

Do We Care About Poll Manipulation in Political Elections?

Vincent Mousseau^a, Henri Surugue^{a,*} and Anaëlle Wilczynski^a

^aMICS, CentraleSupélec, Université Paris-Saclay

Abstract. We consider the problem of poll manipulation in political elections. In the context of strategic voting, we are interested in whether a polling institute can manipulate the information it communicates to voters in order to influence the outcome of the election. We start with a version of the problem where the polling institute is allowed to send any score to voters. Then, for realistic reasons, we investigate a restricted version in which the polling institute cannot announce scores which are too far from the truthful ones. While we show that both decision problems are computationally hard, we go beyond this worst-case complexity analysis by using probabilistic tools to address the possibility of successful and efficient manipulation in practice, w.r.t. several natural preference distributions.

1 Introduction

Strategic voting [20] is a major issue in political elections. Ideally, one would like to avoid such a strategic behavior. However, by the Gibbard-Satterthwaite theorem [14, 26], no reasonable voting rule is immune to voter manipulation. Since we cannot escape strategic voting, one way to tackle manipulation is to study it via a game-theoretical analysis. This approach has been followed by several works in computational social choice [5]. Notably, in *iterative voting* [19], the idea is to analyze the convergence and the quality of a sequential process where voters are allowed to make successive strategic deviations. In the classical iterative voting framework [21], complete knowledge is assumed, in the sense that voters are aware of all others' current ballot, sometimes even of their full preferences. However, this assumption is highly unrealistic and does not capture real scenarios with large electorates, such as political elections.

The question of the information available to the voters is key and has a strong impact on the manipulability of voting processes [10, 25]. To deal with partial information in voting, one can naturally follow a Bayesian approach by considering a probability distribution over a set of possible preference orders for other voters [23, 17]. Alternatively, a set of possible preference profiles can be derived from partial votes [7, 8] or from a given maximum distance to the voters' actual preferences [1]. Another possibility is to assume local information for the voters, which is captured by a social network [16]. Finally, an aggregated global information coming from opinion polls can be communicated to voters [3, 10, 25, 29].

Following this latter line of research, in this article, inspired by political elections, we assume that voters receive only a global information about the voting intentions within the population, which is communicated through *opinion polls*. Voters trust the information

communicated in the polls and compute their best response ballot on the basis of this information. This confidence in the polls grants an important power to the polling institute which disseminates it, raising the natural question of *poll manipulation*. Indeed, a polling institute might have its own interests in the election and try to orient votes toward them. This problem is close to the question of election control [12], where an external agent aims to alter the outcome of the election, but here no structural change is made on the election.

In the line of seminal works analyzing the complexity of voter manipulation [2], one can analyze the complexity of the poll manipulation problem. However, computational intractability may not constitute a relevant barrier to manipulation, as it relies on worst-case analysis [11]. Therefore, to complement complexity results, an average-case study using a probabilistic approach is relevant, as it has been widely investigated for voter manipulation (see, e.g., [13, 18, 24, 31]). In particular, the asymptotic study is meaningful since political elections are characterized by a large number of voters. Considering election control problems, as far as we know, this approach has been surprisingly neglected. A notable exception is a recent work by Xia [30] which investigates the likelihood of manipulability for several coalition influence problems, including control by adding or deleting votes. Up to our best knowledge, no such study has been conducted so far for the poll manipulation problem.

In this article, we study the constructive poll manipulation problem where the polling institute wishes to favor a specific candidate by broadcasting manipulated candidates' scores. This problem has been introduced by Wilczynski [29] and further extended by Baumeister et al. [3], who also consider the destructive variant where the polling institute aims to prevent the election of a given candidate. While both works consider a framework where voters are embedded in a social network and analyze the complexity of the problem with respect to the structure of the graph, we consider a simpler model with no social network, which clarifies the role of the opinion polls. In particular, we analyze the following two versions of the problem. In the unrestricted problem, the polling institute is free to send any score information. The restricted problem considers a more realistic context where only score information that would be close enough to truthful scores are allowed. The idea for this second problem is for the polling institute to lie in a reasonable manner, by submitting realistic scores, not too far from a ground truth that may correspond to the results of a past election, or another poll. Such restrictions help to gain the trust and confidence from the voters. We prove that both versions of the problem are computationally hard, answering an open question from Baumeister et al. [3], but also analyze the probability of existence of a successful and efficiently computable poll manip-

* Corresponding Author. Email: henri.surugue@centralesupelec.fr.

ulation. For this latter purpose, we introduce a natural condition on statistical cultures which is satisfied by most natural preference distributions [27]. In fact, we exhibit a simple heuristic and prove its success for the unrestricted problem, which means that, without restriction, the polling institute can almost always efficiently influence large elections. For the restricted manipulation problem, we prove that if the allowed distance is negligible with respect to the number of voters, then no manipulation is possible. However, when this distance becomes significant, e.g., when it is a fixed proportion of the number of voters, easy manipulation is almost always successful in large elections. Finally, we show that most results still hold when assuming a more general strategic behavior for voters [29].

2 The Model

We first present the poll manipulation problem in the context of strategic voting.

2.1 A voting system

For any positive integer k , let $[k]$ denote the set $\{1, \dots, k\}$. Let N be a set of voters where $N = [n]$, and M be a set of candidates where $M = [m]$. Since the goal is to study strategic voting in large political elections, we naturally assume that $n > m > 2$ (by the Gibbard-Satterthwaite theorem, voting rules are susceptible to manipulation only when there are more than two candidates). Each voter $i \in N$ has preferences over candidates represented by a linear order \succ_i over candidates. Let top_{\succ_i} and $worst_{\succ_i}$ denote the most preferred and worst candidates, respectively, of voter $i \in N$, i.e., $top_{\succ_i} \succeq_i x$ and $x \succeq_i worst_{\succ_i}$ for every candidate $x \in M$. The rank of a given candidate x in a preference order \succ_i is denoted by $r_{\succ_i}(x)$, i.e., $r_{\succ_i}(x) := |\{y \in M : y \succeq_i x\}|$. We use the Kendall tau distance to evaluate the similarity between two preference orders, by counting the number of pairwise comparisons on which the two orders disagree, i.e., $dist_{KT}(\succ_i, \succ_j) = |\{(x, y) \in M^2 : x \succ_i y \text{ and } y \succ_j x\}|$. The set of linear orders \succ_i for all voters $i \in N$ is called a preference profile and is denoted by \succ . Let us denote by $N^{x \succ y}$ the set of voters who prefer x over y , i.e., $N^{x \succ y} := \{i \in N : x \succ_i y\}$.

The winner of an election is determined by the Plurality voting rule where ties are broken lexicographically. Let $b_i \in M$ denote the ballot of voter i and $b \in M^n$ denote the ballot profile. The ballot profile b from which b_i is excluded is denoted by b_{-i} . The winner under Plurality of the ballot profile b is $\mathcal{W}_P(b) \in \arg \max_{x \in M} s_x(b)$, where $s_x(b) := |\{i \in N : b_i = x\}|$ and a lexicographic tie-breaking, denoted by \triangleright , is used if necessary. By abuse of notation, we sometimes directly write $\mathcal{W}_P(s)$ to refer to the winner of a score vector s . Let b^T denote the truthful ballot profile, i.e., $b_i^T \succeq_i x$ for every candidate $x \in M$ and voter $i \in N$, and s^T denote the candidates' scores in b^T .

An election is given by the tuple $(N, M, \succ, \triangleright)$.

