

Voter Response to Iterated Poll Information

Annemieke Reijngoud and Ulle Endriss

Institute for Logic, Language and Computation (ILLC)
University of Amsterdam

Email: annemieke@averechts.nl, ulle.endriss@uva.nl

ABSTRACT

We develop a formal model of opinion polls in elections and study how they influence the voting behaviour of the participating agents, and thereby election outcomes. This approach is particularly relevant to the study of collective decision making by means of voting in multiagent systems, where it is reasonable to assume that we can precisely model the amount of information available to agents and where agents can be expected to follow relatively simple rules when adjusting their behaviour in response to polls. We analyse two settings, one where a single agent strategises in view of a single poll, and one where multiple agents repeatedly update their voting intentions in view of a sequence of polls. In the single-poll setting we vary the amount of information a poll provides and examine, for different voting rules, when an agent starts and stops having an incentive to manipulate the election. In the repeated-poll setting, using both analytical and experimental methods, we study how the properties of different voting rules are affected under different sets of assumptions on how agents will respond to poll information. Together, our results clarify under which circumstances sharing information via opinion polls can improve the quality of election outcomes and under which circumstances it may have negative effects, due to the increased opportunities for manipulation it provides.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Theory, Economics

Keywords

Computational Social Choice, Voting Theory

1. INTRODUCTION

Voting theory has recently come to play an important role in the study of multiagent systems [2, 12]. One of the most

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

intriguing questions in voting is how agents strategise when casting their vote, in view of both their personal preferences and their beliefs about the strategies followed by others. In political elections, voters will often form their beliefs about the strategies of others on the basis of *opinion polls*. In this paper, we start from the same idea and propose a simple formal model for representing relevant information in a poll and the ways in which agents may respond to that information when deciding on what strategy to follow.

Most political elections are based on the *plurality rule*, under which a candidate obtains a point for every voter ranking her first, and the candidate with the largest number of points gets elected. In the context of this rule, the most natural type of information to publish in an opinion poll is the expected number of points for each candidate. Despite its widespread use, the plurality rule has been severely criticised for being overly simplistic and not allowing voters to adequately express their preferences. In multiagent systems, we have the opportunity to instead work with the full range of voting rules that have been proposed and studied in social choice theory [11]. This widening of the scope as far as the voting rule is concerned suggests to also consider a generalisation of the concept of opinion poll. To clarify this point, consider, for instance, the *Copeland rule*, under which you vote by submitting a strict linear order over the candidates and the score of a candidate is computed as the difference between the number of opponents she will beat in a one-to-one majority contest and the number of opponents she will lose to in such a contest (the candidate with the maximal Copeland score wins). In a poll, we could publish the (expected) Copeland score of each candidate or we could record how many copies of each possible ballot (linear order) were received. Alternatively, we could publish the majority graph (the directed graph on the set of candidates in which we include an edge from x to y if a majority of voters prefer x over y) or the weighted majority graph (in which each edge is annotated with the strength of the relevant majority).

To formally capture these ideas, we shall define a *poll information function* as a function mapping the ballots received from the voters in an opinion poll to an appropriate structure (e.g., a majority graph or a list of scores). The output of this function is then communicated to the voters, thereby providing them with partial information on the voting intentions declared by the others.

We shall analyse two scenarios. In the first, we study the incentives of a voter, who is provided with the output of a poll information function, to vote truthfully in a subsequent election. Intuitively, if a poll provides a lot of information,

then this will increase the opportunities of our voter to benefit from voting untruthfully, while a poll carrying very little information might be expected to reduce such opportunities and thereby induce truthful voting. Our results clarify, for several voting rules and several types of opinion polls, to what extent this basic intuition is in fact correct.

In the second scenario, voters participate in a sequence of polls, and in each round one of the voters can update her ballot in view of the poll information received. We consider several types of policies that a voter might use to perform this update: a *strategist* will choose a ballot such that no other ballot provides a strictly better outcome for some and at least as good an outcome for all possible ways of the other voters voting that would be consistent with the poll information received; a *pragmatist* will support her favourite candidate from a small set of, say, two front-runners; and a *truth-teller* will always vote truthfully. For different voting rules and for different combinations of these policies, we analyse whether such a system will converge to a stable state, i.e., whether it will terminate. We then observe that the “rule” we obtain when this kind of game is played for a specific voting rule and a specific set of response policies may be considered a voting rule in its own right, and we study how the properties of the original voting rule relate to the properties of this induced rule. An example for such a property is the frequency of electing a *Condorcet winner*, i.e., a candidate that would beat any other candidate in a one-to-one majority contest.

