Complexity of Conformant Election Manipulation

Zack Fitzsimmons[1(✉)] and Edith Hemaspaandra[2]

[1] College of the Holy Cross, Worcester, MA 01610, USA
zfitzsim@holycross.edu
[2] Rochester Institute of Technology, Rochester, NY 14623, USA

Abstract. It is important to study how strategic agents can affect the outcome of an election. There has been a long line of research in the computational study of elections on the complexity of manipulative actions such as manipulation and bribery. These problems model scenarios such as voters casting strategic votes and agents campaigning for voters to change their votes to make a desired candidate win. A common assumption is that the preferences of the voters follow the structure of a domain restriction such as single peakedness, and so manipulators only consider votes that also satisfy this restriction. We introduce the model where the preferences of the voters define their own restriction and strategic actions must "conform" by using only these votes. In this model, the election after manipulation will retain common domain restrictions. We explore the computational complexity of conformant manipulative actions and we discuss how conformant manipulative actions relate to other manipulative actions.

1 Introduction

The computational study of election problems is motivated by the utility of elections to aggregate preferences in multiagent systems and to better understand the computational tradeoffs between different rules. A major direction in this area has been to study the computational complexity of manipulative actions on elections (see, e.g., Faliszewski and Rothe [15]).

The problems of manipulation [1] and bribery [14] in elections represent two important ways that agent(s) can strategically affect the outcome of an election. Manipulation models the actions of a collection of strategic voters who seek to ensure that their preferred candidate wins by casting strategic votes. Bribery models the actions of an agent, often referred to as the briber, who sets the votes of a subcollection of the voters to ensure that the briber's preferred candidate wins. These problems each relate nicely to real-world scenarios such as how voters may attempt to work together to strategically vote, or the actions of a campaign manager looking to influence the preferences of a group of voters to ensure their candidate wins.

In the manipulation problem each manipulator can cast any strategic vote, and similarly for the bribery problem the votes can be set to any collection

H. Fernau and K. Jansen (Eds.): FCT 2023, LNCS 14292, pp. 176–189, 2023.
https://doi.org/10.1007/978-3-031-43587-4_13

of preferences. However, this is not always a reasonable assumption to make. Voters may have preferences that satisfy a domain restriction such as single-peaked preferences [6] or single-crossing preferences [27] where the manipulator or the briber are restricted to using votes that also satisfy the restriction. We introduce new models of manipulation and bribery where the votes cast by the manipulators or set by the briber must have already been stated in the election, i.e., the votes must *conform* to the views of the electorate. In these conformant models of manipulation and bribery the election after the given manipulative action will retain common domain restrictions such as being single-peaked or single-crossing.

We consider how the computational complexity of our conformant models of manipulation and bribery compare to the standard models. Specifically, we show that there are settings where manipulation and bribery are easy in the standard model, but computationally difficult in the conformant model, and vice versa. This shows that there is no reduction in either direction between the standard and conformant cases (unless P = NP).

Conformant manipulation and bribery are also each related to electoral control. The study of electoral control was introduced by Bartholdi, Tovey, and Trick [2], and it models the actions of an agent who can modify the structure of an election to ensure a preferred candidate wins (e.g., by adding or deleting voters). We explore the connection between conformant manipulation and the exact variant of voter control as well as conformant bribery and the model of control by replacing voters introduced by Loreggia et al. [26]. This includes showing reductions between these problems as well as showing when such reductions cannot exist. Exact versions of electoral control problems can model scenarios where the election chair seeks to ensure their preferred outcome by adding exactly the number of voters required to meet the quorum for a vote. This is in line with the standard motivation for control by replacing voters which includes settings such as voting in a parliament where the chair may replace some of the voters, but makes sure to leave the total number the same to avoid detection [26].

Our main contributions are as follows.

- We introduce the problems of conformant manipulation and conformant bribery, which model natural settings for manipulative attacks on elections.
- We show that there is no reduction in either direction between the standard and conformant versions of manipulation (Theorems 1 and 2) and between the standard and conformant versions of bribery (Theorem 8 and Corollary 2) (unless P = NP).
- We show conformant manipulation reduces to exact control by adding voters (Theorem 3), conformant bribery reduces to control by replacing voters (Theorem 11), and that reductions do not exist in the other direction (Theorems 4 and 12) (unless P = NP).
- We obtain a trichotomy theorem for the complexity of exact control by adding voters problem (Theorem 5) for the important class of pure scoring rules.

Due to space constraints some of our proofs have been deferred to the full version [17].