2.2 A strategic voting framework

In this model, we consider an iterative voting process where voters are strategic with respect to the information they get, which only consists in the score broadcast by the polling institute. Like some previous works in iterative voting [3, 29], we make the assumption that voters trust the announced result by the polling institute but, in contrast, we assume that voters have no other local information on which to rely. Let us describe more in details the voters' strategic behavior.

Initially, all voters vote sincerely since they have no information yet, therefore the initial ballot profile b^0 is exactly the truthful ballot profile b^T . Then, the polling institute sends the results of this

initial election by announcing a score vector s of size m such that $\sum_{j \in M} s_j = n$, where s_j is supposed to stand for the score of candidate $j \in M$ in the initial election. Let s^{-i} denote the score vector s where the truthful ballot of voter i has been removed, i.e., $s_j^{-i} = s_j$, for every $j \in M \setminus \{b_i^T\}$, and $s_{b_i^T}^{-i} = s_{b_i^T} - 1$. After the score is announced, every voter considers possible moves from her initial truthful ballot with respect to that information. Possible deviations are captured by the notion of potential winners: A candidate y is a *potential winner* for voter i w.r.t. announced score vector s if i believes that voting for y will make candidate y the new winner, i.e., $s_{\mathcal{W}_P(s^{-i})}^{-i} - s_y^{-i} + \mathbb{1}_{\mathcal{W}_P(s^{-i}) \triangleright y} \leq 1$. Let PW_i^s denote the set of potential winners for voter i w.r.t. announced scores s .

We say that a candidate y is a potential winner if there exists a voter i such that $y \in PW_i^s$. By definition, the announced winner $\mathcal{W}_P(s)$ is also a potential winner. A voter is said to be *pivotal* if her considered set of potential winners contains more candidates than the winner exclusively. We consider the following best response for each voter i w.r.t. announced score s : i deviates from her current ballot b_i^T to another ballot supporting candidate $y \in PW_i^s \setminus \{\mathcal{W}_P(s)\}$ if y is her most preferred candidate within PW_i^s . Each voter can deviate at most once since she only gets the information about the scores s provided by the polling institute, and cannot see the deviations from other voters (thus the order of voters' deviations does not matter, they could even be simultaneous). Hence, the deviation process ends after at most n steps and converges to a final ballot profile denoted by b^s .

2.3 A decision problem of manipulation

In this model, a polling institute sends out a score and then each voter votes strategically w.r.t. that information, and finally the winner of the election is computed. We want to describe the behavior of the polling institute who may have its own interest in the election. Let x^* be the target candidate of the polling institute, i.e., it wants x^* to be elected. Let I be the space of all possible scores that the polling institute can announce, i.e., $I := \{s \in \mathbb{N}^m \mid \sum_{j=1}^m s_j = n\}$. We consider the following poll manipulation problem:

UNRESTRICTED MANIPULATION PROBLEM

Instance:	Election $(N, M, \succ, \triangleright)$, target candidate $x^* \in M$
Question:	Does there exist a score $s \in I$ to announce such that $\mathcal{W}_P(b^s) = x^*$?

However, the fact that the polling institute is allowed to send any score is not very realistic. We use a restricted version of the decision problem where the distance between the truthful poll and the one sent by the polling institute is bounded. We use the number of vote changes to evaluate the distance between possible scores, i.e., $d(s, s') = \frac{1}{2} \sum_{j \in M} |s_j - s'_j|$, for every scores $s, s' \in I$. Note that this distance is equivalent to the restriction of the ℓ_1 distance on I divided by 2 and sometimes called the "earth mover distance" in the literature [20]. We let $I_k := \{s \in \mathbb{N}^m \mid d(s, s^T) \leq k \text{ and } \sum_{j \in M} s_j = n\}$ be the restricted space of action of the polling institute. We then analyze the following poll manipulation problem:

RESTRICTED MANIPULATION PROBLEM

Instance:	Election $(N, M, \succ, \triangleright)$, target candidate $x^* \in M$, integer k
Question:	Does there exist a score $s \in I_k$ to announce such that $\mathcal{W}_P(b^s) = x^*$?

A poll manipulation is illustrated in the next example.

Example 1. Let us consider an election $(N, M, \succ, \triangleright)$ where $N = \{1, \dots, 8\}$, $M = \{a, b, c, d, e\}$, the tie-breaking \triangleright follows the alphabetical order and the preferences \succ are as follows:

$$\begin{array}{ll} 1: & a \succ c \succ d \succ e \succ b \\ 2: & e \succ b \succ d \succ c \succ a \\ 3, 4: & a \succ d \succ b \succ c \succ e \\ 5, 6: & b \succ d \succ e \succ c \succ a \\ 7: & c \succ d \succ a \succ b \succ e \\ 8: & d \succ e \succ a \succ b \succ c \end{array}$$

The initial truthful scores are given by $s^0 = (3, 2, 1, 1, 1)$ (candidates are indexed w.r.t. their alphabetical position). Suppose that the polling institute communicates the following score vector $s^M = (0, 2, 2, 3, 1)$, at distance 3 to the truthful one. The set of potential winners w.r.t. s^M is equal to $PW_i^{s^M} = \{b, c, d\}$ for every voter $i \in \{1, 2, 3, 4, 8\}$, while $PW_5^{s^M} = PW_6^{s^M} = \{c, d\}$ and $PW_7^{s^M} = \{b, d\}$. Voters 3, 4, 5, 6, 7, and 8 do not have an incentive to deviate since the announced winner d is their most preferred candidate among the potential winners. However, voters 1 and 2 have an incentive to deviate to a ballot supporting c and b , respectively. After their deviations, we reach the final scores $s^2 = (2, 3, 2, 1, 0)$ where b is the winner. Hence, the polling institute can enforce the election of b , whereas a would remain the winner without poll manipulation.

3 Preliminaries on Voting Cultures

In this section, we present voting cultures that are commonly used in the literature [27] to represent the distributions of preferences in elections. We will show that most of them satisfy a general condition on cultures, which is very useful for the purpose of our paper. Let Π^m be the set of all preference orders for m candidates, and $\succ_i \in \Pi^m$ be voter i 's preference order for an arbitrary $i \in N$. We denote as $C(n, \Pi_{sub}^m)$ the probability distribution of drawing n preference orders from $\Pi_{sub}^m \subseteq \Pi^m$ to constitute our preference profile. Such a probability distribution $C(n, \Pi_{sub}^m)$ is called a *culture* and simply denoted by C when the context is clear, and its associated probability is denoted by \mathbb{P}_C . We use $\mathbb{P}_C(a \succ_i b)$ when it is clear from the context instead of $\mathbb{P}_C(\succ_i | a \succ_i b)$. In the following of the paper, we consider independent and identical drawings of voters' preferences such that we can either look at the distribution $C(n, \Pi_{sub}^m)$ as a whole object or n drawings of preferences \succ_i . For technical reasons, we suppose that there are more than two different candidates which are ranked first by a preference order with a positive probability to be drawn under the considered culture. Note that this assumption is also natural since we focus on strategic voting and manipulation only occurs with at least three candidates [14, 26].

Definition 1 (Impartial culture). *The impartial culture, called IC, draws every preference order \succ_i from Π^m with uniform probability.*

One can also define variants of impartial cultures which are uniform but only on a given subset of Π^m , e.g., on single-peaked orders. A preference profile \succ is said to be *single-peaked* [4] if there exists an axis $>$ on M such that, for every voter $i \in N$, and each triple of candidates $x > y > z$, we have $y \succ_i x$ or $y \succ_i z$. Let $\Pi_{>}^m$ be the set of single-peaked preference orders w.r.t. a given axis $>$ on M .