Similar phenomena have been studied before. Brams and Fishburn [1] give several examples that show, for both the plurality rule and another system known as *approval voting*, that executing a series of polls before the actual election can have both positive and negative effects in view of electing a Condorcet winner. Chopra et al. [4] give further examples, showing that a sequence of polls may or may not terminate. Meir et al. [7] identify conditions, in case the plurality rule is used, under which termination can be guaranteed. Finally, the work of Conitzer et al. [5] on the problem of strategic manipulation under partial information is closely related to the first scenario we study: an opinion poll is one way to model the partial information available to a manipulator. For comparison, most research on opinion polls in political science, such as the work of Irwin and Van Holsteyn [6], typically focuses on other concerns, e.g., the question of how polls affect the *expectations* of voters regarding election outcomes.

We proceed as follows. In Section 2 we review basic concepts from voting theory and define our model. Our results on strategic manipulation under limited information as provided by an opinion poll are presented in Section 3, and our results on voter response to iterated poll information are summarised in Section 4. Section 5 concludes. In the interest of space, we omit the proofs of some of our results. These proofs, as well as additional results, may be found in the Master’s thesis of the first author [10].

2. THE MODEL

In this section we introduce our model. We first recall relevant concepts from voting theory [11], and then define the central notion of *poll information function*.

2.1 Voting Theory

Let \mathcal{N} be a finite set of n voters (or *agents*), and let \mathcal{X} be a finite set of m alternatives (or *candidates*). Each voter i

is endowed with a *preference order* \succ_i on \mathcal{X} . To vote, each voter i submits a *ballot* b_i (which may or may not be identical to \succ_i , i.e., she may or may not vote truthfully). We adopt the standard assumption that both preferences and ballots are strict linear orders on \mathcal{X} . Let $\mathcal{L}(\mathcal{X})$ be the set of all such orders. A *profile* $\mathbf{b} = (b_1, \dots, b_n) \in \mathcal{L}(\mathcal{X})^{\mathcal{N}}$ is a vector of ballots, one for each voter. A *voting rule* $F : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$ is a function mapping ballot profiles to nonempty sets of alternatives, the election winners. Most natural voting rules may produce ties, and thus a set of winners rather than a single winner. We can obtain a *resolute* voting rule, i.e., a rule that always returns a single winner, by pairing our voting rule of choice with a *tie-breaking rule*. We restrict attention to tie-breaking rules under which ties are broken according to some fixed but arbitrary order over the alternatives. W.l.o.g., we shall assume that this fixed order is the lexicographic order $a \succ b \succ c \succ \dots$, i.e., we shall work with the *lexicographic tie-breaking rule*.

The following are examples for common voting rules [11]:

- *Positional scoring rules*: A PSR is defined by a scoring vector (s_1, \dots, s_m) with $s_1 \geq \dots \geq s_m$ and $s_1 > s_m$. An alternative receives s_j points for each voter who ranks it at the j th position. The alternative(s) with the most points win(s) the election. Important PSRs are *plurality* with scoring vector $(1, 0, \dots, 0)$, *antiplurality* (or *veto*) with scoring vector $(1, \dots, 1, 0)$, and *Borda* with scoring vector $(m-1, m-2, \dots, 0)$.
- *Copeland*: An alternative’s score is the number of pairwise majority contests it wins minus the number it loses. The alternative(s) with the highest score win(s). A *pairwise majority contest* between alternatives x and y is won by x if a majority of voters ranks x above y .
- *Maximin* (also known as *Simpson*): An alternative’s score is the lowest number of voters preferring it in any pairwise contest. The alternative(s) with the highest score win(s).
- *Bucklin*: An alternative’s score is the smallest k such that a majority of voters rank the alternative in their top k . The alternative(s) with the lowest score win(s).
- *Single transferable vote*: An STV election proceeds in rounds. In each round the alternative(s) ranked first by the fewest voters get(s) eliminated. This process is repeated until only one alternative remains (or until all remaining alternatives are ranked first equally often).