2 Related Models

Since the conformant model of manipulation uses only the votes stated in the initial election, it is related to the possible winner problem with uncertain weights introduced by Baumeister et al. [3], which was recently extended by Neveling et al. [29]. In this problem weights of voters are initially unset and it asks if there exists a way to set the weights such that a given candidate is a winner. As Baumeister et al. [3] mention, this generalizes control by adding/deleting voters (rather than manipulation or bribery).

There are other models for manipulative actions that have a similar motivation to our conformant models, i.e., to have the manipulative action not stand out. Examples are bribery restrictions where the briber cannot put the preferred candidate first in the bribed votes (negative bribery [14]), restrictions on how much the votes of voters are changed [34], and restrictions on which voters can be bribed [10]. However, it is easy to see that in these models all sorts of new votes can be used by the briber, not just votes appearing in the initial election.

Another model in which the votes of the strategic agents are restricted is that of manipulative actions on restricted domains such as single-peaked [6] and single-crossing preferences [27]. For example, for manipulation of a single-peaked election the manipulators must all cast votes that are single-peaked with respect to the rest of the electorate [32]. Notice that in the conformant models we typically keep the domain restriction, but manipulative actions for domain restrictions are quite different since in those settings the manipulators can cast a vote not stated by any of the nonmanipulators as long as it satisfies the given restriction.

3 Preliminaries

An election (C, V) consists of a set of candidates C and a collection of voters V. Each voter $v \in V$ has a corresponding vote, which is a strict total order preference over the set of candidates.

A voting rule, $\mathcal{E}$, is a mapping from an election to a subset of the candidate set referred to as the winner(s).

3.1 Scoring Rules

Our results focus on (polynomial-time uniform) pure scoring rules. A scoring rule is a voting rule defined by a family of scoring vectors of the form $\langle \alpha_1, \alpha_2, \ldots, \alpha_m \rangle$ with $\alpha_i \geq \alpha_{i+1}$ such that for a given election with m candidates the m-length scoring vector is used and each candidate receives α_i points for each vote where they are ranked ith. The candidate(s) with the highest score win. We use the notation score(a) to denote the score of a candidate a in a given election. Important examples of scoring rules are

– k-Approval, $\langle \underbrace{1, \ldots, 1}_{k}, 0, \ldots, 0 \rangle$

– k-Veto, $\langle 0, \ldots, 0, \underbrace{-1, \ldots, -1}_{k} \rangle$
– Borda, $\langle m-1, m-2, \ldots, 1, 0 \rangle$
– First-Last, $\langle 1, 0, \ldots, 0, -1 \rangle$[1]

Note that there are uncountably many scoring rules, that scoring vectors may not be computable, and that the definition of scoring rule does not require any relationship between the scoring vectors for different numbers of candidates. To formalize the notion of a natural scoring rule, we use the notion of *(polynomial-time uniform) pure scoring rules* [4]. These are families of scoring rules where the scoring vector for $m+1$ candidates can be obtained from the scoring vector for m candidates by adding one coefficient and for which there is a polynomial-time computable function that outputs, on input 0^m, the scoring vector for m. Note that the election rules above are all pure scoring rules. Also note that manipulative action problems for pure scoring rules are in NP.

3.2 Manipulative Actions

Two of the most commonly-studied manipulative actions on elections are manipulation [1] and bribery [14]. We consider conformant variants of these standard problems by requiring that the strategic votes cast by the manipulators or set by the briber must have appeared in the initial election. We define these problems formally below.

Name: $\mathcal{E}$-Conformant Manipulation
Given: An election (C, V), a collection of manipulative voters W, and a preferred candidate p.
Question: Does there exist a way to set the votes of the manipulators in W using only the votes that occur in V such that p is an $\mathcal{E}$-winner of the election $(C, V \cup W)$?

Name: $\mathcal{E}$-Conformant Bribery
Given: An election (C, V), a bribe limit k, and a preferred candidate p.
Question: Does there exist a way to set the preferences of a subcollection of at most k voters in V to preferences in V such that p is an $\mathcal{E}$-winner?

3.3 Computational Complexity

We assume the reader is familiar with the complexity classes P and NP, polynomial-time many-one reductions, and what it means to be complete for a given class. Our NP-completeness proofs utilize reductions from the well-known NP-complete problem 3-Dimensional Matching [23].

Name: 3-Dimensional Matching (3DM)

[1] We will see that this rule exhibits very unusual complexity behavior. This rule has also been referred to as "best-worst" in social choice (see, e.g., [24]).

