Manipulation of Opinion Polls to Influence Iterative Elections

Dorothea Baumeister, Ann-Kathrin Selker
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
{d.baumeister,ann-kathrin.selker}@uni-duesseldorf.de

Anaëlle Wilczynski
Technical University of Munich
Munich, Germany
anaelle.wilczynski@in.tum.de

ABSTRACT

In classical elections, voters only submit their ballot once, whereas in iterative voting, the ballots may be changed iteratively. Following the work by Wilczynski [20], we consider the case where a social network represents an underlying structure between the voters, meaning that each voter can see her neighbors' ballots. In addition, there is a polling agency, which publicly announces the result for the initial vote. This paper investigates the manipulative power of the polling agency. Previously, Wilczynski [20] studied constructive manipulation for the plurality rule. We introduce destructive manipulation and extend the study to the veto rule. Several restricted variants are considered with respect to their parameterized complexity. The theoretical results are complemented by experiments using different heuristics.

KEYWORDS

Iterative voting; Opinion polls; Manipulation; Complexity theory

ACM Reference Format:
Dorothea Baumeister, Ann-Kathrin Selker and Anaëlle Wilczynski. 2020. Manipulation of Opinion Polls to Influence Iterative Elections. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 9 pages.

1 INTRODUCTION

In the field of computational social choice there have been a lot of studies on elections, see the book by Brandt et al. [2]. The usual assumption is that voters once submit their ballot and then the winner is determined. This assumption completely neglects the reasoning about how the voters come to their individual decision. Especially in the ages of digital democracy, opinion polls may be executed efficiently and also repeatedly, which may lead the voters to strategically think about their ballot. We focus on iterative elections where voters can update (i.e., manipulate) their ballots. Following the seminal paper by Bartholdi III et al. [1], the issue of manipulation through strategic voting has been studied intensively in the computational social choice context. The most common, but often criticized, assumption is that the manipulator has complete knowledge over all ballots. There are different approaches to tackle this issue, for example the study of incomplete information settings like the possible winner problem introduced by Konczak and Lang [8]. However, we assume that voters only have partial knowledge about the other ballots. They have two sources of information. The first one is the result of some opinion poll, while the second one is the information they get from their neighbors in a social network. We then assume that every voter can update her ballot with respect to this information. In this process, the opinion poll is critical, and hence the polling agency has a lot of power. In this paper, we investigate the manipulative power of the polling agency with respect to different situations. In contrast to the manipulation by the voters, it is reasonable to assume complete information for the polling agency, since it collects the votes.

Wilczynski [20] recently introduced the problem of constructive manipulation, where the polling agency tries to make some distinguished candidate win by announcing some strategic opinion poll. The only condition the opinion poll has to satisfy is that no voter may directly detect the manipulation since the poll contradicts with her actual information. Formally this is modeled through a likelihood condition. The influence of opinion polls on the behavior of voters has also been studied by Reijngoud and Endriss [16] and Endriss et al. [3]. Their focus is on the strategic response of voters to different types of information communicated by the opinion polls, without an underlying social network, whereas we focus on the manipulative actions of the opinion poll itself. Our work analyzes poll manipulation in the setting of iterative voting, a widely studied topic in social choice (see Meir [10] for a recent survey), where voters can successively change their ballot in a strategic way. In this context, Sina et al. [18] have previously investigated election control in the presence of a social network. However, they focus on manipulation by an external agent who can add or remove links in the social network whereas we study manipulation of the initial scores communicated by the polling agency.

Recently, many works have investigated voting where voters are embedded in a social network. Tsang and Larson [19] analyze the consequences, in the strategic behavior of voters, of inferring the outcome of the election from the votes of neighbors in a social network. Alternatively, Gourvès et al. [7] study how voters can manipulate by coalitions which come from a social network structure. However, both works do not consider election control questions. Our work is also related to the study of opinion diffusion in graphs. Faliszewski et al. [4] study the effects of campaigning for manipulating election outcomes in an opinion diffusion process with voter clusters. In a similar context, Wilder and Vorobeychik [21, 22] investigate the game-theoretic properties of a game where an attacker tries to influence the election outcome by diffusing fake news in a social network and a defender aims to limit their impact.

We extend the study by Wilczynski [20] in several different ways. First, the current results are restricted to the use of the plurality rule, and we will also consider the veto rule. Although the rules are very similar, the results differ for restricted cases. Second, we introduce a destructive variant, where the opinion poll aims to prevent the victory of some designated candidate by announcing a strategic poll. Third, we analyze this problem for different parameters

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

and with respect to various distance restrictions. This is particularly important when considering real-world problems. Usually, the voters have some rough idea about how the opinion poll will look like, so the announcement should not deviate too much from the original outcome. We show that the corresponding decision problems are NP-hard for acyclic networks under both plurality and veto, and in P for empty networks under veto, whereas under plurality, the exact complexity for empty networks depends on the restriction model. Furthermore, we prove that all decision problems are tractable for a small number of candidates. Additionally, we design efficient heuristics for the manipulation of both voting rules and compare the manipulation results in our experimental section.

2 POLL-CONFIDENT ITERATIVE MODEL

We first describe the poll-confident iterative voting model.

2.1 Basic notations

Let $N = \{1, \dots, n\}$ be a set of n agents (or voters) and $M = \{x_1, \dots, x_m\}$ a set of m candidates. Each voter i has strict ordinal preferences over the candidates, represented by a linear order $\succ_i$ over M. The preference profile is denoted by $\succ = (\succ_1, \dots, \succ_n)$. Let $\overrightarrow{M}$ (respectively, $\overleftarrow{M}$) be a shorthand for $x_1 \succ \cdots \succ x_m$ (respectively, $x_m \succ \cdots \succ x_1$), and let $N_{x \succ y}$ denote the set of voters who prefer candidate x to candidate y, i.e., $N_{x \succ y} = \{i \in N : x \succ_i y\}$. The winner of the election is determined by a voting rule $\mathcal{F}$. We focus on single-winner elections and use a deterministic tie-breaking rule based on a linear order $\rhd$ over the candidates, in case of ties.

In this article, we focus on two voting rules, namely plurality, denoted by $\mathcal{F}_P$, and veto, denoted by $\mathcal{F}_V$. Under both voting rules, each voter i is asked to submit a ballot $\mathbf{b}_i \in M$ corresponding to a single candidate, i.e., voter i approves candidate $\mathbf{b}_i$ under plurality, whereas voter i vetoes candidate $\mathbf{b}_i$ under veto. A profile of ballots is said to be *truthful* if each agent submits as a ballot her most preferred candidate under plurality and her least preferred candidate under veto. Given a profile of ballots $\mathbf{b} \in M^n$, the score $\mathbf{s_b}(x)$ of each candidate x is computed as follows: $\mathbf{s_b}(x) = |\{i \in N : \mathbf{b}_i = x\}|$. Then the winner under plurality $\mathcal{F}_P(\mathbf{b})$ maximizes the number of approvals, i.e., $\mathcal{F}_P(\mathbf{b}) \in \arg\max_{x \in M} \mathbf{s_b}(x)$, whereas the winner under veto $\mathcal{F}_V(\mathbf{b})$ minimizes the number of vetoes, i.e., $\mathcal{F}_V(\mathbf{b}) \in \arg\min_{x \in M} \mathbf{s_b}(x)$. For the sake of simplicity, we may use $\mathcal{F}(\mathbf{s})$ to denote the winner of a profile of ballots whose score function corresponds to $\mathbf{s}$.