Definition 2 (Single-peaked culture). *For a given axis $>$ over M , a culture $C(n, \Pi_{sub}^m)$ is said to be single-peaked if $C(n, \Pi_{sub}^m) = C(n, \Pi_{>}^m)$.*

We might sample single-peaked preference orders by drawing them uniformly on the restricted space of single-peaked preference orders $\Pi_{>}^m$. The associated culture then refers to Walsh's model [28]. Another way to impartially sample single-peaked preference orders is to use Conitzer's model [6] which draws preference orders in $\Pi_{>}^m$ so that each candidate has the same chance to be ranked first.

Definition 3 (Mallows culture). *For given $\sigma \in \Pi^m$ and $\phi \in [0, 1]$, the Mallows culture, called $\mathcal{M}^{\phi, \sigma}$, draws every preference order with a probability related to its distance to the reference ranking σ , more precisely, $\mathbb{P}_{\mathcal{M}^{\phi, \sigma}}(\succ_i) = \frac{1}{Z} \phi^{dist_{KT}(\succ_i, \sigma)}$ where $Z = \sum_{\succ_i \in \Pi^m} \phi^{dist_{KT}(\succ_i, \sigma)}$.*

Note that culture $\mathcal{M}^{1, \sigma}$ corresponds to the impartial culture.

We introduce below a simple property on cultures which will be key in the poll manipulation analysis.

Definition 4 (Balanced culture). *A distribution $C(n, \Pi^m)$ is said to be balanced for a given candidate $c \in M$ if there exists another sufficiently worst candidate $\ell \in M \setminus \{c\}$, in the sense that $\mathbb{P}_C(c \succ_i \ell) \geq \frac{1}{2}$. The set of such candidates ℓ for x is denoted by $B_C(x)$. In general, a distribution C is said to be balanced if it is balanced for every candidate $c \in M$.*

It turns out that all cultures that we consider are balanced.

Proposition 1. *The impartial culture is balanced.*

Proof. Let $i \in N$ be a voter. The impartial culture is balanced for every candidate because for any pair of candidates x and y , we have $\mathbb{P}_{IC}(x \succ_i y) = \mathbb{P}_{IC}(y \succ_i x) = \frac{1}{2}$ since each preference order in Π^m has the same probability to be drawn. \square

For a given axis $>$ on M , let $e_1^>$ and $e_2^>$ denote the two extreme candidates of $>$.

Proposition 2. *If $x \in M \setminus \{e_1^>, e_2^>\}$, then every single-peaked culture $C(n, \Pi_{>}^m)$ is balanced for x . If $x \in \{e_1^>, e_2^>\}$, then every single-peaked culture $C(n, \Pi_{>}^m)$ which also satisfies $\mathbb{P}_C(\{\succ_i | worst_{\succ_i} = x\}) \leq \frac{1}{2}$, is balanced for x .*

Sketch of proof. We mostly use the following fact: $\mathbb{P}_C(\{\succ_i | worst_{\succ_i} = e_1^>\} \cup \{\succ_i | worst_{\succ_i} = e_2^>\}) = 1$. \square

In particular, the previous proposition shows that both Walsh's [28] and Conitzer's [6] cultures are balanced.

Proposition 3. *Any Mallows culture $\mathcal{M}^{\phi, \sigma}$ is balanced for every candidate $x \in M \setminus \{worst_{\sigma}\}$.*

Sketch of proof. Consider any candidate $x \in M \setminus \{worst_{\sigma}\}$ and the candidate $\ell := worst_{\sigma}$. Let $\Pi_{y \succ z}^m$ denote the set of all preference orders where y is ranked before z , i.e., $\Pi_{y \succ z}^m := \{\succ_i \in \Pi^m : y \succ_i z\}$. Consider the bijection $\tau : \Pi_{\ell \succ x}^m \rightarrow \Pi_{x \succ \ell}^m$, where for every $\succ_i \in \Pi_{\ell \succ x}^m$, we construct the preference order $\tau(\succ_i) \in \Pi_{x \succ \ell}^m$ which is the same as \succ_i except that the positions of x and ℓ are swapped. One can show that $dist_{KT}(\sigma, \succ_i) \geq dist_{KT}(\sigma, \tau(\succ_i))$, for every $\succ_i \in \Pi_{\ell \succ x}^m$, by analyzing the differences between \succ_i and $\tau(\succ_i)$ in terms of agreement on pairwise comparisons with σ . By definition of the Mallows culture $\mathcal{M}^{\phi, \sigma}$, we thus have $\mathbb{P}_{\mathcal{M}^{\phi, \sigma}}(\tau(\succ_i)) \geq \mathbb{P}_{\mathcal{M}^{\phi, \sigma}}(\succ_i)$. Hence, we conclude that $\mathbb{P}_{\mathcal{M}^{\phi, \sigma}}(\{\succ'_i : x \succ'_i \ell\}) \geq \mathbb{P}_{\mathcal{M}^{\phi, \sigma}}(\{\succ'_i : \ell \succ'_i x\})$, and thus $\mathbb{P}_{\mathcal{M}^{\phi, \sigma}}(\{\succ'_i : x \succ'_i \ell\}) \geq \frac{1}{2}$, implying that $\mathcal{M}^{\phi, \sigma}$ is balanced for x . \square

4 The Unrestricted Poll Manipulation Problem

This section is devoted to the study of the unrestricted manipulation problem where the polling institute can send any score in I . We first give some results on the computational complexity of the problem then we continue our work with a probabilistic approach of the problem to capture what can happen in practice.

We first prove that, even in the unrestricted case, the poll manipulation problem is NP-complete. Our result answers an open question from Baumeister et al. [3].

Theorem 4. *The unrestricted manipulation problem is NP-complete.*

Sketch of proof. Membership to NP is straightforward: given communicated scores, we can efficiently derive the possible unique deviation of each voter and compute the winner in the deviating profile.

For hardness, we perform a reduction from a variant of EXACT COVER BY 3-SETS (X3C) known to be NP-complete [15]: Given a set $X = \{x_1, x_2, \dots, x_{3q}\}$ and a set $S = \{S_1, S_2, \dots, S_{3q}\}$ of 3-element subsets of X , where each element x_i occurs in exactly three subsets of S , we ask whether there exists an exact cover, i.e., a subset $S' \subseteq S$ such that every element of X occurs in exactly one member of S' . From an instance (X, S) of X3C, we construct an instance of our unrestricted manipulation problem as follows.

For each element x_i , for $i \in [3q]$, we create a candidate y_i , and for each subset S_j where $j \in [3q]$, we create a candidate c_j . We add three candidates w, z , and t where t is our target candidate.

There are $12q + 7$ voters: for each element x_i , for $i \in [3q]$, we create one voter Y_i , for each subset S_j , for $j \in [3q]$, we create three voters C_j^r where $r \in [3]$, and we finally add two voters T^ℓ , two voters Z^ℓ , two voters W^ℓ for $\ell \in [2]$, and one voter D .