Voting rules can be categorised by their formal properties, often referred to as *axioms* [11]. A voting rule is *anonymous* if it treats all voters symmetrically. A resolute voting rule is *surjective* if for every alternative there exists a profile under which that alternative wins. A *constant* voting rule is a rule that always elects the same unique winner. If there is a voter such that her top-ranked alternative is always the unique winner, then the voting rule is *dictatorial*. Otherwise it is *non-dictatorial*. A voting rule is *unanimous* if it elects (only) alternative x whenever x is ranked first by all voters. A voting rule satisfies the *Pareto condition* if it does not return an alternative y that is ranked below some other alternative x by all voters. Note that any Pareto efficient rule is also unanimous (but not *vice versa*).

Finally, a voting rule is *Condorcet-consistent* if it elects (only) the Condorcet winner whenever it exists, and it is *strongly Condorcet-consistent* if it elects (only) the full set of weak Condorcet winners whenever that set is nonempty.

A *weak Condorcet winner* is an alternative that does not lose any pairwise majority contest, although it may tie some. A *Condorcet winner* wins any pairwise majority contest. Note that (weak) Condorcet winners only exist for some profiles. If a Condorcet winner does exist, then it must be unique, while there can be several weak Condorcet winners.

2.2 Polls and Poll Information Functions

In our model of an *opinion poll*, each voter is asked for her ballot.¹ We call the resulting ballot profile a *poll profile*. Often we would not want to communicate the whole poll profile to the electorate, e.g., to respect the privacy of voters or because it would be computationally too expensive to do so. Let \mathcal{I} be the set of all possible pieces of poll information that we might want to communicate to the electorate in view of a given poll profile. A *poll information function* (PIF) is a function $\pi : \mathcal{L}(\mathcal{X})^N \rightarrow \mathcal{I}$ mapping poll profiles to elements of \mathcal{I} . Here are some natural choices for \mathcal{I} and the corresponding PIF π :

- *Profile*: The profile-PIF simply returns the full input profile: $\pi(\mathbf{b}) = \mathbf{b}$.
- *Ballot*: The ballot-PIF returns a vector recording how often each ballot occurs in the input profile.
- *(Weighted) Majority Graph*: The (W)MG-PIF returns the (weighted) majority graph of the input profile. A *majority graph* is a directed graph in which each node represents an alternative. There is an edge (x, y) from x to y if x wins their pairwise majority contest. In a *weighted majority graph*, each edge (x, y) is labelled with the difference in number between voters ranking x above y and voters ranking y above x .
- *Score*: Given a voting rule F , the corresponding score-PIF returns for each alternative its score under the input profile according to F . F should assign points to each alternative for this PIF to be well-defined.
- *Rank*: Given a voting rule F , the corresponding rank-PIF returns the rank of each alternative under the input profile according to F . F should rank all alternatives for this PIF to be well-defined.
- *Winner*: Given a voting rule F , the corresponding winner-PIF returns the winning alternative(s) under the input profile according to F : $\pi(\mathbf{b}) = F(\mathbf{b})$.
- *Zero*: The zero-PIF does not provide any information, i.e., it simply returns a constant value: $\pi(\mathbf{b}) = 0$.

Upon receiving the signal $\pi(\mathbf{b})$, and assuming she knows how π is defined, what can voter i infer about the poll profile \mathbf{b} ? Of course, she knows her own ballot b_i with certainty. So, what can she infer about the remainder of the profile, \mathbf{b}_{-i} ? We call the set of (partial) profiles that voter i must consider possible in view of the information she holds after receiving $\pi(\mathbf{b})$ her *information set*. It is defined as follows:

$$\mathcal{W}_i^{\pi(\mathbf{b})} := \{\mathbf{c}_{-i} \in \mathcal{L}(\mathcal{X})^{N \setminus \{i\}} \mid \pi(b_i, \mathbf{c}_{-i}) = \pi(\mathbf{b})\}$$

We may think of $\pi(\mathbf{b})$ as the actual world and of $\mathcal{W}_i^{\pi(\mathbf{b})}$ as the set of possible worlds that are consistent with i 's knowledge in world $\pi(\mathbf{b})$. Indeed, \mathcal{W} satisfies the basic properties we would expect from a knowledge operator:

¹In most real-world opinion polls, pollsters do not ask *all* voters for their opinion. We can simulate this in our abstract model by simply dropping information on some of the voters before communicating the poll to the electorate.