Given: Pairwise disjoint sets X, Y, and Z such that $\|X\| = \|Y\| = \|Z\| = k > 0$ and set $\mathcal{M} \subseteq X \times Y \times Z$.
Question: Does there exist a $\mathcal{M}' \subseteq \mathcal{M}$ of size k such that each $a \in X \cup Y \cup Z$ appears exactly once in $\mathcal{M}'$.

3-Dimensional Matching remains NP-complete when each element $a \in X \cup Y \cup Z$ appears in exactly three triples (Garey and Johnson [19] show this for at most three triples, which can be adapted to exactly three by the approach from Papadimitriou and Yannakakis [30]). Note that in that case $\|\mathcal{M}\| = 3k$.

For our polynomial-time algorithms, we will reduce to polynomial-time computable (edge) matching problems. We define the most general version we use, Max-Weight b-Matching for Multigraphs, below. The version of this problem for simple graphs was shown to be in P by Edmonds and Johnson [12], and as explained in [20, Section 7], it is easy to reduce such problems to Max-Weight Matching, which is well-known to be in P [11], using the construction from [31]. (Note that we can assume that the b-values are bound by the number of edges in the graph.)

Name: Max-Weight b-Matching for Multigraphs
Given: An edge-weighted multigraph $G = (V, E)$, a function $b : V \to \mathbb{N}$, and integer $k \geq 0$.
Question: Does there exist an $E' \subseteq E$ of weight at least k such that each vertex $v \in V$ is incident to at most $b(v)$ edges in E'?

In addition to NP-hardness and polynomial-time results, we have results that link the complexity of voting problems to the complexity of Exact Perfect Bipartite Matching [30].

Name: Exact Perfect Bipartite Matching
Given: A bipartite graph $G = (V, E)$, a set of red edges $E' \subseteq E$, and integer $k \geq 0$.
Question: Does G contain a perfect matching that contains exactly k edges from E'?

This problem was shown to be in RP by Mulmuley, Vazirani, and Vazirani [28], but it is a 40-year-old open problem whether it is in P.

4 Conformant Manipulation

The problem of manipulation asks if it is possible for a given collection of manipulative voters to set their votes so that their preferred candidate wins. This problem was first studied computationally by Bartholdi, Tovey, and Trick [1] for the case of one manipulator and generalized by Conitzer, Sandholm, and Lang [8] for the case of a coalition of manipulators.

In our model of conformant manipulation the manipulators can only cast votes that at least one nonmanipulator has stated. As mentioned in the introduction, this is so that the manipulators vote realistic preferences for the given

election by conforming to the preferences already stated. Since this modifies the standard model of manipulation, we will consider how the complexity of these problems relate to one another.

Manipulation is typically easy for scoring rules. Hemaspaandra and Schnoor [22] showed that manipulation is in P for every pure scoring rule with a constant number of different coefficients. However, our model of conformant manipulation is NP-complete for even the simple rule 4-approval. To show hardness, we use a construction similar to the construction that shows hardness for control by adding voters from Lin [25].

Theorem 1. *4-Approval Conformant Manipulation is* NP*-complete.*

Proof. Let the pairwise disjoint sets X, Y, and Z such that $\|X\| = \|Y\| = \|Z\| = k > 0$ and $\mathcal{M} \subseteq X \times Y \times Z$ be an instance of 3DM where each $a \in X \cup Y \cup Z$ appears in exactly three triples. Note that $\|\mathcal{M}\| = 3k$. We construct an instance of conformant manipulation as follows.

Let the candidate set C consist of preferred candidate p, and for each a in $X \cup Y \cup Z$, we have candidate a and three padding candidates a_1, a_2, and a_3. We now construct the collection of nonmanipulators.

- For each (x, y, z) in $\mathcal{M}$, we have a voter voting $p > x > y > z > \cdots$.
- For each a in $X \cup Y \cup Z$, we have $4k-4$ voters voting $a > a_1 > a_2 > a_3 > \cdots$.

We have k manipulators.

Note that we have the following scores from the nonmanipulators. $\text{score}(p) = 3k$ and for $a \in X \cup Y \cup Z$, $\text{score}(a) = 4k - 4 + 3 = 4k - 1$ and $\text{score}(a_1) = \text{score}(a_2) = \text{score}(a_3) = 4k - 4$.

If there is a matching, let the k manipulators vote corresponding to the matching. Then p's score increases by k, for a total of $4k$, and for each $a \in X \cup Y \cup Z$, $\text{score}(a)$ increases by 1 for a total of $4k$. The scores of the dummy candidates remain unchanged. Thus, p is a winner.