Their preferences are defined below, where $y(s_j^r)$ denotes the candidate y_i associated with the r^{th} element of subset S_j , and when a subset of candidates is mentioned, the candidates are ranked according to the increasing order of their indices.

$$\begin{array}{ll} Y_i: & w \succ y_i \succ z \succ \{y_{i'}\}_{i' \neq i} \succ \{c_j\}_j \succ t & \text{for } i \in [3q] \\ C_j^r: & y(s_j^r) \succ c_j \succ z \succ w \succ \{y_{i'}\}_{i' \neq i} \succ \{c_j\}_j \succ t & \text{for } j \in [3q], r \in [3] \\ T^\ell: & t \succ z \succ w \succ \{y_i\}_i \succ \{c_j\}_j & \text{for } \ell \in [2] \\ Z^\ell: & z \succ w \succ \{y_i\}_i \succ \{c_j\}_j \succ t & \text{for } \ell \in [2] \\ W^\ell: & w \succ z \succ \{y_i\}_i \succ \{c_j\}_j \succ t & \text{for } \ell \in [2] \\ D: & w \succ t \succ z \succ \{y_i\}_i \succ \{c_j\}_j & \end{array}$$

Finally, the tie-breaking rule is as follows: $w \triangleright t \triangleright z \triangleright y_1 \triangleright \dots \triangleright y_{3q} \triangleright c_1 \triangleright \dots \triangleright c_{3q}$.

The winner of the election with the truthful ballot profile is candidate w . The details of the scores for this truthful ballot profile are given in the second column of Table 1.

Table 1. Candidates' scores in the complexity proof of Theorem 4

candidate	initial score	announced score	score after manipulation
y_i ($i \in [3q]$)	3	3	3
c_j ($j \in [3q]$)	0	3 if $S_j \in S'$ 0 otherwise	3 if $S_j \in S'$ 0 otherwise
w	$3q + 3$	2	2
t	2	2	3
z	2	3	2
winner	w	z	t

We claim that there exists an exact cover in (X, S) iff we can force the election of candidate t in the constructed instance.

Suppose first that there exists a subset $S' \subseteq S$ such that every element of X occurs in exactly one element of S' . Let us consider manipulated communicated scores which differ from the sincere ones by taking $3q + 1$ votes initially given to w to give one additional vote to z and three votes to c_j for each $S_j \in S'$. These scores are summarized in the third column of Table 1. One can prove that these communicated scores trigger deviations which lead to the final scores presented in the fourth column of Table 1, where t is the winner.

Suppose now that there exist communicated scores such that the target candidate t becomes the winner after deviations from the voters. One can show that the only possibility for communicated scores

to lead to the victory of t is to announce candidate z the winner and, as potential winners, the target candidate t and exactly q candidates c_j which correspond to subsets S_j forming an exact cover of X . \square

Note that even though we have proved that the problem is NP-complete, we know from Baumeister et al. [3] that it is FPT when parameterized by the number of candidates m . Another way to go beyond the NP-hardness result, which focuses on worst-case complexity, is to analyze the actual possibility of poll manipulation using a probabilistic approach which works even when m is large. We will see that the poll manipulation problem is often easy to tackle in a probabilistic point of view, following natural statistical cultures as defined in Section 3. We will start by considering a balanced culture.

For a given target candidate x^* the polling institute wants to elect, we say its poll manipulation is successful if after all strategic moves from voters, the desired candidate x^* is elected. Let us denote by S the associated event of success, which corresponds to the yes-instances of the unrestricted poll manipulation problem.

Let $2PW\text{-}H(x^*, \ell)$ be the heuristic which announces a score with exactly two potential winners x^* and ℓ , with x^* the target candidate and ℓ the announced winner. For realistic conditions, one point is given to candidates with a positive score in the truthful ballot profile. Assuming $n > m + 5$ is sufficient to guarantee the possibility of making any pair of candidates the only potential winners (this hypothesis is rather weak since we focus on large elections in terms of voters). It then suffices to check whether the associated communicated polling score leads to the victory of x^* . This heuristic can be called by a global heuristic, which tests it with different candidates ℓ . Our heuristics are computable in polynomial time and are inspired from the heuristics of Wilczynski [29] and Baumeister et al. [3], where the idea is to find a candidate ℓ , which is a threatening winner, i.e., enough voters prefer x^* to ℓ , while x^* is the only credible alternative to ℓ , in order to incentivize voters to deviate to x^* .

Let $S_{2PW\text{-}H(x^*, \ell)}$ denote the event of success for heuristic $2PW\text{-}H(x^*, \ell)$. Let $X_{-\ell}$ be the random variable which counts the number of voters who prefer x^* over ℓ , i.e., $X_{-\ell} = |N^{x^* \succ \ell}|$.

Similarly, let $Y_{\ell, j}$ be the random variable which counts the number of voters who prefer ℓ over x^* while their most preferred candidate is j , i.e., $Y_{\ell, j} = |\{i \in N^{\ell \succ x^*} : \text{top}_{\succ_i} = j\}|$. If our heuristic $2PW\text{-}H(x^*, \ell)$ indeed succeeds to announce exactly two potential winners x^* and ℓ with ℓ as a winner, then only voters who prefer x^* over ℓ and currently vote for another candidate, will deviate and they will do so in favor of x^* . Note that voters already having x^* as their top choice would keep this vote because there is no other potential winner. Therefore, in total, after deviations, x^* obtains a number of votes which is equal to the numbers of voters who prefer x^* over ℓ . It follows that x^* would win only if the number of voters preferring x^* over ℓ is greater than the number of voters who keep their vote for another candidate, implying that for a given culture $C(n, \Pi^m)$, $\mathbb{P}_C(S_{2PW\text{-}H(x^*, \ell)}) = \mathbb{P}_C(\forall j \in M \setminus \{x^*\}, Y_{\ell, j} \leq X_{-\ell})$.

Our first theorem provides a high lower bound on the probability of success of the poll manipulation heuristic.

Theorem 5. *For a balanced culture $C(n, \Pi^m)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of the sub-heuristic $2PW\text{-}H(x^*, \ell)$, is as follows: $\mathbb{P}_C(S_{2PW\text{-}H(x^*, \ell)}) \geq 1 - 2(m - 2)(e^{-2n(p_{x^*, \ell} - q_{x^*, \ell, j^*})^2})$ where:*

- $p_{x^*, \ell} := \mathbb{P}_C(x^* \succ_i \ell)$,
- $r_{x^*, \ell, j} := \mathbb{P}_C(\{\ell \succ_i x^*\} \cap \{j = \text{top}_{\succ_i}\})$, for $j \neq \ell$,
- $q_{x^*, \ell, j} := \frac{p_{x^*, \ell} + r_{x^*, \ell, j}}{2}$,
- $j^* := \arg \min_{j \in M \setminus \{x^*, \ell\}} \mathbb{P}_C(Y_{\ell, j} \leq X_{-\ell})$.

In particular, the probability of success of the global heuristic satisfies the same lower bound.