We consider a strategic game called *iterative voting* [10] where, starting from an initial voting profile, agents can successively deviate from their current submitted ballot in order to get a better outcome at the next election. The strategy profile at step t is denoted by $\mathbf{b}^t$. We assume that the initial profile $\mathbf{b}^0$ is truthful, indeed the agents do not have any information to enable them to deviate yet. A single voter is assumed to deviate between two consecutive steps. In case of several voters having incentive to deviate at the same step, one of them is arbitrarily chosen, unless a particular *turn function* τ is specified for choosing the deviator. Note that the choice of the turn function might influence the election outcome.

In the classical iterative voting setting, there is common knowledge of the current strategy profile (or at least the associated scores). For realistic reasons we consider, following Wilczynski [20], that the voters only get partial information about the current strategy profile which is determined by a *social network* and an *opinion poll* and thus which can be biased. More precisely, we assume that the agents are embedded in a social network represented by a directed graph $G = (N, E)$ such that for each arc $(i, j) \in E$, agent i is able to observe the current ballot of agent j. The social network is said to be *empty* if $E = \emptyset$ and *acyclic* if there is no directed cycle in G. The set of agents that a given agent i can observe is denoted by $\Gamma(i) := \{j \in N : (i, j) \in E\} \cup \{i\}$. For a voting profile $\mathbf{b}$, the score of candidate x that agent i is able to observe is denoted by $\mathbf{s}^i_{\mathbf{b}}(x) = |\{j \in \Gamma(i) : \mathbf{b}_j = x\}|$. Moreover, as a prior information, the agents know the scores of the initial profile, given by a polling agency, through the vector of scores $\Delta = (\Delta_1, \dots, \Delta_m)$ where Δ_j stands for the score of candidate x_j which is announced by the polling agency. By abuse of notation, we may also use $\Delta(x)$ for the announced score of candidate x. To summarize, an instance of the poll-confident iterative voting model is a tuple $\mathcal{I} = (N, M, \succ, G, \rhd, \tau)$.

2.2 Manipulation by voters

The manipulation moves of voters are conditioned by the information they get, which is determined by the deviations that they are able to observe. Each agent i has a specific belief regarding the scores of the strategy profile at step t which is given by a believed score vector $B^t_i = (B^t_i(x_1), \dots, B^t_i(x_m))$. The voters trust the results communicated by the polling agency, and thus $B^0_i = \Delta$ for every agent $i \in N$. The believed score vector for both the plurality and veto rules is updated at each step as follows.

Definition 2.1 (Score Belief Update). At step $t+1$, after the deviation of an agent j from ballot $\mathbf{b}^t_j = x$ to ballot $\mathbf{b}^{t+1}_j = y$ at step t, the score of candidate z that agent i believes is given by

$$B^{t+1}_i(z) = \begin{cases} B^t_i(z) - 1 & \text{if } z = x \text{ and } j \in \Gamma(i) \\ B^t_i(z) + 1 & \text{if } z = y \text{ and } j \in \Gamma(i) \\ B^t_i(z) & \text{otherwise} \end{cases}$$

According to the belief of agent i, the *current believed winner* at step t is candidate $\mathcal{F}(B^t_i)$. We assume that the voters only deviate when they believe that they are pivotal, i.e., they believe that their deviation changes the winner of the election.[1] In such a context, identifying the *potential winners* which are the candidates that an agent can make win is essential. However, this mainly depends on the belief of the agents.

Definition 2.2 (Potential winner). A candidate x is a potential winner for agent i at step t, i.e., $x \in PW^t_i$, if, without considering the current ballot $\mathbf{b}^t_i$ of agent i, agent i believes that one more vote in favor of x under plurality or one more veto against another candidate under veto, will make candidate x the new winner.

Observe that the two voting rules under consideration are not symmetric with respect to the set of potential winners. Under plurality, for a given agent, there may be several potential winners other than the current believed winner and it seems rational that

[1] Introducing thresholds to relax the assumption of strict pivot, in the spirit of the works of Meir et al. [11], Obraztsova et al. [14] or Wilczynski [20], also makes sense. We do not make such an assumption for the sake of simplicity and to especially focus on the impact of the social environment of the agents (social network, opinion poll).

the agent will choose to favor the candidate that she prefers. In contrast, under veto, vetoing candidates other than the current believed winner would not produce any direct change according to the belief of an agent. Therefore, there is only one potential winner other than the believed winner, i.e., the one which becomes the new winner after one more veto for the current believed winner. This difference strongly conditions the dynamics of deviations that we consider for each voting rule. While best response deviations are considered under plurality, deviations consisting of vetoing the current believed winner are considered under veto.

Definition 2.3 (Best response deviation (plurality)). A voter i deviates from ballot $\mathbf{b}_i^t$ to ballot $\mathbf{b}_i^{t+1} := y$ at step t following a best response if $y \in PW_i^t \setminus \{\mathcal{F}_P(B_i^t)\}$ and $y \succ_i z$ for any $z \in PW_i^t \setminus \{y\}$.

Definition 2.4 (Veto-winner deviation (veto)). A voter i deviates from ballot $\mathbf{b}_i^t$ to ballot $\mathbf{b}_i^{t+1}$ at step t following a veto-winner deviation if $\mathbf{b}_i^{t+1} = \mathcal{F}_V(B_i^t)$ and $\mathcal{F}_V(B_i^{t+1}) \succ_i \mathcal{F}_V(B_i^t)$.

Both best response and veto-winner dynamics are proved to converge under plurality and veto, respectively, when the social network is complete, i.e., the scores of the current strategy profile are common knowledge [9, 12, 17]. Moreover, convergence is also satisfied when the social network is acyclic or transitive [20].

When the dynamics converges, it reaches a stable state where no voter has an incentive to deviate according to her belief. In this article, we are interested in the identity of the *iterative winner*, i.e., the winner of the stable state reached by the dynamics. We aim to analyze how the polling agency can influence the outcome of the dynamics by manipulating the scores of the initial poll which is communicated to the voters.

2.3 Manipulation by the polling agency

In order for the polling agency not to be detected manipulating the initial poll, it is important that the manipulated poll meets the following criterion, first introduced by Wilczynski [20].

Definition 2.5 (Likelihood condition). A polling vector Δ is *plausible* if $n = \sum_{j=1}^{m} \Delta_j$ and it gives for each candidate at least the highest score that an agent can observe, i.e., $\Delta_j \geq \max_{i \in N} \mathbf{s}^i_{\mathbf{b}^0}(x_j)$.

Note that checking whether a poll satisfies the likelihood condition is possible in polynomial time.

In this paper, we will study whether the polling agency is able to influence the outcome of the iterative election via the following decision problem for voting rule $\mathcal{F} \in$ {plurality, veto}.

$\mathcal{F}$-{Constructive / Destructive}-Election-Enforcing:
Instance: Instance $(N, M, \succ, G, \rhd, \tau)$, target candidate p.
Question: Can the polling agency announce a plausible poll Δ so that p {is / is not} the iterative winner?

In reality, the likelihood condition as shown in Definition 2.5 might be too weak and give the polling agency too much power. Especially in cases where some organizations keep an eye on the polling agency or where there have been recent election results, the polling agency should not announce a poll that extremely differs from the correct poll. The motivation is similar to the one presented by Obraztsova and Elkind [13] for optimal manipulation in voting. They propose to bound the manipulative action by some distance, which makes manipulation possibly harder to detect in real-world instances. Therefore, we introduce the following distance-restricted problems, where the *Manhattan distance* between two m-vectors Δ and Δ' is defined as $dist(\Delta, \Delta') = \sum_{i=1}^{m} |\Delta_i - \Delta'_i|$.