For the converse, suppose the manipulators vote such that p is a winner. Since $k > 0$, after manipulation there is a candidate a in $X \cup Y \cup Z$ with score at least $4k$. The highest possible score for p after manipulation is $4k$, and this happens only if p is approved by every manipulator. It follows that every manipulator approves p and for every a in $X \cup Y \cup Z$, a is approved by at most one manipulator. This implies that the votes of the manipulators correspond to a cover. □

We just saw a case where the complexity of conformant manipulation is harder than the standard problem (unless P = NP). This is not always the case. One setting where it is clear to see how to determine if conformant manipulation is possible or not is when there are only a fixed number of manipulators. We have only a polynomial number of votes to choose from (the votes of the nonmanipulators) and so a fixed number of manipulators can brute force these choices in polynomial time as long as determining the winner can be done in polynomial time.

Theorem 2. *Conformant Manipulation is in* P *for every voting rule with a polynomial-time winner problem when there are a fixed number of manipulators.*

This behavior is in contrast to what can occur for the standard model of manipulation. One well-known example is that manipulation for the Borda rule is NP-complete even when there are only two manipulators [5,9]. Intuitively the hardness of manipulation in this case is realized by the choice of different votes that the manipulator(s) have.

Corollary 1. *For Borda, Manipulation with two manipulators is* NP*-complete, but Conformant Manipulation with two manipulators is in* P*.*

In some ways conformant manipulation acts more like the problem of electoral control introduced by Bartholdi, Tovey, and Trick [2], specifically *Control by Adding Voters*, which asks when given an election, a collection of unregistered voters, add limit k, and preferred candidate p, if p is a winner of the election after adding at most k of the unregistered voters. In conformant manipulation we can think of the nonmanipulative voters as describing the different votes to choose from for the manipulators.

At first glance it may appear that there is a straightforward reduction from conformant manipulation to control by adding voters, but in conformant manipulation *all* k of the manipulators must cast a vote, while in control by adding voters *at most* k votes are added. In this way conformant manipulation is closer to the "exact" variant of control by adding voters where *exactly* k unregistered voters must be added. The following is immediate.

Theorem 3. *Conformant Manipulation polynomial-time many-one reduces to Exact Control by Adding Voters.*

Below we show that conformant manipulation is in P for the voting rule 3-approval, but exact control by adding voters is NP-complete. And so there is no reduction from exact control by adding voters to conformant manipulation (unless P = NP).

Theorem 4. *For 3-Approval, Exact Control by Adding Voters is* NP*-complete, but Conformant Manipulation is in* P*.*

Proof. Given an election (C, V), k manipulators, and a preferred candidate p, we can determine if conformant manipulation is possible in the following way.

If there is no nonmanipulator that approves p then p is a winner if and only if there are no voters.

If there is at least one nonmanipulator that approves p, the manipulators will all cast votes that approve of p, and so we know the final score of p after manipulation: $fs_p = \text{score}(p) + k$. However, these manipulator votes will each also approve of two other candidates. To handle this we can adapt the approach used by Lin [25] to show control by adding voters is in P for 3-approval elections, which constructs a reduction to Max-Cardinality b-Matching for Multigraphs.

For each candidate $a \neq p$, let $b(a) = fs_p - \text{score}(a)$, i.e., the maximum number of approvals that a can receive from the manipulators without beating p. If $b(a)$ is negative for any candidate $a \neq p$ then conformant manipulation is not possible. For each distinct nonmanipulative vote of the form $\{p, a, b\} > \ldots$, add k edges between a and b. Conformant manipulation is possible if and only if there is a b-edge matching of size at least k.

We now consider the complexity of exact control by adding voters. Given an instance of 3-Dimensional Matching: pairwise disjoint sets X, Y, and Z such that $\|X\| = \|Y\| = \|Z\| = k$, and $\mathcal{M} \subseteq X \times Y \times Z$ with $\mathcal{M} = \{M_1, \ldots, M_{3k}\}$ we construct the following instance of exact control by adding voters.

Let the candidate set $C = \{p, d_1, d_2\} \cup X \cup Y \cup Z$, the preferred candidate be p, and the add limit be k.

Let there be one registered voter voting $p > d_1 > d_2 > \ldots$, and let the set of unregistered voters consist of one voter voting $x > y > z > \ldots$ for each $M_i = (x, y, z)$.