Sketch of proof. For our target candidate x^* and a balanced culture $C(n, \Pi^m)$, let us consider a candidate $\ell \in B_C(x^*)$. Since $\ell \in B_C(x^*)$, we always have $Y_{\ell, \ell} \leq X_{-\ell}$ and thus $\mathbb{P}_C(S_{2PW-H}(x^*, \ell)) = \mathbb{P}_C(\forall j \in M \setminus \{x^*, \ell\}, Y_{\ell, j} \leq X_{-\ell})$. We will show a lower bound to this latter quantity. Using Bonferroni's inequality, we have: $\mathbb{P}_C(\forall j \in M \setminus \{x^*, \ell\}, Y_{\ell, j} \leq X_{-\ell}) \geq \sum_{j \in M \setminus \{x^*, \ell\}} \mathbb{P}_C(Y_{\ell, j} \leq X_{-\ell}) - (m-3) \geq (m-2) \cdot \min_{j \in M \setminus \{x^*, \ell\}} \mathbb{P}_C(Y_{\ell, j} \leq X_{-\ell}) - (m-3)$. By considering $j^* := \arg \min_{j \in M \setminus \{x^*, \ell\}} \mathbb{P}_C(Y_{\ell, j} \leq X_{-\ell})$, we then have that $\mathbb{P}_C(\forall j \in M \setminus \{x^*, \ell\}, Y_{\ell, j} \leq X_{-\ell}) \geq (m-2) \cdot (\mathbb{P}_C(Y_{\ell, j^*} \leq X_{-\ell}) - 1) + 1$. Let us now treat the term $\mathbb{P}_C(Y_{\ell, j^*} \leq X_{-\ell})$. We remark that $X_{-\ell}$ follows a binomial distribution of parameters n and $p_{x^*, \ell}$, and Y_{ℓ, j^*} follows a binomial distribution of parameters n and r_{x^*, ℓ, j^*} , where $p_{x^*, \ell}$ and $r_{x^*, \ell, j}$ are defined as $p_{x^*, \ell} = \mathbb{P}_C(x^* \succ_i \ell)$ and $r_{x^*, \ell, j} = \mathbb{P}_C(\{\ell \succ_i x^*\} \cap \{j = \text{top}_{\succ_i}\})$, for every $j \neq \ell$. We introduce $q_{x^*, \ell, j^*} := \frac{p_{x^*, \ell} + r_{x^*, \ell, j^*}}{2}$ to lower bound our probability as follows: $\mathbb{P}_C(Y_{\ell, j^*} \leq X_{-\ell}) \geq \mathbb{P}_C(\{Y_{\ell, j^*} < q_{x^*, \ell, j^*} \cdot n\} \cap \{X_{-\ell} > q_{x^*, \ell, j^*} \cdot n\})$. We use again Bonferroni's inequality and an inclusion of events to get: $\mathbb{P}_C(Y_{\ell, j^*} \leq X_{-\ell}) \geq \mathbb{P}_C(\{Y_{\ell, j^*} < q_{x^*, \ell, j^*} \cdot n\}) + \mathbb{P}_C(\{X_{-\ell} > q_{x^*, \ell, j^*} \cdot n\}) - 1$. We then use Hoeffding's inequality to find these lower bounds: $\mathbb{P}_C(Y_{\ell, j^*} < q_{x^*, \ell, j^*} \cdot n) \geq 1 - e^{-2n(q_{x^*, \ell, j^*} - r_{x^*, \ell, j^*})^2}$ and $\mathbb{P}_C(X_{-\ell} > q_{x^*, \ell, j^*} \cdot n) \geq 1 - e^{-2n(p_{x^*, \ell} - q_{x^*, \ell, j^*})^2}$. Putting the last two inequalities together we get: $\mathbb{P}_C(Y_{\ell, j^*} \leq X_{-\ell}) \geq 1 - e^{-2n(p_{x^*, \ell} - q_{x^*, \ell, j^*})^2} - e^{-2n(q_{x^*, \ell, j^*} - r_{x^*, \ell, j^*})^2}$. Coming back to the first work of the proof we have: $\mathbb{P}_C(S_{2PW-H}(x^*, \ell)) \geq 1 - (m-2)(e^{-2n(p_{x^*, \ell} - q_{x^*, \ell, j^*})^2} + e^{-2n(q_{x^*, \ell, j^*} - r_{x^*, \ell, j^*})^2})$. Finally, since q_{x^*, ℓ, j^*} is defined as the middle between $p_{x^*, \ell}$ and r_{x^*, ℓ, j^*} , we can simplify the inequality: $\mathbb{P}_C(S_{2PW-H}(x^*, \ell)) \geq 1 - 2(m-2)e^{-2n(p_{x^*, \ell} - q_{x^*, \ell, j^*})^2}$. \square

We can thus deduce the same lower bound for the probability of existence of a successful unrestricted poll manipulation.

Corollary 6. For a balanced culture $C(n, \Pi^m)$, the probability of success of an unrestricted poll manipulation is as follows: $\mathbb{P}_C(S) \geq 1 - 2(m-2)(e^{-2n(p_{x^*, \ell} - q_{x^*, \ell, j^*})^2})$.

Our next theorem considers the asymptotic case and shows the convergence of the lower bound probability toward 1 when n becomes large. Since the number of voters is typically large in political elections, this shows an important susceptibility to poll manipulation.

Theorem 7. For a balanced culture $C(n, \Pi^m)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of the sub-heuristic $2PW-H(x^*, \ell)$, and thus of the Global Heuristic, tends toward 1, i.e., $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_{2PW-H}(x^*, \ell)) = 1$ and thus $\lim_{n \rightarrow \infty} \mathbb{P}_C(S) = 1$.

Proof. We use the lower bound from Theorem 5 to deduce the convergence toward 1 of this probability. In fact it is enough to pass to the limit on both sides in n the number of voters. The only tricky point might be when $p_{x^*, \ell} = q_{x^*, \ell, j^*}$. However, this situation can happen only when the culture puts positive probability only on preference orders whose top can only be x^* or j and in an equal manner, which is not possible by natural assumption on the culture. \square

Observe that the quantities $p_{x^*, \ell}$ and $q_{x^*, \ell, j}$ from Theorem 5 are constants and different, we thus have exponentially fast convergence toward 1 for the probability of success of $2PW-H(x^*, \ell)$ w.r.t. the number of voters. To give a quick intuition, observe that for $m = 5$

and $n = 50$, we get a lower bound of 0.82 and for $m = 5$ and $n = 100$, we already have a lower bound of 0.99 which is very fast!

Beyond this general result on balanced cultures, the goal would be to capture realistic cultures regarding real elections [27]. From Propositions 1–3, we can derive the following corollary which shows that our general result covers very natural concrete cultures.

Corollary 8. For a culture $C(n, \Pi^m)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of $2PW-H(x^*, \ell)$, and thus of the Global Heuristic, tends toward 1, i.e., $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_{2PW-H}(x^*, \ell)) = 1$ and thus $\lim_{n \rightarrow \infty} \mathbb{P}_C(S) = 1$, when:

- C corresponds to the impartial culture, or
- C is a single-peaked culture and x^* is not an extreme candidate or x^* is extreme but $\mathbb{P}_C(\succ_i | \text{worst}_{\succ_i} = x^*) \leq \frac{1}{2}$, which includes Walsh's and Conitzer's cultures, or
- C corresponds to a Mallows culture $\mathcal{M}^{\phi, \sigma}$ where $x^* \neq \text{worst}_{\sigma}$.

Our results show that even if the poll manipulation problem is hard, it is very likely for the polling institute to efficiently and successfully control the election, under natural preference distributions. However, the hypothesis that allows to send any score is questionable since the polling institute might be forced to meet some legal quality standards or to maintain voter trust by sending a reasonable score.

5 The Restricted Poll Manipulation Problem

This section is devoted to the study of the manipulation problem in its restricted version i.e., the polling institute is restricted in its ability to lie about the scores and can only send a score vector from I_k .

The restricted poll manipulation problem is known to be NP-hard [3]. However, one could hope to get a fixed-parameter tractable algorithm w.r.t. the maximum allowed distance k to the truthful scores. We show below that such an efficient algorithm is unlikely to exist since we prove that the problem is W[1]-hard.

Theorem 9. The restricted manipulation problem is W[1]-hard.

Sketch of proof. From an instance $(G = (V, E), k)$ of k -Clique, known to be W[1]-complete [9], where $n := |V|$ and, w.l.o.g., $2 < k < n - 1$, we construct an instance of the restricted manipulation problem as follows.

For each vertex $v_i \in V$, we create a candidate v_i , and for each edge $\{v_i, v_j\} \in E$, we create a candidate e_{ij} (we suppose $i < j$ for this notation). We add three other candidates w , t , and z .