$\mathcal{F}$-Poll-Restricted-{Constr. / Destr.}-Election-Enforcing:
Instance: $(N, M, \succ, G, \rhd, \tau)$, target candidate p, distance d.
Question: Can the polling agency announce a plausible poll Δ so that p {is / is not} the iterative winner and $dist(\Delta, \mathbf{s}(\mathbf{b}^0)) \leq d$?

Example 2.6. Let us consider an instance with 6 voters and 4 candidates where $G = (N, \{(1, 2), (3, 4)\})$ and $x_3 \rhd x_2 \rhd x_1 \rhd x_4$. The preferences are as follows.

$$\begin{aligned} 1, 2, 3 &: \quad x_1 \succ x_2 \succ x_3 \succ x_4 \\ 4 &: \quad x_3 \succ x_1 \succ x_4 \succ x_2 \\ 5, 6 &: \quad x_1 \succ x_4 \succ x_2 \succ x_3 \end{aligned}$$

Under veto, the truthful winner is x_1. If $\Delta = \mathbf{s}(\mathbf{b}^0)$, there is no deviation: x_1 is the top candidate of all voters except voter 4, but she cannot deviate, otherwise her worst candidate x_2 will be elected. Suppose that the polling agency aims to avoid the election of x_1. By the likelihood condition, it must hold that $\Delta(x_2) \geq 1$, $\Delta(x_3) \geq 1$ and $\Delta(x_4) \geq 2$. If $\Delta = (0, 3, 1, 2)$, then voter 4 believes that x_1 is the winner and x_3 a potential winner. She thinks that she can safely deviate without making x_2 elected, so she deviates for vetoing x_1 and makes x_2 the new winner. Voter 3 observes this deviation and then deviates to veto x_3 that she believes to be the winner. However, x_2 remains the real iterative winner. This is the only successful poll manipulation, thus if the distance to the truthful scores is limited to less than 4, there is no poll-restricted manipulation.

Voters and their current votes are visible for their neighbors. Especially when candidates can be positioned on a spectrum, voters might be inclined to vote for candidates that do not clash with their preference order, either for ideological reasons or because they are worried about what their friends might think of them. Therefore, we introduce the following problem, where the distance between a ballot submitted by agent i approving (resp., vetoing) candidate x under $\mathcal{F}_P$ (resp., $\mathcal{F}_V$) and her truthful ballot is given by the number of swaps between two consecutive candidates in ranking $\succ_i$ that are necessary to put x at the top (resp., bottom) of $\succ_i$.

$\mathcal{F}$-Voter-Restricted-{Constr./Destr.}-Election-Enforcing:
Instance: $(N, M, \succ, G, \rhd, \tau)$, target candidate p, distance d.
Question: Can the polling agency announce a plausible poll Δ so that p {is / is not} the iterative winner when voters can only submit a ballot at distance at most d from their truthful ballot?

We assume our reader to be familiar with the complexity classes P, NP, para-NP, FPT, and the W-hierarchy, as well as the concepts of polynomial-time many-one reducibility and fpt-reducibility (see, e.g., Papadimitriou [15] and Flum and Grohe [5]).

The winner determination for the considered iterative elections might exceed polynomial time, even for converging elections and acyclic networks. However, for each of the constructed instances in our hardness proofs, the winner determination is possible in polynomial time, therefore proving the intractability of the problems does not depend on the complexity of the winner determination.

3 MANIPULATING POLL PLURALITY SCORES

In this section, we investigate the election enforcing problem under the plurality rule and best response dynamics. It turns out that most of the variants of the problem are intractable, except when the number of candidates is relatively small.

All hardness results in this section hold even when the social network is acyclic and the turn function is constructed so that each voter changes her vote at most once. We use the following NP-complete decision problem to prove our first result. Hitting Set asks—given a universe $X = \{x_1, \dots, x_m\}$, a collection $\mathcal{S} = \{S_1, \dots, S_n\}$ of subsets over X, and a nonnegative integer k—whether there exists a hitting set of size k, i.e., a set $X' \subseteq X$ of size k such that $S \cap X' \neq \emptyset$ for all $S \in \mathcal{S}$. Note that Hitting Set is also W[2]-complete when parameterized by the size of the hitting set k.

Theorem 3.1. *$\mathcal{F}_P$-Destr.-Election-Enforcing is* NP*-hard.*

Sketch of proof. Let $(X, \mathcal{S}, k)$ be an instance of Hitting Set where $X = \{x_1, \dots, x_m\}$ and $\mathcal{S} = \{S_1, \dots, S_n\}$. Without loss of generality, we assume that $k > 3$. Construct an instance of $\mathcal{F}_P$-Destructive-Election-Enforcing as follows:

Let $X \cup Y \cup \{p, z\}$ be the set of candidates, where p is the target candidate and $Y = \{y_0, y_1, \dots y_m\}$. The table below shows the preferences of the voters, partitioned into parts A to F.

Part	Name	Preference	for
A	a_1 :	$y_0 > p > \overleftarrow{X} > \dots > z$	
	a_2 :	$y_0 > z > \overrightarrow{X} > \dots > p$	
	a_3 :	$y_0 > p > z > \overrightarrow{X} > \dots$	
	a_4 :	$p > z > \overrightarrow{X} > \dots$	
B	b_i :	$y_i > x_i > z > p > \overrightarrow{X \setminus \{x_i\}} > \dots$	$1 \le i \le m$
C	c_j	$p > z > \overrightarrow{X} > \dots$	$1 \le j \le n$
D	d_j :	$z > p > \overrightarrow{X} > \dots$	$1 \le j \le n$
E	e_j :	$p > z > \overrightarrow{X} > \dots$	$1 \le j \le k$
F	$f_{i,j}$:	$x_i > z > p > \overrightarrow{X \setminus \{x_i\}} > \dots$	$1 \le i \le m, 1 \le j \le n$

The complete set of arcs in the social network is as follows. There is an arc from a_2 and a_3 to a_1, and an arc from each voter in B to a_2, a_3, and a_4. Each voter c_j in C sees the voters in B corresponding to the variables in S_j, and additionally has an arc to a_2. All voters in D see each voter in B and additionally the voter a_3. The voters in E each have arcs to each voter in B, C, and D, and additionally see the voters a_2 and a_3. Finally, each voter $f_{i,j}$ has an edge to a_1, a_2, and a_3, and each voter $f_{i,n}$ has arcs to the voters $f_{i,1}$ to $f_{i,n-1}$.

We base the turn function on the order $\overrightarrow{A} > \overrightarrow{B} > \overrightarrow{C} > \overrightarrow{D} > \overrightarrow{E} > \overrightarrow{F}$ and use the order $z \rhd p \rhd \overrightarrow{X} \rhd \dots$ for tie-breaking.

The following table shows the correct initial poll Δ the polling agency should announce (line 1), and the minimum number of points the polling agency has to give each candidate in a manipulated poll due to the likelihood condition in Definition 2.5 (line 2). All in all, the polling agency has a contingent of (only) k points.

	p	z	$x \in X'$	$x \in X \setminus X'$	$y \in Y$
Δ	$n+k+1$	n	n	n	3/1
min	$n+1$	n	n	n	3/1
Δ'	$n+1$	n	$n+1$	n	3/1
final	$n+1$	$n+k+1$	$n+1$	n	2/$\le$ 1

We claim that there is a hitting set of size k, i.e., a set $X' \subseteq X$ of size k so that $X' \cap S \neq \emptyset$ for each $S \in \mathcal{S}$, iff the polling agency can publish a plausible Δ' that results in p not winning the election.

