$ with $d > t_j$. Thus, t_j can loose at most $\sum_{q=1}^{3} a_{j_q}$ points by candidates from $D \cup E$. Since $|T| = n/3$, in every partial vote at most $n/3 - 1$ candidates of T can be inserted before t_j. Thus, the score of t_j in V_1, V_2, V_3 is at least $3(m-1) - \sum_{q=1}^{3} a_{j_q} - n + 3$. By assumption, $\sum_{q=1}^{3} a_{j_q} < B$ and since all $a_{j_q} \in A$ are multiples of n,

$\sum_{q=1}^{3} a_{j_q} \leq B - n$. Then, in total t_j will make at least $3(m-1) - B + n - n + 3 = 3(m-1) - B + 3 > s_p^{\max}(t_j)$ points in the partial votes, and t_j thus beats c.

Now, assume that there is a t_j with $\sum_{q=1}^{3} a_{j_q} > B$. We consider the amount of points all remaining candidates of $T \backslash \{t_j\}$ together can loose against c by candidates from $D \cup E$. Recall that i candidates of T must be inserted before any e_i (property (I)). Clearly, inserting all candidates as far right as possible maximizes the amount of points the candidates of $T \backslash \{t_j\}$ can loose. Due to property (I) this amount is at most $\sum_{i=1}^{n} a_i - \sum_{q=1}^{3} a_{j_q} = (n/3) \cdot B - \sum_{q=1}^{3} a_{j_q} < (n/3 - 1) \cdot B$. Further, this amount can be "contributed" to the candidates only in multiples of n since a_i differs from a_j at least by n. Then at least one candidate $t \in T \backslash \{t_j\}$ can only loose less than the average amount of points. More precisely, t must loose less than B points and, thus, can loose at most $B-n$ points by candidates from $D \cup E$. Again, we have that t can loose at most $n-3$ additional points by inserting candidates of T. Thus, the minimum score that t will make in V_1, V_2, and V_3 is $3(m-1) - B + n - n + 3 > s_p^{\max}(t)$ and t will beat c.

Second, it remains to show that every number in A is selected exactly once. We cannot select e_n twice without violating property (I). For $i < n$, assume that a candidate e_i is selected twice. Due to property (I), we must have selected at least $i+1$ candidates corresponding to $a_s, s \leq i < n$. Hence, in total, the sum of the numbers corresponding to the selected candidates is at most $\sum_{s=1}^{i} a_i + a_i + \sum_{s=i+2}^{n} a_i < \sum_{s=1}^{n} a_i = (n/3) \cdot B$. Since we have shown before that every candidate t_j selects candidates that sum up exactly to B and we have $|T| = n/3$, this is a contradiction.

Summarizing, in any extension where c wins, the selected candidates of T correspond to a solution for 3-n-PARTITION. □

4 Total number of undetermined pairs

Xia and Conitzer [2008] showed for five common voting rules that the POSSIBLE WINNER problem is NP-complete even if each partial order only contains a constant number of undetermined pairs of candidates. As a consequence, there is no hope for showing fixed-parameter tractability with respect to this parameter. To chart the border of tractability, we consider the parameter k denoting the "total number of undetermined pairs". More precisely, let P_O be a partial profile over C. For $O \in P_O$, let $u(O)$ denote the number of undetermined pairs in O, that is, $u(O) := |\{\{c_1, c_2\} \in C : (c_1 \succ c_2) \notin O \text{ and } (c_2 \succ c_1) \notin O\}|$. Then, $k := \sum_{O \in P_O} u(O)$.

A general search tree approach. Consider a partial order where the candidates c_1 and c_2 form an undetermined pair. To extend this partial order into a linear order, one has to decide whether either $c_1 \succ c_2$ or $c_2 \succ c_1$. Clearly, it may happen that not each of these options is compatible with already fixed pairwise rankings within the given partial orders. Then, this option can be discarded. However, in the worst case, one faces a branching into two valid cases, in each branch decreasing the parameter denoting the total number of undetermined pairs by one. Clearly, this yields a search tree of size $O(2^k)$. For an arbitrary voting rule r, let $f_r(n, m)$ denote the running time needed to compute a winner when given linear orders. Then, for every leaf one can check whether c is a winner for the corresponding extension in $f_r(n, m)$ time, giving the following theorem.