Let $K := (n-k)k$. For each vertex $v_i \in V$, we create k voters U_i^ℓ for $\ell \in [k]$, and $K - 1 - \delta(v_i)$ voters D_i^ℓ for $\ell \in [K - 1 - \delta(v_i)]$ (by our assumption on k , this quantity cannot be negative), where $\delta(v_i)$ denotes the degree of vertex v_i in G .

For each edge $\{v_i, v_j\} \in E$, we create two voters F_{ij}^i and F_{ij}^j , and $K - 2$ voters E_{ij}^ℓ for $\ell \in [K - 2]$. Finally, we add K voters T^ℓ for $\ell \in [K]$ and $K - 1$ voters Z^ℓ for $\ell \in [K - 1]$.

The preferences of the voters over the candidates are described below, for each $i \in [n]$, and each $\{v_p, v_q\} \in E$:

$$\begin{array}{ll} U_i^\ell: & w \succ v_i > z \succ \{v_j\}_{j \neq i} \succ \{e_{r,s}\}_{\{r,s\} \ni i} \succ t & \text{for } \ell \in [k] \\ F_{pq}^p: & v_p > e_{pq} > z \succ w \succ \{v_j\}_{j \neq p} \succ \{e_{r,s}\}_{\{r,s\} \neq \{p,q\}} \succ t \\ F_{pq}^q: & v_q > e_{pq} > z \succ w \succ \{v_j\}_{j \neq q} \succ \{e_{r,s}\}_{\{r,s\} \neq \{p,q\}} \succ t \\ D_i^\ell: & v_i > z \succ w \succ \{v_j\}_{j \neq i} \succ \{e_{r,s}\}_{\{r,s\} \ni i} \succ t & \text{for } \ell \in [K - 1 - \delta(v_i)] \\ T^\ell: & t \succ z \succ w \succ \{v_j\}_j \succ \{e_{r,s}\}_{\{r,s\}} & \text{for } \ell \in [K] \\ E_{pq}^\ell: & e_{pq} > z \succ w \succ \{v_j\}_j \succ \{e_{r,s}\}_{\{r,s\} \neq \{p,q\}} \succ t & \text{for } \ell \in [K - 2] \\ Z^\ell: & z \succ w \succ \{v_j\}_j \succ \{e_{r,s}\}_{\{r,s\}} \succ t & \text{for } \ell \in [K - 1] \end{array}$$

Finally, the tie-breaking rule is as follows: $z \triangleright t \triangleright \dots \triangleright w$.

The winner of the election with the truthful ballot profile is candidate w . The details of the scores for this truthful ballot profile are given in the second column of Table 2.

Table 2. Candidates' scores in the complexity proof of Theorem 9

candidate	initial score	announced score	score after manipulation
v_i ($i \in [n]$)	$K - 1$	K if $v_i \in S$ $K - 1$ otherwise	K if $v_i \in S$ $K - 1$ otherwise
e_{ij} ($\{v_i, v_j\} \in E$)	$K - 2$	K if $v_p, v_q \in S$ $K - 2$ otherwise	K if $v_p, v_q \in S$ $K - 2$ otherwise
w	kn	K	K
t	K	$K - 1$	K
z	$K - 1$	K	$K - 1$
winner	w	z	t

We claim that G admits a clique of size k iff we can force the election of candidate t by announcing scores which differ from the truthful ones by at most $k^2 + 1$ vote changes.

Suppose first that there exists a subset of vertices $S \subseteq V$ such that S is a k -Clique of G , i.e., $|S| = k$ and $\{v_i, v_j\} \in E$ for every $v_i, v_j \in S$. Consider manipulated communicated scores which differ from the sincere ones by taking k^2 votes initially given to w in order to give one additional vote to each $v_i \in S$ (there are k such candidates) and two additional votes to each e_{ij} such that $v_i, v_j \in S$ (there are $\frac{k(k-1)}{2}$ such candidates), and finally by taking one vote initially given to t in order to give it to z . In total, these scores differ from the sincere ones by exactly $k^2 + 1$ vote changes, they are summarized in the third column of Table 2. One can prove that these communicated scores trigger deviations which lead to the final scores presented in the fourth column of Table 2, where t is the winner.

Suppose now that there exist communicated scores such that candidate t becomes the winner after all voters' deviations. One can show that the only possibility for communicated scores to lead to the victory of the target candidate t is to announce candidate z the winner and, as potential winners, k candidates v_i , as well as $k(k-1)/2$ candidates e_{pq} , such that for each potential winner v_i , there are $k-1$ potential winners e_{ij} (or e_{ji}) corresponding to edges incident to v_i . \square

Nevertheless, we prove below that the restricted poll manipulation problem can be efficiently solved if the parameter k of the maximum distance to the truthful scores is a constant.

Proposition 10. *The restricted manipulation problem is in XP w.r.t. the maximum distance k to the truthful scores. More precisely, it can be solved by an algorithm which runs in time $\Theta(m^{2k+1} \cdot n)$.*

Sketch of proof. We provide an upper bound of m^{2k} to $|I_k|$. It is then enough to visit every score vector of I_k to check whether it leads to the victory of the target candidate. \square

However, the previous result cannot be used if k is large and does not tell whether there actually exists a successful manipulation. We thus use a probabilistic approach to analyze the possibility of poll manipulation. Let S_k denote the event of success for the restricted poll manipulation where k denotes the maximum allowed distance to the truthful scores. We first prove that when k is small compared to \sqrt{n} , the restricted poll manipulation tends to be impossible.

Theorem 11. *For any culture $C(n, \Pi^m)$, if the maximum distance k to the truthful scores is such that $k = o(\sqrt{n})$ and the target candidate x^* is not winning in the initial score, then the probability of existence of a successful poll manipulation to elect x^* tends toward zero, i.e., $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_k) = 0$.*

Sketch of proof. We first identify the law of the truthful scores.

Observation 12. *For a culture $C(n, \Pi^m)$, the truthful scores s^T follow a multinomial law $\text{Multi}(q, n)$ where $q = (q_1, \dots, q_m)$ and $q_j := \mathbb{P}_C(\{\mathcal{W}_P(s^T) = j\})$, for every $j \in M$.*

The truthful scores follow a multinomial law because there are n voters' preferences drawn independently at random with the same law, and we have m possibilities for the most preferred candidate of each voter, and these are the only necessary elements to compute scores s^T . Let c^* be the truthful winner, i.e., $c^* := \mathcal{W}_P(b^T)$. Informally, a necessary condition for the existence of a successful manipulation with the two-candidate heuristic is that there is at least one candidate that is sufficiently close to the winner. The pair of candidates would then be this candidate and the current winner. Of course, this is not necessarily sufficient, as the pair may not be the right one. However, we will see that this necessary condition occurs with probability 0, and that is enough for us to conclude. We then write $S_k \subset \{\bigcup_{z \neq c^*} \{|s_{c^*}^T - s_z^T + \mathbb{1}_{c^* > z}| \leq k\}\}$.