($\Rightarrow$) Suppose $(X, \mathcal{S})$ is a yes-instance of Hitting Set and let X' be a hitting set of size k. The polling agency can publish the manipulated initial poll Δ' as described in the table.

The election then proceeds as follows. Voter a_2 changes her ballot to z, whereas the remaining voters in A cannot achieve a better outcome than the current winner p. The voters in B observe the change from a_2 to z and are now convinced that z is winning due to tie-breaking. The k voters corresponding to an $x_i \in X'$ change their ballot to x_i to give the respective x_i the missing point to win, whereas the other $m - k$ voters in B do not think they can change the outcome to their advantage. The voters in C observe the changes in A and B and—since X' is a hitting set—each sees (at least) one x_i gaining a point, so they react by collectively changing their ballot to z to make z the plurality winner by tie-breaking. The voters in D also think an x_i is currently winning by one point after observing the voters in B, and collectively switch to p. After observing all changes made up to this point, the voters in E switch to z—they observe p winning and losing exactly n points for a total of $n+1$ points, z gaining $n+1$ and losing n points for a total of $n+1$ points, and the $x \in X'$ gaining one point for a total of $n+2$ points. Finally, none of the voters in F change their ballot because they all see z winning and are unable to reach a more favorable result.

All in all, z wins the election. The final scores can be seen in the last line of the table.

($\Leftarrow$) Suppose that each $X' \subseteq X$ of size at most k is disjoint to an $S \in \mathcal{S}$. That means that it is not possible to convince *all* voters in C to change their ballot from p to another candidate. Due to the space constraints, we omit detailed explanations for each possible manipulated poll. However, regardless of how the manipulated poll is set up, p remains the winner of the election. ❑

Note that the above proof also shows that plurality election enforcing is W[2]-hard for both the poll-restricted and unrestricted constructive variant as well as for the destructive variants when parameterized by the distance between the original and the manipulated initial poll. In the constructive cases, z is the target candidate.

The following theorem shows that even a highly restricted acyclic social network is sufficient to show hardness of manipulation for the restricted problem variants. We use a network where the longest path is of length 1 and—in the voter-restricted problem variant—where the maximum outdegree of a node is 6. Furthermore, in the voter-restricted variant, the voters are only inclined to vote for their two most preferred candidates.

Theorem 3.2. *(1) $\mathcal{F}_P$-Voter-Restricted-{Constr., Destr.}-Election-Enforcing is* para-NP*-hard when parameterized by the number of swaps and the length of a longest path in the network.*
(2) $\mathcal{F}_P$-Poll-Restricted-{Constructive, Destructive}-Election-Enforcing is para-NP*-hard when parameterized by the length of a longest path and the maximum outdegree of the social network.*

Due to space constraints, we omit the proof of this theorem.

Next, we investigate whether manipulation becomes easy if we restrict our allowed instances even further.

Proposition 3.3. *If the winner determination is possible in polynomial time, then $\mathcal{F}_P$-{Unrestricted, Poll-Restricted, Voter-Restricted}-{Constr., Destr.}-Election-Enforcing is in* FPT *when parameterized by the number of candidates m.*

Proof. Construct plausible initial polls for each subset $M' \subseteq M$ of the m candidates in the following way. Set M' corresponds to the initial set of potential winners. For each subset $M^* \subseteq M'$, create a plausible poll Δ' if possible so that $\Delta'(x) = \alpha$ for $x \in M^*$, $\Delta'(x) = \alpha - 1$ for $x \in M' \setminus M^*$, and $\Delta'(x) < \alpha - 2$ for $x \in M \setminus M'$, where α is an integer that can differ from poll to poll. For fixed M' and M^*, each poll meeting these requirements will yield the same election result regardless of the value of α and the exact scores of the candidates in $M \setminus M'$. Note that constructing such a poll (resp., ascertaining that a plausible poll satisfying the requirements does not exist) is possible in polynomial time for all problem variants, as the value of α is bounded by the number of voters. Since we construct at most $2^m \cdot 2^m$ initial polls and testing whether they fulfill our requirements and yield the desired election outcome is possible in polynomial time for each poll, our algorithm is an fpt-algorithm when parameterized by m. ❑

A possible further restriction for the network is an empty graph, i.e., a network where voters only rely on the opinion poll. However, this does not seem to simplify the constructive manipulation problem: it can be necessary to include arbitrary many candidates in the initial set of potential winners. While we conjecture that this problem remains NP-hard for all our considered variants for an empty graph, we can only prove this for the poll-restricted variant, leaving the exact complexity open for the unrestricted and voter-restricted variants. The proof uses a reduction from an NP-complete restricted version of the problem X3C [6], where we are given a universe $X = \{x_1, \ldots, x_{3m}\}$ and a collection $\mathcal{S} = \{S_1, \ldots, S_{3m}\}$, $S_j \subseteq X$, $|S_j| = 3$, so that each $x \in X$ is contained in exactly three sets S_j, and we ask whether there exists an exact cover $\mathcal{S}' \subseteq \mathcal{S}$ of size m so that the union of all sets in $\mathcal{S}'$ equals X.

Theorem 3.4. *$\mathcal{F}_P$-Poll-Restricted-Constructive-Election-Enforcing remains NP-hard when the social network is empty.*

Sketch of proof. Let $(X, \mathcal{S})$ be an instance of X3C where $X = \{x_1, \ldots, x_{3m}\}$, $\mathcal{S} = \{S_1, \ldots, S_{3m}\}$ and $S_j = \{x_{j,1}, x_{j,2}, x_{j,3}\}$. Construct an instance of $\mathcal{F}_P$-Poll-Restricted-Constr.-Election-Enforcing as follows. Let $X \cup \mathcal{S} \cup \{w, p, z\}$ be the set of candidates, where p is the target candidate. The following table shows the preferences of the voters, partitioned into the parts A to G.

Part	Name	Preference	for
A	a_i :	$w > x_i > z > \overrightarrow{\mathcal{S}} > \overrightarrow{X \setminus \{x_i\}} > p$	$1 \le i \le 3m$
B	$b_{j,1}$:	$x_{j,1} > S_j > z > w > \overrightarrow{\mathcal{S} \setminus \{S_j\}} > \overrightarrow{X \setminus \{x_{j,1}\}} > p$	$1 \le j \le 3m$
	$b_{j,2}$:	$x_{j,2} > S_j > z > w > \overrightarrow{\mathcal{S} \setminus \{S_j\}} > \overrightarrow{X \setminus \{x_{j,2}\}} > p$	$1 \le j \le 3m$
	$b_{j,3}$:	$x_{j,3} > S_j > z > w > \overrightarrow{\mathcal{S} \setminus \{S_j\}} > \overrightarrow{X \setminus \{x_{j,3}\}} > p$	$1 \le j \le 3m$
C	c_i	$x_i > z > w > \overrightarrow{\mathcal{S}} > \overrightarrow{X} > p$	$1 \le i \le 3m$
D	d_k :	$p > z > \overleftarrow{X} > \overleftarrow{\mathcal{S}} > w$	$1 \le k \le 5$
E	e_k :	$z > w > \overrightarrow{\mathcal{S}} > \overrightarrow{X} > p$	$1 \le k \le 3$
F	f_k :	$w > z > \overrightarrow{\mathcal{S}} > \overrightarrow{X} > p$	$1 \le k \le 4$
G	g_j :	$S_j > z > w > \overrightarrow{\mathcal{S} \setminus S_j} > \overrightarrow{X} > p$	$1 \le j \le 3m$

We use the tie-breaking order $z \rhd w \rhd \overrightarrow{\mathcal{S}} \rhd \overrightarrow{X} \rhd p$ and a maximum allowed distance between the correct and the manipulated initial poll of $3m + 1$.