Theorem 5. *For a partial n-voter profile over m candidates and a voting rule r,* POSSIBLE WINNER *can be decided in $O(2^k \cdot (m + f_r(n, m)) + nm^2)$ time, where k denotes the total number of undetermined pairs.*

Theorem 5 is based on a pure worst-case analysis. Significant practical improvements are conceivable. For instance, it is promising to select the order in which the undetermined pairs are processed in a more clever way. Subsequently, for the Borda rule we demonstrate that a provable improvement over the straightforward search tree size of $O(2^k)$ is possible by using a refined search strategy.

An improved search tree for Borda's rule. A central observation concerning Borda for getting a search tree asymptotically smaller than $O(2^k)$ lies in the detection of a polynomial-time solvable special case. To this end, for a partial order O, we define an *isolated undetermined pair of candidates* to be an undetermined pair where both candidates do not form an undetermined pair in O with any other candidate.

Theorem 6. *For k being the total number of undetermined pairs, in case of Borda* POSSIBLE WINNER *can be decided in $O(nm^2 + k^2)$ time if all undetermined pairs are isolated.*

Proof. If the distinguished candidate c is contained in an undetermined pair, then, in the linear order, c is always placed in front of the second candidate of the pair. After that, one can assume that none of the undetermined pairs contains c. For an isolated undetermined pair $\{c_1, c_2\}$ of candidates, the relative order of c_1 and c_2 with respect to all other candidates is already determined. More precisely, it is not hard to see that c_1 and c_2 must have the same relative order with respect to each of the remaining candidates, and, thus, they must be direct neighbors in the final linear order. Then, their scores in this final order will differ by exactly one point. Thus, for every candidate $c' \neq c$ being in at least one undetermined pair, one can compute the minimum number of points $l(c')$ that c' will make in every possible extension. That is, $l(c')$ is the sum over the scores for c' obtained by choosing $c'' \succ c'$ in all undetermined pairs $\{c', c''\}$ that contain c'. Note that the score $s(c)$ for the distinguished candidate c is already fixed. Clearly, if $l(c') \geq s(c)$ for some c', then c cannot become possible winner. Otherwise, let $b(c') := s(c) - l(c') - 1 \geq 0$ denote the *balance* of c' with respect to c. The balance counts the number of partial orders where c' may be placed better than the other candidate in an undetermined pair without defeating c.

Using the balance $b(c')$ for all candidates $c' \in C \backslash \{c\}$, one can decide POSSIBLE WINNER with the help of a maximum flow computation as follows. Consider a four-level directed, arc-weighted s-t-network with distinguished vertices s and t (see Figure 1). The first level only consists of vertex s. The second level consists of vertices one-to-one representing all undetermined pairs. Note that the same two candidates may induce more than one undetermined pair because they may

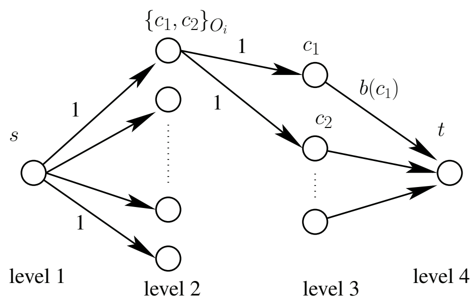

Figure 1: Flow network for Borda with isolated undetermined pairs

occur in more than one partial order. The vertex s is connected by arcs of weight one to all level-two vertices. The third level of vertices one-to-one represents all candidates occurring in at least one undetermined pair. Every level-two vertex representing an undetermined pair is connected by two weight-one arcs to the two vertices corresponding to the two candidates contained in the undetermined pair. The fourth level only consists of vertex t. Every level-three vertex representing a candidate c' is connected by one arc to t which is assigned the weight $b(c')$.