It is easy to see that p can be made a winner by adding exactly k unregistered voters if and only if there is a matching of size k. □

The standard case of control by adding voters for 3-approval is in P [25], but as shown above the exact case is NP-complete. Related work that mentions exact control has results only where the exact variant is also easy [13,16]. Note that, as observed in [16], control polynomial-time reduces to exact control, since, for example, p can be made a winner by adding at most k voters if and only if p can be made a winner by adding 0 voters or 1 voter or 2 voters or ...), and so if exact control is easy, the standard case will be as well, and if the standard case is hard, then the exact case will be hard. Note that we are using a somewhat more flexible notion of reducibility than many-one reducibility here, since we are allowing the disjunction of multiple queries to the exact control problem. Such a reduction is called a *disjunctive truth-table (dtt) reduction*. This type of reduction is still much less flexible than a Turing reduction.

Is 3-approval special? For the infinite and natural class of pure scoring rules, Table 1 completely classifies the complexity of exact control by adding voters and compares this behavior to the complexity of control by adding voters [21] and control by deleting voters [22]. In particular, and in contrast to earlier results, we obtain a *trichotomy* theorem for exact control by adding voters.[2]

Theorem 5. *For every pure scoring rule f,*

1. *If f is ultimately (i.e., for all but a finite number of candidates) equivalent to 0-approval, 1-approval, 2-approval, 1-veto, or 2-veto, exact control by adding voters is in P.*
2. *If f is ultimately equivalent to first-last, then exact control by adding voters is (dtt) equivalent to the problem Exact Perfect Bipartite Matching [30].*

[2] Exact Perfect Bipartite Matching [30] is defined in Sect. 3.3. As mentioned there the complexity of this problem is still open. And so Theorem 5 is a trichotomy theorem unless we solve a 40-year-old open problem.

Table 1. This table classifies the complexity of all pure scoring rules for the specified control action. A scoring rule entry represents all pure scoring rules that are ultimately equivalent to that scoring rule. The dichotomy for control by adding voters is due to [21], the dichotomy for control by deleting voters to [22], and the result for exact control by adding voters for first-last is due to [16]. EPBM stands for the Exact Perfect Bipartite Matching, which is defined in Sect. 3.3.

	P	eq. to EPBM	NP-complete
Exact Control by Adding Voters	0/1/2-approval, 1/2-veto	first-last	all other cases
Control by Adding Voters	0/1/2/3-approval, 1/2-veto, first-last, $\langle \alpha, \beta, 0, \ldots, 0 \rangle$		all other cases
Control by Deleting Voters	0/1/2-approval, 1/2/3-veto, first-last, $\langle 0, \ldots, 0, -\beta, -\alpha \rangle$		all other cases

3. In all other cases, exact control by adding voters is NP-complete (under dtt reductions).

Proof.

1. The case for 0-approval is trivial, since all candidates are always winners. The 1-approval and 1-veto cases follow by straightforward greedy algorithms. The case for 2-approval can be found in [16] and the case for 2-veto is similar.
2. The case for first-last can be found in [16].
3. Note that the remaining cases are hard for control by adding voters or for control by deleting voters. We have already explained how we can dtt reduce control by adding voters to exact control by adding voters. Similarly, we can dtt reduce control by deleting voters to exact control by adding voters, since p can be made a winner by deleting at most k voters if and only p can be made a winner by adding to the empty set $n - k$ voters or $n - k + 1$ voters or $n - k + 2$ voters or $\ldots$ or n voters, where n is the total number of voters. It follows that all these cases are NP-complete (under dtt reductions). □

In this section, we compared conformant manipulation with manipulation and conformant manipulation with exact control by adding voters. We can also look at conformant manipulation versus control by adding voters. Here we also find voting rules where the manipulative actions differ. In particular, control by adding voters for first-last is in P [21], but we show in the full version that conformant manipulation for first-last is equivalent to exact perfect bipartite matching.

Theorem 6. *First-Last Conformant Manipulation is equivalent to Exact Perfect Bipartite Matching (under dtt reductions).*

We also show in the full version that there exists a voting rule where control by adding voters is hard and conformant manipulation is easy.

Theorem 7. *There exists a voting rule where Control by Adding Voters is* NP*-complete, but Conformant Manipulation is in* P*.*

This shows that a reduction from control by adding voters to conformant manipulation does not exist (unless P = NP).

5 Conformant Bribery

We now turn to our model of conformant bribery. The standard bribery problem introduced by Faliszewski, Hemaspaandra, and Hemaspaandra [14] asks when given an election, a bribe limit k, and a preferred candidate p, if there exists a subcollection of at most k voters whose votes can be changed such that p is a winner. In our model of conformant bribery the votes can only be changed to votes that appear in the initial election. As with conformant manipulation, this is so that the votes are changed to preferences that are still realistic with respect to the preferences already stated. Notice how this also nicely models how a voter can convince another to vote their same vote. In the same way as we did with manipulation in the previous section, we can compare the complexity behavior of our conformant model with respect to the standard model.