Note that w is currently winning with a score of $3m + 4$, whereas p only has a score of 5 and cannot gain any points regardless of the broadcasted initial poll because each voter but the ones in D rank p last. The only way to make p win the election is therefore convincing the voters approving of w to approve other candidates, but in a way so that p does not lose approvals and so that each other candidate has a final score of at most 4 due to the tie-breaking.

($\Rightarrow$) Suppose $(X, \mathcal{S})$ is a yes-instance of X3C and let $\mathcal{S}'$ be the exact cover of size m. Assign $3m$ points from w in a way so that each $S \in \mathcal{S}'$ gains three points, and assign one point from p to z. This results in each candidate but the $S \notin \mathcal{S}'$ to have a score of 4 so that the voters think z is winning the election. Then the voters in A collectively change their ballot from w to the respective x_i and the voters in B that correspond to the $S_j \in \mathcal{S}'$ change their vote to S_j. Since $\mathcal{S}'$ is an exact cover, each of the $x \in X$ gain and lose exactly one point. None of the remaining voters change their ballot because they cannot improve the election result. Therefore, in the final result, p has kept a score of 5 whereas each other candidate has a score of at most 4, resulting in p winning the election.

($\Leftarrow$) Suppose that there does not exist an exact cover $\mathcal{S}'$ of size at most m and recall that p cannot gain any points. Then for each plausible initial poll, either w does not lose at least $3m$ points, there is an x that is not covered by $\mathcal{S}'$ and therefore receives a final score of at least 5, or the voters in D collectively change their ballot from p to another candidate, all resulting in p losing the election. We omit the details due to space constraints. ❑

In the destructive case, the manipulation problem becomes easy, at least for the unrestricted and poll-restricted variants. Note that without arcs, winner determination is possible in polynomial time.

Proposition 3.5. *$\mathcal{F}_P$-{Unrestricted, Poll-Restricted}-Destr.-Election-Enforcing is in* P *when the social network is empty.*

Our algorithm that solves the election enforcing instance creates plausible initial polls similar to the way of the algorithm in Proposition 3.3, but only uses pairs of candidates (x, y) as potential winners so that we only construct a polynomial number of polls. Nevertheless, this suffices in an empty graph because for each initial poll, at most half of the voters still vote for the target candidate p. If none of the constructed initial polls yield another winner than p, then it is not possible to make p lose the election. However, for the voter-restricted variant, this proof does not work anymore because we do not know how many voters will change their ballot even when at least half of them prefer x to y. In fact, it can be necessary to include up to all candidates in the initial set of potential winners. Therefore, we conjecture that $\mathcal{F}_P$-Voter-Restricted-Destructive-Election-Enforcing remains NP-hard even when the social network is empty.

4 MANIPULATING POLL VETO SCORES

In this section, we investigate the problems of election enforcing for the polling agency under the veto rule and veto-winner dynamics. Due to the nature of the veto-winner deviations, the results differ a bit from those under the plurality rule. In particular, for each agent at each step, the set of potential winners other than the current believed winner is composed of at most one candidate.

Let us denote by V_x the set of voters vetoing candidate x at the initial truthful profile, i.e., $V_x := \{i \in N : \mathbf{b}_i^0 = x\}$. Observe that if the number of vetoes for candidate x announced by polling vector Δ is not sufficiently large, then the voters in V_x will not deviate at step 0, because they would think that they make their worst candidate x win by removing their veto against it, i.e., $PW_i^0 = \{\mathcal{F}_V(\Delta), x\}$

for $i \in V_x$. Therefore, the global idea of the manipulation of the polling agency under veto is to announce enough vetoes against a candidate whose vetoers must deviate. Let us denote by PW the second best candidate announced by Δ with a score difference of one with the announced winner (advantage w.r.t $\triangleright$ included), i.e., PW is a potential winner for all voters $i \in V_x$ such that $x \notin PW_i^0$. The problem of enforcing the election of a given candidate p (respectively, ensuring candidate p does not win the election) is intractable even if the social network is relatively sparse. The proof of the following theorem uses a social network where the longest path is only of length 2 (respectively, of length 1 for the voter-restricted variant). Furthermore, in the voter-restricted variant, the voters can even only veto their two least preferred candidates.

Theorem 4.1. *$\mathcal{F}_V$-{Unrestricted, Poll-Restricted, Voter-Restricted}-{Constructive, Destructive}-Election-Enforcing is para-NP-hard when parameterized by the length of the longest path and—for the voter-restricted variant—by the number of swaps.*

Sketch of proof. We first prove that $\mathcal{F}_V$-Destructive-Election-Enforcing is NP-hard even when the longest path is of length 2. The hardness of the poll-restricted variant with the same parameters and the hardness of the constructive variants immediately follow. We just need to set the maximum Manhattan distance between the original and the manipulated initial poll to at least $6m + 2$ and—in the constructive variants—the target candidate to z. After the proof, we give a slight modification for the voter-restricted variant that reduces the length of the longest path to 1.

We reduce from X3C. Let $(X, \mathcal{S})$ be an instance of X3C where $X = \{x_1, \ldots, x_{3m}\}$ and $\mathcal{S} = \{S_1, \ldots, S_{3m}\}$ so that $S_j \subseteq X$, $|S_j| = 3$, and each $x \in X$ is contained in exactly three sets $S \in \mathcal{S}$. Construct an instance of $\mathcal{F}_V$-Destructive-Election-Enforcing as follows.

Let $\mathcal{S} \cup \{p, z, d_1, d_2\}$ be the set of candidates, where p is the target candidate. The table below shows the preferences of the voters, partitioned into parts A to G.

Part	Name	Preference	for
A	a_k :	$d_1 \succ p \succ \cdots \succ z \succ d_2$	$1 \leq k \leq m$
	a_{m+1} :	$d_1 \succ z \succ \overrightarrow{\mathcal{S}} \succ p \succ d_2$	
B	b_j :	$\cdots \succ \overrightarrow{\mathcal{S} \setminus \{S_j\}} \succ p \succ z \succ S_j$	$1 \leq j \leq 3m$
C	c_i	$\cdots \succ \overrightarrow{\mathcal{S}} \succ p \succ z$	$1 \leq i \leq 3m$
D	d_k :	$p \succ \cdots \succ d_2 \succ z$	$1 \leq k \leq m$
E	e_j :	$p \succ \cdots \succ z \succ d_2 \succ S_j$	$1 \leq j \leq 3m$
F	f_k :	$p \succ \cdots \succ z \succ d_1 \succ d_2$	$1 \leq k \leq 9m - 1$
G	g_k :	$p \succ \cdots \succ z \succ d_2 \succ d_1$	$1 \leq k \leq 10m$

The social network has the following set of arcs. Each voter c_i in C has an arc to a_1, a_{m+1}, and to each of the three voters in B that correspond to the sets S_j that contain x_i. Furthermore, voter d_m sees each of the other voters in D, the voters in E see a_1 and additionally e_{3m} sees each of the other voters in E, and the voters in F and G have an arc to each voter in A and additionally f_{9m-1} (resp., g_{10m}) has an arc to each of the other voters in F (resp. G).