- (REF) $\mathbf{b}_{-i} \in \mathcal{W}_i^{\pi(a_i, \mathbf{b}_{-i})}$
- (SYM) if $\mathbf{b}_{-i} \in \mathcal{W}_i^{\pi(a_i, \mathbf{c}_{-i})}$, then $\mathbf{c}_{-i} \in \mathcal{W}_i^{\pi(a_i, \mathbf{b}_{-i})}$
- (TRA) if $\mathbf{b}_{-i} \in \mathcal{W}_i^{\pi(a_i, \mathbf{c}_{-i})}$ and $\mathbf{c}_{-i} \in \mathcal{W}_i^{\pi(a_i, \mathbf{d}_{-i})}$, then $\mathbf{b}_{-i} \in \mathcal{W}_i^{\pi(a_i, \mathbf{d}_{-i})}$

Axiom (REF) simply states that the actual poll profile is always part of every voter's information set. Axioms (SYM) and (TRA) together express that whenever a voter considers some ballot profile possible, then that profile would also induce her current information set. For a discussion of the knowledge-theoretic properties of polls in view of strategic voting we refer to the work of Chopra et al. [4].

We define the degree of "informativeness" of a PIF in terms of the information sets it induces:

DEFINITION 1. A PIF π is said to be **at least as informative** as another PIF σ , if $\mathcal{W}_i^{\pi(\mathbf{b})} \subseteq \mathcal{W}_i^{\sigma(\mathbf{b})}$ for all poll profiles $\mathbf{b} \in \mathcal{L}(\mathcal{X})^N$ and all voters $i \in N$.

We note that Conitzer et al. [5] work with a similar notion of information set, except that they do not require an information set to be induced by poll information, but rather permit any set of conceivable profiles to form the information set of a given voter. Moreover, they do not include a voter's own ballot in her information set. There are also interesting connections to the work of Chevaleyre et al. [3] on the compilation complexity of voting rules: their *compilation functions* are the same types of functions as our PIFs.

3. RESPONSE TO A SINGLE POLL

In this section we analyse the case where, on the basis of an opinion poll, a single voter chooses to vote strategically.

3.1 Manipulation wrt. Poll Information

Suppose we run an opinion poll and communicate the result to voter i using PIF π (and suppose i did vote truthfully in the poll) and we then run the actual election. Will i have an incentive to manipulate and vote untruthfully, assuming she believes that the other voters will not change their ballot? We assume she does, if there is a scenario consistent with the poll information she received that will result in a better election outcome for her and there is no scenario where she will do worse than when voting truthfully.

DEFINITION 2. Let π be a PIF. Given a resolute voting rule F , a voter i , and a profile \mathbf{b} with $b_i = \succ_i$, we say that i has an **incentive to π -manipulate** using ballot c_i^* , if $F(c_i^*, \mathbf{c}_{-i}) \succ_i F(\succ_i, \mathbf{c}_{-i})$ for some profile $\mathbf{c}_{-i} \in \mathcal{W}_i^{\pi(\mathbf{b})}$ and $F(c_i^*, \mathbf{c}_{-i}) \succeq_i F(\succ_i, \mathbf{c}_{-i})$ for all other profiles $\mathbf{c}_{-i} \in \mathcal{W}_i^{\pi(\mathbf{b})}$.

In above definition, \succeq_i is the reflexive closure of \succ_i and both relations are extended from elements of \mathcal{X} to singleton sets over \mathcal{X} in the natural manner. $F(\succ_i, \mathbf{c}_{-i})$ denotes the winning singleton under voting rule F when everyone votes as in profile \mathbf{c} , while voter i votes according to \succ_i , etc.

A voting rule is **susceptible to π -manipulation** if there is a voter who has an incentive to π -manipulate (for some \mathbf{b} and some c_i^*). If a resolute voting rule is not susceptible to π -manipulation, then it is **immune to π -manipulation**.

LEMMA 1. If a PIF π is at least as informative as another PIF σ , then any resolute voting rule that is susceptible to σ -manipulation is also susceptible to π -manipulation.

PROOF. See Reijngoud [10]. \square

As an immediate corollary to Lemma 1, we obtain that, if π is at least as informative as σ , then any resolute voting rule that is immune to π -manipulation is also immune to σ -manipulation. In the sequel, we shall prove several susceptibility and immunity results for specific PIFs. Lemma 1 shows how such results can be transferred to other PIFs.

3.2 Susceptibility Results

When π is the profile-PIF, returning the full poll profile, then our notion π -manipulation reduces to the standard notion of manipulability used in voting theory. The seminal Gibbard-Satterthwaite Theorem [11] states that any resolute voting rule for three or more alternatives that is surjective and nondictatorial will be susceptible to manipulation. We shall now prove a simple generalisation of this theorem, relating the notion of π -manipulability applied and the informational requirements of the voting rule used.