We will analyze the probability of the second event to get an upper bound on the probability of success of the restricted poll manipulation problem. By using Observation 12 and a central limit theorem for multinomial law s^T , we get: $\frac{1}{\sqrt{n}}(s^T - nq) \xrightarrow[n \rightarrow +\infty]{} \mathcal{N}(0; K)$, where $K_{i,j} = q_i \delta_{i,j} - q_i q_j$, for every $1 \leq i, j \leq m$, with $\delta_{i,j} = 1$ if $i = j$ and $\delta_{i,j} = 0$ otherwise. We denote $\mathcal{N}(0; K) = (\mathcal{N}_1, \dots, \mathcal{N}_m)$ and remark that each \mathcal{N}_j follows a Gaussian law. Using the previous point we show that for any $z \in M \setminus \{x^*\}$, we have: $\lim_{n \rightarrow +\infty} \mathbb{P}_C(|s_{c^*}^T - s_z^T + \mathbb{1}_{c^* > z}| \leq k) = 0$. It follows for the probability of the success event that $\lim_{n \rightarrow +\infty} \mathbb{P}_C(S_k) \leq \lim_{n \rightarrow +\infty} \mathbb{P}_C(\bigcup_{z \neq c^*} |s_{c^*}^T - s_z^T + \mathbb{1}_{c^* > z}| \leq k) \leq \lim_{n \rightarrow +\infty} \sum_{z \neq c^*} \mathbb{P}_C(|s_{c^*}^T - s_z^T + \mathbb{1}_{c^* > z}| \leq k) = 0$. We then get: $\lim_{n \rightarrow +\infty} \mathbb{P}_C(S_k) = 0$, which concludes the proof. \square

Then, we get immediately the following corollary if we include the case where x^* might win in the initial poll, because it is always possible to communicate scores that keep the same winner.

Corollary 13. *For any culture $C(n, \Pi^m)$, if the maximum distance k to the truthful scores is such that $k = o(\sqrt{n})$, then $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_k) = \mathbb{P}_C(\{\mathcal{W}_P(s^T) = x^*\})$.*

We might note that, e.g., $\mathbb{P}_C(\{\mathcal{W}_P(s^T) = x^*\}) \approx \frac{1}{m}$ when considering the impartial culture.

We now focus on a case where poll manipulation can be successful, and prove that we can even efficiently compute it, thanks to an adaptation of the global heuristic where the sub-heuristic to call is Restricted 2PW-H(x^*, ℓ) which, starting from s^T , tries to announce ℓ as the winner and x^* as the only other potential winner, while taking into account the maximum allowed distance k . Let $S_{\text{Restr-2PW-H}(x^*, \ell)}$ denote the event of success of this sub-heuristic.

Theorem 14. *For a balanced culture C , if the maximum distance k to the truthful scores is such that $\mathcal{K} \cdot n = o(k)$ where $\mathcal{K} = \max_z (q_{c^*} - q_z)$ and $c^* := \mathcal{W}_P(b^T)$, then there exists $\ell \in B_C(x^*)$ such that the probability of success of Restricted 2PW-H(x^*, ℓ) tends toward 1, i.e., $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_{\text{Restr-2PW-H}(x^*, \ell)}) = 1$ and thus $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_k) = 1$.*

Proof. We can first observe that each score the polling institute may send can be summarized by its set of potential winners and its winner, since two announced scores with the same potential winners and winner produce the same voters' deviations. A type $\mathcal{T}(s)$ for a score vector s is thus defined as a pair $(PW, w) \in 2^M \times M$ where $w \in PW$, representing its potential winners and its winner. The set of all possible score types is denoted by \mathcal{T} . We will then show that: $\mathbb{P}_C(\{\bigcup_{s \in I_k} \mathcal{T}(s) = \mathcal{T}\}) = 1$. Let c^* be the truthful winner, i.e., $c^* := \mathcal{W}_P(b^T)$. Informally, a sufficient condition for the existence of a strategy of each type is that all candidates are sufficiently close to

the winner. More precisely, we would like them all to be closer than $\frac{k}{m}$, so that the cost of making potential winners any pair of candidates never exceeds k . We then get: $\{\bigcup_{z \neq c^*} \{ |s_{c^*}^T - s_z^T + \mathbb{1}_{c^* \triangleright z}| < \frac{k}{m} \}\} \subset \{\bigcup_{s \in I_k} \mathcal{T}(s) = \mathcal{T}\}$. We again use the same technique adding and subtracting $n \cdot q_{c^*}$ and $n \cdot q_z$ and a central limit theorem on the truthful scores s^T following a multinomial law (Observation 12). However, we have this time a remaining term $\sqrt{n}(q_{c^*} - q_z)$ that is bounded by assumption ($\mathcal{K} \cdot n = o(\sqrt{k})$). We then get: $\lim_{n \rightarrow +\infty} \mathbb{P}_C(|s_{c^*}^T - s_z^T + \mathbb{1}_{c^* \triangleright z}| < \frac{k}{m}) = 1$. Since a countable intersection of events of probability 1 is of probability 1, we have: $\lim_{n \rightarrow +\infty} \mathbb{P}_C(\bigcap_{z \neq c^*} \{ |s_{c^*}^T - s_z^T + \mathbb{1}_{c^* \triangleright z}| < \frac{k}{m} \}) = 1$. Then, we have: $\mathbb{P}_C(\bigcup_{s \in I_k} \mathcal{T}(s) = \mathcal{T}) = 1$. Using the fact that a successful strategy exists in the unrestricted case and since all strategies are accessible, we get: $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_{\text{Restr-2PW-H}}(x^*, \ell)) = 1$. Therefore, we have also: $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_k) = 1$. \square

The case $k = \alpha \cdot n$ with $\alpha \in]0, 1]$, when α is large enough, is included in Theorem 14. This has a clear interpretation: if the polling institute is allowed to lie by a fraction α on scores then we will fall in the manipulation regime for a sufficiently large number of voters.

Like for the unrestricted problem, the general result of Theorem 14 holds for the concrete cultures mentioned in Section 3.

Corollary 15. *For a culture C and $\mathcal{K} \cdot n = o(k)$ where $\mathcal{K} = \max_z(q_{c^*} - q_z)$ and $c^* := \mathcal{W}_P(b^T)$, there exists $\ell \in B_C(x^*)$ such that the probability of success of Restricted 2PW-H(x^*, ℓ) tends toward 1, i.e., $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_{\text{Restr-2PW-H}}(x^*, \ell)) = 1$ and thus $\lim_{n \rightarrow \infty} \mathbb{P}_C(S_k) = 1$, when:*

- C corresponds to the impartial culture, or
- C is a single-peaked culture and x^* is not an extreme candidate or x^* is extreme but $\mathbb{P}_C(\succ_i | \text{worst}_{\succ_i} = x^*) \leq \frac{1}{2}$, which includes Walsh's and Conitzer's cultures, or
- C corresponds to a Mallows culture $\mathcal{M}^{\phi, \sigma}$ where $x^* \neq \text{worst}_{\sigma}$.

6 Toward a Generalization of Strategic Behavior

Until now, we only considered strategic moves from pivotal voters. However, one can argue that voters might want to deviate when they are close enough to be pivotal. Such a strategic behavior can be captured by considering pivotal thresholds $p_i \in \mathbb{N}$ for each voter i , as done by Wilczynski [29] in an idea close to local-dominance [22]. This slightly modifies the definition of potential winners: A candidate y is a *general potential winner* for voter i w.r.t. score s if i believes that adding p_i votes to y will make candidate y the new winner, i.e., $s_{\mathcal{W}_P(s-i)}^{-i} - s_y^{-i} + \mathbb{1}_{\mathcal{W}_P(s-i) \triangleright y} \leq p_i$. We denote PW_i^{s, p_i} the set of general potential winners for i w.r.t. score s . The definition of best response naturally follows by considering general potential winners. Our initial setting corresponds to the case where $p_i = 1$.