We base the turn function on the order $\overrightarrow{A} > \overrightarrow{B} > \overrightarrow{C} > \overrightarrow{D} > \overrightarrow{E} > \overrightarrow{F} > \overrightarrow{G}$ and use the order $z \triangleright p \triangleright \overrightarrow{\mathcal{S}} \triangleright \ldots$ for tie-breaking.

The following table shows the correct initial poll Δ the polling agency should announce (line 1), and the minimum number of vetoes the polling agency has to give each candidate in a manipulated poll due to the likelihood condition in Definition 2.5 (line 2).

	p	z	$S \in \mathcal{S}'$	$S \notin \mathcal{S}$	d_1	d_2
Δ	0	$4m$	$3m + 1$	$3m + 1$	$10m$	$10m$
min	0	m	$3m$	$3m$	$10m$	$10m$
Δ'	$3m$	$3m$	$3m + 1$	$3m$	$10m$	$10m$
final	$3m$	$3m$	$3m$	$3m + 1$	$10m$	$9m$

($\Rightarrow$) Suppose $(X, \mathcal{S})$ is a yes-instance of X3C and let $\mathcal{S}' \subseteq \mathcal{S}$ be the exact cover of size m. The polling agency can publish the poll Δ' described in the table. The election then proceeds as follows. All voters think that z is the winner. Therefore, voters a_1 to a_m in A change their ballot to z to make p the winner of the election by tie-breaking. Voters b_j also want to hinder z from winning. However, they only veto z in the case that S_j is part of the exact cover, because otherwise the loss of a veto for S_j would result in S_j being the veto winner with only $3m - 1$ vetoes. Since $\mathcal{S}'$ is an exact cover, each voter in C observes z gaining two vetoes—one from a_1 and one from a voter in C—and reacts by vetoing p. This is possible because z now has enough vetoes to not become the veto winner after losing a veto. None of the voters in D, E, and F change their ballot. All in all, z gains $2m$ vetoes from the voters in A and B and loses $3m$ vetoes from the voters in C, resulting in z winning the election with $3m$ vetoes due to tie-breaking (see last line of the table).

($\Leftarrow$) Suppose that there is no exact cover $\mathcal{S}'$ of size at most m. Then, regardless of the initial poll, there is at least one voter in C who does not change her veto to p so that p does not obtain the necessary number of vetoes to lose the election. We omit the details due to space constraints.

For the voter-restricted variant, the depicted proof obviously works (for a maximum number of $3m + 2$ swaps), but we can even strengthen the result by tightening the parameters: Set the maximum number of swaps to 1, i.e., only allow the voters to veto their two least preferred candidates. Delete the arcs between parts E and A, F and A, as well as G and A. The resulting social network has a longest path of length 1. ❑

However, the manipulation problem can be solved efficiently if the number of candidates is small.

Proposition 4.2. *If the winner determination is possible in polynomial time, then $\mathcal{F}_V$-{Unrestricted, Poll-Restricted, Voter-Restricted}-{Constructive, Destructive}-Election-Enforcing is in FPT when parameterized by the number of candidates m.*

The proof works analogously to the proof of Proposition 3.3.

In contrast to plurality, manipulation is easy under veto when the social network is empty, even in the constructive case.

Proposition 4.3. *$\mathcal{F}_V$-Constructive-Election-Enforcing is solvable in polynomial time when the social network is empty.*

Sketch of proof. The idea of our algorithm is to communicate a polling vector Δ that makes the voters removing vetoes against p, or that prevents many deviations from agents vetoing other candidates. Since the set of potential winners other than the current believed winner is composed of at most one candidate, we try all the $O(m^2)$ combinations of pairs of distinct candidates (ω, y) such that ω is the announced winner of Δ and y is the other announced potential winner PW. For any pair of candidates (ω, y), we create a polling vector Δ which fulfills the likelihood condition (Def. 2.5) with a minimum number of vetoes, and then we add the minimum number of vetoes in order to get ω and y the winner and PW of Δ, respectively. The rest of available vetoes is distributed as follows:

- Case $p \notin \{\omega, y\}$: If p is not at least the second winning candidate in $\mathbf{b}^0$, then vetoes against p must be removed, so we add in Δ the minimum number of vetoes to p in order to "authorize" V_p, i.e., in order to have $p \notin PW_i^0$ (and thus $y \in PW_i^0$) for every $i \in V_p$. Otherwise, we test the two options: authorize V_p or not (still polynomial). With the remaining available vetoes, we "block", as much as possible, the deviation of voters V_x for all other candidates $x \notin \{\omega, y\}$ such that $|V_x \cap N_{\omega>y}| + \mathbb{1}_{\{p \triangleright x\}} < \mathbf{s}_{\mathbf{b}^*}(p)$ where $\mathbf{s}_{\mathbf{b}^*}(p) = |V_p|$ if V_p is not authorized or $\mathbf{s}_{\mathbf{b}^*}(p) = |V_p \cap N_{\omega>y}|$ otherwise. The goal is to avoid that the final score of another candidate is lower than the final score of p. For blocking voters V_x, we add the minimum number of vetoes to candidates other than x in order to make a voter $i \in V_x$ believe that $x \in PW_i^0$. If some available vetoes remain, we use them to authorize as much as possible the other voters V_z for all candidates $z \notin \{\omega, y\}$ such that $|V_z \cap N_{\omega>y}| + \mathbb{1}_{\{p \triangleright z\}} > |V_p|$, by choosing first the candidates which maximize $|V_z \cap N_{y>\omega}|$.
- Case $y = p$: V_p is already blocked, therefore no veto against p can be removed. To become the iterative winner, p must be the second best in $\mathbf{b}^0$ and the deviations must add enough vetoes against ω, which must be $\mathcal{F}_V(\mathbf{b}^0)$, while the deviations of V_x must be blocked if x could have a smaller score than p. Therefore, we block V_x for all candidates $x \notin \{\omega, y\}$ such that $|V_x \cap N_{\omega>y}| + \mathbb{1}_{\{p \triangleright x\}} < |V_p|$. Then, we authorize as much as possible the other voters V_z for all candidates $z \notin \{\omega, y\}$ such that $|V_z \cap N_{\omega>y}| + \mathbb{1}_{\{p \triangleright z\}} > |V_p|$ by choosing first the candidates which maximize $|V_z \cap N_{y>\omega}|$.
- Case $\omega = p$: The only possible deviations are vetoes against p and no veto against p can be removed. Therefore, it must hold that p is the actual winner, i.e., $p = \mathcal{F}_V(\mathbf{b}^0)$, and the deviations must be limited as much as possible. We block as many V_x as possible for candidates $x \notin \{\omega, y\}$ by choosing first the candidates which minimize $\min\{\min_{z \neq p}(|V_z| + \mathbb{1}_{\{p \triangleright z\}}); |V_x \cap N_{\omega>y}| + \mathbb{1}_{\{p \triangleright x\}}\} - |V_x \cap N_{y>\omega}|$. If at some point there are not enough available vetoes to block another set of voters V_x, then we assign the rest of available vetoes to candidate x which maximizes this quantity. ❑

The idea of the algorithm for the destructive variant is similar. The details are omitted due to space restrictions.