Bribery is in P for the voting rule 3-veto [25], but we show below that our model of conformant bribery is NP-complete for this rule.

Theorem 8. *3-Veto Conformant Bribery is* NP*-complete.*

Proof. Let X, Y, and Z be pairwise disjoint sets such that $\|X\| = \|Y\| = \|Z\| = k$, and $\mathcal{M} \subseteq X \times Y \times Z$ with $\mathcal{M} = \{M_1, \ldots, M_{3k}\}$ be an instance of 3DM where each element $a \in X \cup Y \cup Z$ appears in exactly three triples. We construct an instance of conformant bribery as follows.

Let the candidate set $C = \{p\} \cup X \cup Y \cup Z \cup \{p_1, p_2\}$. Let p be the preferred candidate and let k be the bribe limit. Let there be the following voters.

- For each $M_i = (x, y, z)$,
  - One voter voting $\cdots > x > y > z$
- $k + 4$ voters voting $\cdots > p > p_1 > p_2$

We view the corresponding scoring vector for 3-veto as $\langle 0, \ldots, 0, -1, -1, -1\rangle$ to make our argument more straightforward. And so, before bribery $\text{score}(p) = \text{score}(p_1) = \text{score}(p_2) = -k - 4$ and for each $a \in X \cup Y \cup Z$, $\text{score}(a) = -3$.

If there exists a matching $\mathcal{M}' \subseteq \mathcal{M}$ of size k, for each $M_i \in \mathcal{M}'$ such that $M_i = (x, y, z)$ we can bribe one of the voters voting $\cdots > p > p_1 > p_2$ to vote $\cdots > x > y > z$. Since $\mathcal{M}'$ is a matching the score of each candidate $a \in X \cup Y \cup Z$ decreases by 1 to be -4, and since k of the voters vetoing p are bribed the score of p increases by k to -4 and p is a winner.

For the converse, suppose there is a successful conformant bribery. Only the voters vetoing p should be bribed, and so the score of p after bribery is -4. The score of each candidate $a \in X \cup Y \cup Z$ must decrease by at least 1, and so it is easy to see that a successful conformant bribery of at most k voters must correspond to a perfect matching. □

We now consider the case where bribery is hard for the standard model, but easy in our conformant model. Since the briber is restricted to use only votes that appear in the initial election we have the same behavior as stated in Theorem 2 for conformant manipulation.

Theorem 9. *Conformant Bribery is in* P *for every voting rule with a polynomial-time winner problem when there is a fixed bribe limit.*

There are generally fewer results looking at a fixed bribe limit than there are looking at a fixed number of manipulators. One example is that for Single Transferable Vote (STV), bribery is NP-complete even when the bribe limit is 1 [33], but our focus is on scoring rules. Brelsford at al. [7] show bribery is NP-complete for Borda, but do not consider a fixed bribe limit. However, it is easy to adapt the NP-hardness proof for Borda manipulation with two manipulators from Davies et al. [9]. The main idea is to add two voters that are so bad for the preferred candidate that they have to be bribed.

Theorem 10. *Borda Bribery with a bribe limit of 2 is* NP*-complete.*

Proof. We need to following properties of the instance of Borda Manipulation with two manipulators constructed in Davies et al. [9]. (For this proof we use the notation from Davies et al. [9].) The constructed election has a collection of (nonmanipulative) voters V and $q+3$ candidates. Preferred candidate p scores C and candidate a_{q+1} scores $2(q+2)+C$. This implies that in order for p to be a winner, the two manipulators must vote p first and a_{q+1} last. We add two voters voting $a_{q+1} > \cdots > p$. Note that such votes are very bad for p and that we have to bribe two voters voting $a_{q+1} > \cdots > p$. In order to ensure that we bribe exactly the two added voters, it suffices to observe that we can ensure in the construction from Davies et al. [9] that p is never last in any vote in V. So, p can be made a winner by bribing two voters in $V \cup \{a_{q+1} > \cdots > p, a_{q+1} > \cdots > p\}$ if and only if p can be made a winner by two manipulators in the election constructed by Davies et al. [9]. □

Corollary 2. *For Borda, Bribery is* NP*-complete with a bribe limit of 2, but Conformant Bribery is in* P *with a bribe limit of 2.*

Bribery can be thought of as control by deleting voters followed by manipulation. For conformant bribery we can see that the same will hold, just with conformant manipulation. However, we also have a correspondence to the problem of control by replacing voters introduced by Loreggia et al. [26]. Control by replacing voters asks when given an election, a collection of unregistered voters, parameter k, and preferred candidate p, if p can be made a winner by replacing at most k voters in the given election with a subcollection of the unregistered voters of the same size. It is straightforward to reduce conformant bribery to control by replacing voters (for each original voter v, we have a registered copy of v and k unregistered copies of v), and so we inherit polynomial-time results from this setting.