Not every voting rule requires all information a ballot profile supplies to compute the winners. For the plurality rule, for example, it suffices to give for each alternative the number of ballots in which it is ranked first. For a given PIF $\pi : \mathcal{L}(\mathcal{X})^{\mathcal{N}} \rightarrow \mathcal{I}$, we say that a voting rule F is *computable from π -images* if there exists a function $H : \mathcal{I} \rightarrow 2^{\mathcal{X}} \setminus \{\emptyset\}$ such that $F = H \circ \pi$. We furthermore say that F is *strongly computable from π -images* if it is computable from π -images and $\pi(\mathbf{b}) = \pi(b_i, \mathbf{c}_{-i})$ entails $F(c_i, \mathbf{b}_{-i}) = F(\mathbf{c})$ for any two profiles \mathbf{b} and \mathbf{c} , i.e., upon learning $\pi(\mathbf{b})$ a voter i can compute the winners for *any* way of voting herself (rather than just for b_i). For example, the Copeland rule is computable but not strongly computable from MG-information (i.e., from images under the MG-PIF), while it is strongly computable from WMG-information. Furthermore, any anonymous voting rule is strongly computable from ballot-information.

THEOREM 1. *Let π be a PIF. When $m \geq 3$, any resolute voting rule that is surjective, nondictatorial, and strongly computable from π -images is susceptible to π -manipulation.*

PROOF. Let π be a PIF with range \mathcal{I} and let F be a resolute voting rule meeting above conditions. From the Gibbard-Satterthwaite Theorem it follows that F is susceptible to profile-manipulation, i.e., there exist a profile \mathbf{b} , a voter i , and a ballot c_i^* such that $F(c_i^*, \mathbf{b}_{-i}) \succ_i F(\succ_i, \mathbf{b}_{-i})$. Since F cannot differentiate between profiles that produce the same \mathcal{I} -structure, we get $F(\succ_i, \mathbf{c}_{-i}) = F(\succ_i, \mathbf{b}_{-i})$ for any \mathbf{c}_{-i} with $\pi(\succ_i, \mathbf{c}_{-i}) = \pi(\succ_i, \mathbf{b}_{-i})$. As F is *strongly computable from π -images*, this entails $F(c_i^*, \mathbf{c}_{-i}) = F(c_i^*, \mathbf{b}_{-i})$ for any \mathbf{c}_{-i} with $\pi(\succ_i, \mathbf{c}_{-i}) = \pi(\succ_i, \mathbf{b}_{-i})$. Hence, F is susceptible to π -manipulation. \square

The conditions of Theorem 1 are not *necessary* for susceptibility. There are resolute voting rules that are surjective, nondictatorial, and susceptible to π -manipulation, yet not computable from π -images, as our next result shows.

THEOREM 2. *When $m \geq 3$ and n is even, any strongly Condorcet-consistent voting rule, paired with the lexicographic tie-breaking rule, is susceptible to MG-manipulation.*

PROOF. Let \mathcal{X} , \mathcal{N} and F satisfy above conditions and let π be the MG-PIF. We construct a profile with three weak Condorcet winners such that voter i 's second favourite alternative wins if she votes truthfully, and her first favourite wins if she votes untruthfully.

Fix $a, b, c \in \mathcal{X}$ with $a \neq b \neq c$. Let $\succ_i = b \succ a \succ c \succ \mathcal{X} \setminus \{a, b, c\}$, where alternatives $\mathcal{X} \setminus \{a, b, c\}$ are ranked in any order. And let $c_i^* = b \succ c \succ a \succ \mathcal{X} \setminus \{a, b, c\}$. Let \mathbf{b}_{-i} be a profile in which $\frac{n-2}{2}$ voters submit $b \succ a \succ c \succ \mathcal{X} \setminus \{a, b, c\}$, and $\frac{n-2}{2} + 1$ voters submit $c \succ a \succ b \succ \mathcal{X} \setminus \{a, b, c\}$. Then $F(\succ_i, \mathbf{b}_{-i}) = a$ (as ties are broken in favour of a) and $F(c_i^*, \mathbf{b}_{-i}) = b$. It is not difficult to check that there is no profile $\mathbf{c}_{-i} \in \mathcal{W}_i^{\pi(\succ_i, \mathbf{b}_{-i})}$ such that $F(c_i^*, \mathbf{c}_{-i}) \prec_i F(\succ_i, \mathbf{c}_{-i})$. It follows that F is susceptible to MG-manipulation. \square

Examples for voting rules that are strongly Condorcet-consistent include the maximin-rule, but not, for instance, the (Condorcet-consistent) Copeland rule.