Proposition 4.4. *F_V-Destructive-Election-Enforcing is solvable in polynomial time when the social network is empty.*

Note that the algorithms of Propositions 4.3 and 4.4 can be simply adapted for taking into account restrictions in voter manipulations. For the poll-restricted variant, the principle of the algorithms remains the same. The only specific point is how a candidate y is made a (potential) winner. Instead of starting from the minimal vector of scores which satisfies the likelihood condition, we start from the truthful scores. If the imposed distance is d, then we can change a veto to one candidate from another at most $\lfloor d/2 \rfloor$ times. We switch vetoes against y to other candidates x with a smaller score with a priority to those for which V_x must be authorized to deviate.

Corollary 4.5. *$\mathcal{F}_V$-{Poll-Restr., Voter-Restr.}-{Constr., Destr.}-Election-Enforcing is in* P *when the network is empty.*

5 REAL POLL MANIPULATION: HEURISTICS

Most of our results are complexity results stating that, in the worst case, it may be hard for the polling agency to manipulate. However, it does not prevent manipulation to occur in practice. We thus examine some heuristics for the unrestricted problem of election enforcing and test them by running experiments. All our heuristics follow the same principle: we test all pairs of distinct candidates (ω, y) for announcing them the winner and another potential winner of Δ, respectively, following a given order. The order of test for pairs of candidates varies according to the variant of election enforcing and the voting rule, as described below (we omit further details due to space restrictions):

- Plurality / constructive: We refine a little the heuristic proposed by Wilczynski [20]. The order over pairs is such that $(\omega, y) \geq (\omega', y')$ where $\omega \neq p$ and $\omega' \neq p$ (target candidate p should not lose points) if $y = p$ and $y' \neq p$, or if $y = y' = p$ and $|N_{y>\omega}| \geq |N_{y>\omega'}|$, or if $y \neq p$ and $y' \neq p$ and $|N_{y>\omega}| \leq |N_{y'>\omega'}|$. In such a way, we favor configurations where p can get more points.
- Plurality / destructive: The order is such that $(\omega, y) \geq (\omega', y')$ if $|N_{y>\omega}| \geq |N_{y'>\omega'}|$ where $p \notin \{y, y'\}$ (target candidate p should not get more points). In such a way, we favor configurations where many voters will deviate to favor potential winner y.
- Veto / constructive: The order is such that $(\omega, y) \geq (\omega', y')$ if $|V_p \cap N_{y>\omega}| \geq |V_p \cap N_{y'>\omega'}|$ where $p \notin \{\omega, \omega'\}$ (deviations should not add more vetoes to target candidate p). In such a way, we favor configurations where more vetoes will be removed from p.
- Veto / destructive: For a given pair (ω, y), let x be the candidate which minimizes $\mathbf{s}^*(x) = |V_x \cap N_{\omega>y}|$. The order over pairs is such that $(\omega, y) \geq (\omega', y')$ if $p \notin \{x, x'\}$ and $\mathbf{s}^*(x) \leq \mathbf{s}^*(x')$, or if $\omega = p$ and $\omega' \neq p$, or if $\omega = \omega' = p$ and $|N_{y>\omega} \setminus V_p| \geq |N_{y'>\omega'} \setminus V_p|$. In such a way, we favor configurations where a candidate $x \neq p$ can lose many vetoes, or where many voters will deviate by vetoing p.

We test our heuristics by running 1,000 instances of the poll-confident iterative model with 50 agents and 5 candidates. The preference rankings of the agents are drawn from the impartial culture and the social network is supposed to be acyclic (in order to ensure convergence, to not limit too much manipulation and because our problems are hard for this class of graphs).

We compare the results of heuristics with the results given by the dynamics without manipulation from the polling agency and the results given by the exact algorithm where all possible manipulations of the polling agency that satisfy the likelihood condition are tested. We measure the frequency of election (for the constructive variant) of the target candidate as the iterative winner, or the frequency of non-election (for the destructive variant) of the target candidate as the iterative winner, according to the three different algorithms.

In order to create more challenge for the heuristics, the target candidates for the constructive variant are "bad" candidates: the Condorcet loser, i.e., the candidate which is beaten by all the other candidates in pairwise comparisons (we restrict in this case to a domain where such a candidate exists), or the Borda loser, i.e., the candidate with the lowest Borda score,[2] or the truthful loser, i.e., the candidate with the lowest (resp., highest) score under plurality (resp., veto). In the same vein, the target candidates for the

[2] For computing the Borda score of candidate x, we add $m - j$ points to x for each voter i if x is the j^{th} most preferred candidate of voter i.

destructive variant are "good" candidates: the Condorcet winner, i.e., the candidate which beats all the other candidates in pairwise comparisons (we restrict in this case to a domain where such a candidate exists), or the Borda winner, i.e., the candidate with the highest Borda score, or the truthful winner, i.e., the candidate with the highest (resp., lowest) score under plurality (resp., veto).

The results concerning both variants are presented in Figure 1.

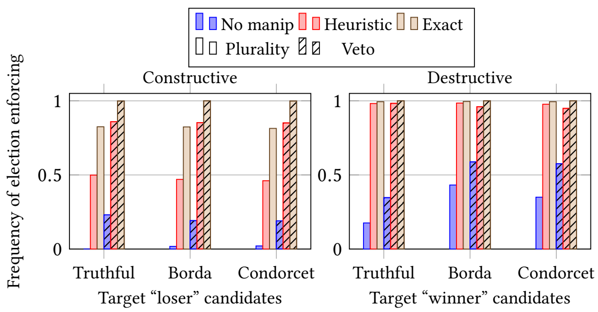

Figure 1: $\mathcal{F}$-{Constr., Destr.}-Election-Enforcing with no poll manipulation, heuristic or exact poll manipulation for different target candidates (truthful/Borda/Condorcet winner/loser) under $\mathcal{F} \in \{plurality, veto\}$ in an acyclic social network for $n = 50$ and $m = 5$.

It turns out that our heuristics for the destructive variant perform very well: the frequency of non-election of the target candidate p is very high and extremely close to the frequency with the exact algorithm. For the constructive variant, the frequency of election of p is very close under veto but the performance of our heuristic is a little bit lower under plurality. Nevertheless, it is always closer to the result of the exact algorithm than to the result where no poll manipulation occurs. This can be explained by the structure of the potential winners set under plurality: in our heuristic we only choose one potential winner to announce as a challenger of the announced winner whereas it could be cleverer to announce as potential winners an appropriate set of candidates.

It seems that even the results with the exact algorithm differ according to the variant of manipulation and the voting rule. In order to have a deeper understanding of this phenomenon, we run further experiments with the exact algorithm where the setting of simulations is the same as previously, except that we vary the number of agents from 10 to 50. The results are presented in Figure 2.

From the results given in Figure 2, two main conclusions can be drawn: (1) the polling agency can successfully manipulate more often for avoiding the election of a candidate than for making a candidate elected, i.e., the frequency of election enforcing is clearly higher for the destructive variant than for the constructive variant, and (2) the polling agency can successfully manipulate more often under veto than under plurality. The highest frequency of successful manipulation occurs for the destructive variant under veto, which seems natural regarding the nature of this voting rule under which a ballot means a disapproval for one candidate.