Let us first analyze the impact of pivotal thresholds on strategic voting. For this purpose, we suppose that the polling institute is sincere and sends truthful scores $s = s^T$, and that all thresholds are equal and denoted by p , i.e., $p_i = p$, for every voter i . Let us define the expected proportion of strategic voters \mathcal{P}_{SV} w.r.t. culture $C(n, \Pi^m)$, n , m , and p . Let U_i^p denote the event where the top candidate of voter i is not a general potential winner for i , i.e., $U_i^p = \{\text{top}_{\succ_i} \notin PW_i^{s, p}\}$, and D_i^p the event where voter i could favor a potential winner other than the current winner, that she prefers to it, i.e., $D_i^p = \{\exists w \in M \setminus \{\text{top}_{\succ_i}\} : w \succ_i \mathcal{W}_P(s) \text{ and } w \in PW_i^{s, p}\}$. By definition, the proportion of strategic voters counts the voters for who the two events are true, i.e., $\mathcal{P}_{SV}(C, n, m, p) = \mathbb{E}_C[\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{U_i^p \cap D_i^p}]$.

The following proposition provides several insights on the proportion of strategic voters at the limits, by showing that the variations of the dependent events U_i and D_i are opposed with respect to p .

- Proposition 16.** 1. $U = \mathbb{E}_C[\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{U_i^p}]$ is decreasing w.r.t. p .
 2. $D = \mathbb{E}_C[\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{D_i^p}]$ is increasing w.r.t. p .
 3. $\mathcal{P}_{SV}(C, n, m, p) \leq \min(U, D)$.
 4. $\lim_{p \rightarrow +\infty} \mathcal{P}_{SV}(C, n, m, p) = 0$ and $\mathcal{P}_{SV}(C, n, m, 0) = 0$.
 5. $\lim_{n \rightarrow +\infty} \mathcal{P}_{SV}(C, n, m, p) = 0$ when p is fixed.

Although the previous proposition helps to better understand the proportion of strategic voters at the limits, it is still difficult to exactly determine the behavior for other values of p , in particular when \mathcal{P}_{SV} is maximum, because of the dependency between U and D .

Let us now analyze the poll manipulation problems. Let S^G (resp., S_k^G) denote the associated event of success for the unrestricted (resp., restricted) problem with generalized strategic behavior.

Proposition 17. *For a balanced culture $C(n, \Pi^m)$, $p > 0$ and $p = o(n)$, we have $\mathbb{P}_C(S^G) \geq \mathbb{P}_C(S)$ and $\mathbb{P}_C(S_k^G) \geq \mathbb{P}_C(S_k)$.*

It follows from Proposition 17 that Theorems 5 and 7, for successful unrestricted poll manipulation, also hold under a generalized strategic behavior with the given weak hypotheses. Similarly, the convergence result toward 1 for the probability of a successful restricted poll manipulation (Theorem 14) and the generalization of Theorem 11 can also be extended to a generalized strategic behavior.

7 Conclusion

In the context of political elections where voters are assumed to be strategic, we have studied the poll manipulation problem: *Can a polling institute lie about candidates' scores it communicates to voters in order to influence the outcome of the election?* Two variants are investigated: an unrestricted one where any scores can be sent, and a restricted one, more realistic, where the polling institute cannot announce scores too far from the reality. We show that both problems are computationally hard and answer an open question from Baumeister et al. [3]. However, we go beyond this worst-case analysis by using probabilistic tools to balance computational hardness. Under a broad condition on cultures, satisfied by many concrete preference distributions, we prove a lower bound on the probability of success of an easily computable heuristic for the unrestricted problem. This enables us to obtain a rapid convergence toward 1 of the manipulation probability, meaning that large elections are highly manipulable when the polling institute can freely manipulate without altering the trust of voters. When it may not be the case, i.e., in a restricted context, our asymptotic results show that manipulability strongly relies on whether the allowed distance to truthful scores depends on the election size. Manipulation tends to fail when this distance is negligible w.r.t. the number of voters. However, when the distance is significant, e.g., is a given proportion of the election size, which appears as a very natural assumption, efficient and successful manipulation tends to be always possible, showing that political elections are highly susceptible to poll manipulation in practice.

Future work should be devoted to studying the generalization of these results in different directions. Considering other voting rules or other types of information communicated in the poll would be natural. Another avenue of work could be to examine different strategic voting behaviors, such as local dominance [20]. Finally, a challenging future direction would be to adapt our analysis to dependent cultures such as the Pólya-Eggenberger urn [27].

References

- [1] A. Anand and P. Dey. Distance restricted manipulation in voting. *Theoretical Computer Science*, 891:149–165, 2021.
- [2] J. J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social choice and welfare*, 6:227–241, 1989.
- [3] D. Baumeister, A.-K. Selker, and A. Wilczynski. Manipulation of opinion polls to influence iterative elections. In *Proceeding of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, pages 132–140, 2020.
- [4] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
- [5] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.
- [6] V. Conitzer. Eliciting single-peaked preferences using comparison queries. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 420–427, 2007.
- [7] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI 2011)*, volume 25, pages 638–643, 2011.
- [8] P. Dey, N. Misra, and Y. Narahari. Complexity of manipulation with partial information in voting. *Theoretical Computer Science*, 726:78–99, 2018.
- [9] R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1-2):109–131, 1995.
- [10] U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. Strategic voting with incomplete information. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, pages 236–242, 2016.
- [11] P. Faliszewski and A. D. Procaccia. AI’s war on manipulation: Are we winning? *AI Magazine*, 31(4):53–64, 2010.
- [12] P. Faliszewski and J. Rothe. Control and bribery in voting. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 7, pages 146–168. Cambridge University Press, 2016.
- [13] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proceeding of the 49th Annual IEEE symposium on Foundations of Computer Science (FOCS 2008)*, pages 243–249, 2008.
- [14] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.
- [15] T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical computer science*, 38:293–306, 1985.
- [16] U. Grandi. Social choice and social networks. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 9, pages 169–184. AI Access, 2017.
- [17] N. Hazon, Y. Aumann, S. Kraus, and M. Wooldridge. On the evaluation of election outcomes under uncertainty. *Artificial Intelligence*, 189:1–18, 2012.
- [18] M. Isaksson, G. Kindler, and E. Mossel. The geometry of manipulation—a quantitative proof of the Gibbard-Satterthwaite theorem. *Combinatorica*, 32(2):221–250, 2012.
- [19] R. Meir. Iterative voting. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.
- [20] R. Meir. Strategic voting. *Synthesis lectures on artificial intelligence and machine learning*, 13(1):1–167, 2018.
- [21] R. Meir, M. Polukarov, J. Rosenschein, and N. Jennings. Convergence to equilibria in plurality voting. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2010)*, volume 24, pages 823–828, 2010.
- [22] R. Meir, O. Lev, and J. S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM conference on Economics and computation (EC 2014)*, pages 313–330, 2014.
- [23] R. B. Myerson and R. J. Weber. A theory of voting equilibria. *American Political Science Review*, 87(1):102–114, 1993.
- [24] A. D. Procaccia and J. S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proceedings of the 6th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, page 718–720, 2007.
- [25] A. Reijngoud and U. Endriss. Voter response to iterated poll information. In *Proceeding of the 4th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 635–644, 2012.
- [26] M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10(2):187–217, 1975.
- [27] S. Szufa, P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Drawing a map of elections in the space of statistical cultures. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, pages 1341–1349, 2020.
- [28] T. Walsh. Generating single peaked votes. *arXiv preprint arXiv:1503.02766*, 2015.
- [29] A. Wilczynski. Poll-confident voters in iterative voting. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2019)*, volume 33, pages 2205–2212, 2019.
- [30] L. Xia. The impact of a coalition: Assessing the likelihood of voter influence in large elections. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC 2023)*, pages 1156–1156, 2023.
- [31] L. Xia and V. Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC 2008)*, pages 109–118, 2008.