Theorem 11. *Conformant Bribery polynomial-time many-one reduces to Control by Replacing Voters.*

It's natural to ask if there is a setting where conformant bribery is easy, but control by replacing voters is hard, and we show in the full version that this is in fact the case.

Theorem 12. *There exists a voting rule where Control by Replacing Voters is* NP*-complete, but Conformant Bribery is in* P*.*

In the related work on control by replacing voters, only the complexity for 2-approval remained open (see Erdélyi et al. [13]). This was recently shown to be in P by Fitzsimmons and Hemaspaandra [18]. This result immediately implies that conformant bribery for 2-approval is also in P.

Theorem 13. *2-Approval Conformant Bribery is in* P*.*

2-approval appears right at the edge of what is easy. For 3-approval, control by deleting voters and bribery are hard [25], control by replacing voters is hard [13], and we show in the full version that conformant bribery is hard as well (recall that for 3-approval, conformant manipulation (Theorem 4) and control by adding voters [25] are easy, essentially because all we are doing is "adding" votes that approve p).

Theorem 14. *3-Approval Conformant Bribery is* NP*-complete.*

As a final note, we mention that for first-last, conformant bribery, like conformant manipulation, is equivalent to exact perfect bipartite matching again showing the unusual complexity behavior of this rule.

Theorem 15. *First-Last Conformant Bribery is equivalent to Exact Perfect Bipartite Matching (under dtt reductions).*

The proof of the above theorem can be found in the full version.

6 Conclusion

The conformant models of manipulation and bribery capture a natural setting for election manipulation. We found that there is no reduction between the standard and conformant models in either direction (unless P = NP), and further explored the connection between these models and types of electoral control.

We found the first trichotomy theorem for scoring rules. This theorem concerns the problem of exact control by adding voters and highlights the unusual complexity behavior of the scoring rule first-last. We show that this unusual complexity behavior also occurs for our conformant models.

We also observed interesting behavior for exact variants of control, including a nontrivial case where the complexity of a problem increases when going from the standard to the exact case.

We see several interesting directions for future work. For example, we could look at conformant versions for other bribery problems (e.g., priced bribery) or for restricted domains such as single-peakedness. We are also interested in further exploring the complexity landscape of problems for the scoring rule first-last.

Acknowledgements. This work was supported in part by grant NSF-DUE-1819546. We thank the anonymous reviewers for their helpful comments and suggestions.

References

1. Bartholdi, J., Tovey, C., Trick, M.: The computational difficulty of manipulating an election. Soc. Choice Welfare **6**(3), 227–241 (1989)
2. Bartholdi, J., Tovey, C., Trick, M.: How hard is it to control an election? Math. Comput. Model. **16**(8/9), 27–40 (1992)
3. Baumeister, D., Roos, M., Rothe, J., Schend, L., Xia, L.: The possible winner problem with uncertain weights. In: Proceedings of the 20th European Conference on Artificial Intelligence, pp. 133–138 (2012)
4. Betzler, N., Dorn, B.: Towards a dichotomy of finding possible winners in elections based on scoring rules. J. Comput. Syst. Sci. **76**(8), 812–836 (2010)
5. Betzler, N., Niedermeier, R., Woeginger, G.: Unweighted coalitional manipulation under the Borda rule is NP-hard. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence, pp. 55–60 (2011)
6. Black, D.: On the rationale of group decision-making. J. Polit. Econ. **56**(1), 23–34 (1948)
7. Brelsford, E., Faliszewski, P., Hemaspaandra, E., Schnoor, H., Schnoor, I.: Approximability of manipulating elections. In: Proceedings of the 23rd National Conference on Artificial Intelligence, pp. 44–49 (2008)
8. Conitzer, V., Sandholm, T., Lang, J.: When are elections with few candidates hard to manipulate? J. ACM **54**(3), 1–33 (2007)
9. Davies, J., Katsirelos, G., Narodytska, N., Walsh, T., Xia, L.: Complexity of and algorithms for the manipulation of Borda, Nanson's and Baldwin's voting rules. Artif. Intell. **217**, 20–42 (2014)
10. Dey, P., Misra, N., Narahari, Y.: Frugal bribery in voting. Theor. Comput. Sci. **676**, 15–32 (2017)
11. Edmonds, J.: Maximum matching and a polyhedron with 0,1-vertices. J. Res. National Bureau Standards–B. Math. Math. Phys. **69B**(1/2), 125–130 (1965)
12. Edmonds, J., Johnson, E.: Matching: a well-solved class of integer linear programs. In: Combinatorial Structures and Their Applications (Gordon and Breach), pp. 89–92 (1970)
13. Erdélyi, G., Neveling, M., Reger, C., Rothe, J., Yang, Y., Zorn, R.: Towards completing the puzzle: complexity of control by replacing, adding, and deleting candidates or voters. Auton. Agent. Multi-Agent Syst. **35**(41), 1–48 (2021)
14. Faliszewski, P., Hemaspaandra, E., Hemaspaandra, L.: How hard is bribery in elections? J. Artif. Intell. Res. **35**, 485–532 (2009)
15. Faliszewski, P., Rothe, J.: Control and bribery in voting. In: Handbook of Computational Social Choice, pp. 146–168. Cambridge University Press (2016)
16. Fitzsimmons, Z., Hemaspaandra, E.: Insight into voting problem complexity using randomized classes. In: Proceedings of the 31st International Joint Conference on Artificial Intelligence, pp. 293–299 (2022)