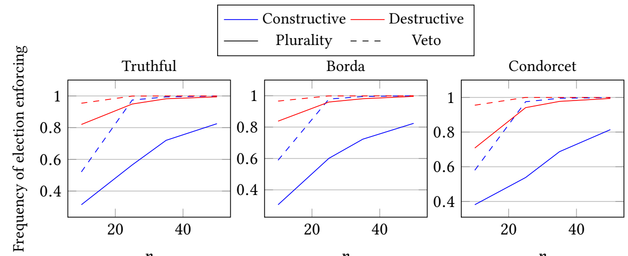

Figure 2: $\mathcal{F}$-{Constr, Destr.}-Election-Enforcing with exact poll manipulation for different target candidates (truthful/Borda/Condorcet winner/loser) under $\mathcal{F} \in \{plurality, veto\}$, in an acyclic social network for $m = 5$.

6 CONCLUSIONS

We have examined the manipulative power of a polling agency announcing preliminary results before an election. The polling agency may manipulate with two different goals in mind: making a given candidate elected (constructive variant) or avoiding the election of a given candidate (destructive variant). However, the polling agency is not totally free regarding how it can manipulate: the announced scores should not be too far from reality to be trusted by voters. Moreover, voters may have a local information by their relatives in a social network, limiting the manipulative power of the polling agency. Our results are summarized in the table below.

Manip.	Variant	Plurality		Veto	
		Acyclic network	Empty network	Acyclic network	Empty network
Constr.	Unrestr.	NP-h ([20]) FPT w.r.t. m (Prop. 3.3)	?	NP-h (Th. 4.1) FPT w.r.t. m (Prop. 4.2)	P (Prop. 4.3)
	Poll-restr.	NP-h (Th. 3.2) FPT w.r.t. m (Prop. 3.3)	NP-h (Th. 3.4)	NP-h (Th. 4.1) FPT w.r.t. m (Prop. 4.2)	P (Cor. 4.5)
	Voter-restr.	NP-h (Th. 3.2) FPT w.r.t. m (Prop. 3.3)	?	NP-h (Th. 4.1) FPT w.r.t. m (Prop. 4.2)	P (Cor. 4.5)
Destr.	Unrestr.	NP-h (Th. 3.1) FPT w.r.t. m (Prop. 3.3)	P (Prop. 3.5)	NP-h (Th. 4.1) FPT w.r.t. m (Prop. 4.2)	P (Prop. 4.4)
	Poll-restr.	NP-h (Th. 3.2) FPT w.r.t. m (Prop. 3.3)	P (Prop. 3.5)	NP-h (Th. 4.1) FPT w.r.t. m (Prop. 4.2)	P (Cor. 4.5)
	Voter-restr.	NP-h (Th. 3.2) FPT w.r.t. m (Prop. 3.3)	?	NP-h (Th. 4.1) FPT w.r.t. m (Prop. 4.2)	P (Cor. 4.5)

When the voters have no local information through the social network, manipulating is easier for the polling agency, especially under veto. Although the manipulative power of the polling agency is mainly computationally limited in theory, we designed efficient heuristics. They perform better for the destructive variant under veto. More generally, it seems that the two variants of manipulation and the two voting rules we consider are not symmetric: the polling agency is more successful in the destructive than in the constructive case, and manipulation is more successful under veto than under plurality. This work can be extended in several directions. Considering more complex voting rules which require the submission of a ranking in ballots could be a challenging perspective. Investigating preference restrictions such as single-peaked preferences could also make sense, as well as supposing that the polling agency only gets partial information about the preferences of the voters.

ACKNOWLEDGMENTS

This work was supported in part by DFG grants BA 6270/1-1 and BR 2312/12-1.

REFERENCES

[1] John J. Bartholdi III, Craig A. Tovey, and Michael A. Trick. 1989. The Computational Difficulty of Manipulating an Election. *Social Choice and Welfare* 6, 3 (1989), 227–241.

[2] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia (Eds.). 2016. *Handbook of Computational Social Choice.* Cambridge University Press.

[3] Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. 2016. Strategic Voting with Incomplete Information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16).* AAAI Press / IJCAI, New York, NY, USA, 236–242.

[4] Piotr Faliszewski, Rica Gonen, Martin Koutecký, and Nimrod Talmon. 2018. Opinion Diffusion and Campaigning on Society Graphs. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18).* IJCAI, Stockholm, Sweden, 219–225.

[5] Jörg Flum and Martin Grohe. 2006. *Parameterized Complexity Theory.* Springer-Verlag.

[6] Teofilo F. Gonzalez. 1985. Clustering to Minimize the Maximum Intercluster Distance. *Theoretical Computer Science* 38 (1985), 293–306.

[7] Laurent Gourvès, Julien Lesca, and Anaëlle Wilczynski. 2016. Strategic Voting in a Social Context: Considerate Equilibria. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI-16).* IOS Press, The Hague, The Netherlands, 1423–1431.

[8] Kathrin Konczak and Jérôme Lang. 2005. Voting Procedures with Incomplete Preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Worshop on Advances in Preference Handling.* Edinburgh, Scotland, 124–129.

[9] Omer Lev and Jeffrey S Rosenschein. 2016. Convergence of Iterative Scoring Rules. *Journal of Artificial Intelligence Research* 57 (2016), 573–591.

[10] Reshef Meir. 2017. Iterative Voting. In *Trends in Computational Social Choice*, Ulle Endriss (Ed.). AI Access, Chapter 4, 69–86.

[11] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. 2014. A Local-Dominance Theory of Voting Equilibria. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC-14).* ACM, Palo Alto, California, USA, 313–330.

[12] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. 2010. Convergence to Equilibria in Plurality Voting. In *Proceedings of the 24th Conference on Artificial Intelligence (AAAI-10).* AAAI Press, Atlanta, Georgia, USA, 823–828.

[13] Svetlana Obraztsova and Edith Elkind. 2012. Optimal Manipulation of Voting Rules. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-12).* IFAAMAS, Valencia, Spain, 619–626.

[14] Svetlana Obraztsova, Omer Lev, Maria Polukarov, Zinovi Rabinovich, and Jeffrey S. Rosenschein. 2016. Nonmyopic Voting Dynamics: An Optimistic Approach. In *Proceedings of the 10th Multidisciplinary Workshop on Advances in Preference Handling (MPREF-16).* New York, NY, USA.

[15] Christos Papadimitriou. 1995. *Computational Complexity* (second ed.). Addison-Wesley.

[16] Annemieke Reijngoud and Ulle Endriss. 2012. Voter Response to Iterated Poll Information. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-12).* IFAAMAS, Valencia, Spain, 635–644.

[17] Reyhaneh Reyhani and Mark C. Wilson. 2012. Best Reply Dynamics for Scoring Rules. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI-12).* IOS Press, Montpellier, France, 672–677.

[18] Sigal Sina, Noam Hazon, Avinatan Hassidim, and Sarit Kraus. 2015. Adapting the Social Network to Affect Elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15).* IFAAMAS, Istanbul, Turkey, 705–713.

[19] Alan Tsang and Kate Larson. 2016. The Echo Chamber: Strategic Voting and Homophily in Social Networks. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-16).* IFAAMAS, Singapore, 368–375.

[20] Anaëlle Wilczynski. 2019. Poll-Confident Voters in Iterative Voting. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI-19).* AAAI Press, Honolulu, Hawaii, USA, 2205–2212.

[21] Bryan Wilder and Yevgeniy Vorobeychik. 2018. Controlling Elections through Social Influence. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-18).* IFAAMAS, Stockholm, Sweden, 265–273.

[22] Bryan Wilder and Yevgeniy Vorobeychik. 2019. Defending Elections against Malicious Spread of Misinformation. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI-19).* AAAI Press, Honolulu, Hawaii, USA, 2213–2220.