The central claim now is as follows (proof omitted): The constructed flow network allows for an integer flow of value k iff the distinguished candidate c is a possible winner of the corresponding Borda instance with k undetermined pairs.

Altogether, we arrive at the overall running time $O(nm^2 + k^2)$. To this end, note that the flow network can be constructed in $O(nm^2)$ time. Further, the number of arcs of the flow network is linear in k. The Ford-Fulkerson algorithm can compute a maximum integer flow in $O(|A| \cdot f)$ time, where $|A|$ denotes the number of arcs and f denotes the value of a maximum flow. Since the value of the maximum flow is bounded by k, the claimed running time follows. □

The basic idea for an improved search tree algorithm for Borda in the general case is as follows. Three candidates $\{c_1, c_2, c_3\} \subseteq C$ form an *undetermined triple* with respect to some partial order $O \in P_O$ if there are at least two undetermined pairs in O, each formed by two candidates from $\{c_1, c_2, c_3\}$. Now, one branches on undetermined triples instead of undetermined pairs. This leads to a search tree of size 1.82^k. Once there are no more undetermined triples, then one can show that all remaining undetermined pairs must be isolated. Hence, in the leaves of the refined search tree the algorithm from Theorem 6 can be applied.

Theorem 7. *For the Borda rule,* POSSIBLE WINNER *can be decided in* $O(1.82^k(nm^2 + k^2))$ *time, where* k *denotes the total number of undetermined pairs.*

5 Conclusion

With our multivariate complexity analysis for the POSSIBLE WINNER problem we complement previous work [Walsh, 2007; Xia and Conitzer, 2008]. We studied parameterizations based on the number of candidates, the number of votes, and the total number of undetermined candidate pairs. It is conceivable that further parameterizations are worth investigation. Moreover, whereas we only dealt with two-dimensional (that is, one parameter at a time) complexity analysis, it also seems prospective to study the dependence on parameter pairs (three-dimensional) or even higher-dimensional parameterizations. All our results also hold for the possible co-winner case. We conclude with a few concrete challenges for future work: Make the fixed-parameter tractability results from Section 2 more practical by replacing integer linear programming with combinatorial algorithms. Concerning the efficient enumeration of extensions, study in how many of all possible extensions a distinguished candidate is a winner. Consider the case where one does not allow all possible extensions of partial orders but restricts these to "CP-nets", see [Xia and Conitzer, 2008].

References

[Conitzer *et al.*, 2007] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.

[Faliszewski *et al.*, 2007] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Llull and Copeland voting broadly resist bribery and control. In *Proc. of 22nd AAAI'07*, pages 724–730, 2007.

[Faliszewski *et al.*, 2008] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Copeland voting: Ties matter. In *Proc. of 7th AAMAS 08*, pages 983–990, 2008.

[Garey and Johnson, 1979] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[Konczak and Lang, 2005] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proc. of IJCAI-2005 Multidisciplinary Workshop on Advances in Preference Handling*, 2005.

[Lang *et al.*, 2007] J. Lang, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Winner determination in sequential majority voting. In *Proc. of 20th IJCAI '07*, pages 1372–1377, 2007.

[Niedermeier, 2006] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.

[Pini *et al.*, 2007] M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Incompleteness and incomparability in preference aggregation. In *Proc. of 20th IJCAI '07*, pages 1464–1469, 2007.

[Walsh, 2007] T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proc. of 22nd AAAI '07*, pages 3–8. AAAI Press, 2007.

[Xia and Conitzer, 2008] L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proc. of 23rd AAAI '08*, pages 196–201. AAAI Press, 2008.

[Xia *et al.*, 2008] L. Xia, V. Conitzer, A. D. Procaccia, and J. S. Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proc. 2nd COMSOC '08*, pages 427–435, 2008.