17. Fitzsimmons, Z., Hemaspaandra, E.: Complexity of conformant election manipulation. Tech. Rep. arXiv:2307.11689 [cs.GT], arXiv.org (2023)
18. Fitzsimmons, Z., Hemaspaandra, E.: Using weighted matching to solve 2-approval/veto control and bribery. In: Proceedings of the 26th European Conference on Artificial Intelligence (2023), to appear
19. Garey, M., Johnson, D.: Computers and Intractability: a guide to the theory of NP-completeness. W. H, Freeman and Company (1979)
20. Gerards, A.: Matching. In: M.B. et al., (ed.) Handbooks in OR and MS Vol. 7, chap. 3, pp. 135–224. Cambridge University Press (1995)
21. Hemaspaandra, E., Hemaspaandra, L., Schnoor, H.: A control dichotomy for pure scoring rules. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence, pp. 712–720 (2014)
22. Hemaspaandra, E., Schnoor, H.: Dichotomy for pure scoring rules under manipulative electoral actions. In: Proceedings of the 22nd European Conference on Artificial Intelligence, pp. 1071–1079 (2016)
23. Karp, R.: Reducibility among combinatorial problems. In: Proceedings of Symposium on Complexity of Computer Computations, pp. 85–103 (1972)
24. Kurihara, T.: Axiomatic characterisations of the basic best-worst rule. Econ. Lett. **172**, 19–22 (2018)
25. Lin, A.: The complexity of manipulating k-approval elections. In: Proceedings of the 3rd International Conference on Agents and Artificial Intelligence, pp. 212–218 (2011)
26. Loreggia, A., Narodytska, N., Rossi, F., Venable, K., Walsh, T.: Controlling elections by replacing candidates or votes. In: Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems, pp. 1737–1738 (2015)
27. Mirrlees, J.: An exploration in the theory of optimum income taxation. Rev. Econ. Stud. **38**(2), 175–208 (1971)
28. Mulmuley, K., Vazirani, U., Vazirani, V.: Matching is as easy as matrix inversion. Combinatorica **7**(1), 105–113 (1987)
29. Neveling, M., Rothe, J., Weishaupt, R.: The possible winner problem with uncertain weights revisited. In: Proceedings of the 23rd International Symposium on Fundamentals of Computation Theory, pp. 399–412 (2021)
30. Papadimitriou, C., Yannakakis, M.: The complexity of restricted spanning tree problems. J. ACM **29**(2), 285–309 (1982)
31. Tutte, W.: A short proof of the factor theorem for finite graphs. Can. J. Math. **6**, 347–352 (1954)
32. Walsh, T.: Uncertainty in preference elicitation and aggregation. In: Proceedings of the 22nd National Conference on Artificial Intelligence, pp. 3–8 (2007)
33. Xia, L.: Computing the margin of victory for various voting rules. In: Proceedings of the 12th ACM Conference on Electronic Commerce, pp. 982–999 (2012)
34. Yang, Y., Shrestha, Y., Guo, J.: On the complexity of bribery with distance restrictions. Theor. Comput. Sci. **760**, 55–71 (2019)