Our final example for a π -susceptibility result concerns PSRs. Observe that a PSR is unanimous if and only if $s_1 > s_2$ holds for the scoring vector defining it.

THEOREM 3. *When $m \geq 3$ and $n \geq 4$, any unanimous positional scoring rule, paired with the lexicographic tie-breaking rule, is susceptible to winner-manipulation.*

PROOF. Let \mathcal{X} , \mathcal{N} and F satisfy above conditions and let π be the winner-PIF wrt. F . We construct a profile where voter i 's third favourite alternative wins if she votes truthfully and her second favourite wins if she votes untruthfully.

Fix $a, b, c \in \mathcal{X}$ with $a \neq b \neq c$. Let $\succ_i = c \succ a \succ b \succ \mathcal{X} \setminus \{a, b, c\}$, where alternatives $\mathcal{X} \setminus \{a, b, c\}$ are ranked in any order. And let $c_i^* = a \succ c \succ b \succ \mathcal{X} \setminus \{a, b, c\}$. If n is odd, let \mathbf{b}_{-i} be a profile in which $\frac{n-3}{2}$ voters submit $a \succ b \succ \mathcal{X} \setminus \{a, b, c\}$, $\frac{n-3}{2}$ voters submit $b \succ a \succ \mathcal{X} \setminus \{a, b, c\}$, and the remaining two voters submit $c \succ b \succ a \succ \mathcal{X} \setminus \{a, b, c\}$ and $b \succ a \succ c \succ \mathcal{X} \setminus \{a, b, c\}$. If n is even, let \mathbf{b}_{-i} be a profile in which $\frac{n-2}{2}$ voters submit $a \succ b \succ \mathcal{X} \setminus \{a, b, c\}$, and $\frac{n-2}{2}$ voters submit $b \succ a \succ \mathcal{X} \setminus \{a, b, c\}$. Since F is unanimous, i.e., $s_1 > s_2$, we get $F(\succ_i, \mathbf{b}_{-i}) = b$ and $F(c_i^*, \mathbf{b}_{-i}) = a$. It is not difficult to check that there is no profile $\mathbf{c}_{-i} \in \mathcal{W}_i^{\pi(\succ_i, \mathbf{b}_{-i})}$ such that $F(c_i^*, \mathbf{c}_{-i}) \prec_i F(\succ_i, \mathbf{c}_{-i})$. It follows that F is susceptible to winner-manipulation. \square

3.3 Immunity Results

We now turn our attention to voting rules that are immune to certain types of manipulation. First, it is not difficult to verify that any *dictatorial* as well as any *constant* voting rule will be immune to *profile-manipulation* (and thus, by Lemma 1, also to any other form of manipulation). At the other extreme, as we shall see next, we can obtain immunity results for two large classes of voting rules with respect to the weakest form of immunity considered here, namely *zero-manipulation*. The next theorem is inspired by (and corrects a minor mistake in) a result due to Conitzer et al. [5].

THEOREM 4. *When $n \geq 3$, any strongly Condorcet-consistent voting rule, paired with the lexicographic tie-breaking rule, is immune to zero-manipulation.*

PROOF. Let \mathcal{N} and F satisfy above conditions and let π be the zero-PIF. Fix any voter i , any poll profile \mathbf{b} with $b_i = \succ_i$, and any untruthful ballot c_i^* . Since $c_i^* \neq \succ_i$, there is a pair of alternatives such that $x \succ_i y$ and $y \succ_{c_i^*} x$. Claim: there exists a profile \mathbf{c}_{-i} such that $F(c_i^*, \mathbf{c}_{-i}) \prec_i F(\succ_i, \mathbf{c}_{-i})$. If n is odd, let \mathbf{c}_{-i} be a profile in which $\frac{n-1}{2}$ voters submit $x \succ y \succ \mathcal{X} \setminus \{x, y\}$, and $\frac{n-1}{2}$ voters submit $y \succ x \succ \mathcal{X} \setminus \{x, y\}$, where alternatives $\mathcal{X} \setminus \{x, y\}$ are ranked

in any order. If n is even, let \mathbf{c}_{-i} be a profile in which $\frac{n-2}{2}$ voters submit $x \succ y \succ \mathcal{X} \setminus \{x, y\}$, and $\frac{n-2}{2}$ voters submit $y \succ x \succ \mathcal{X} \setminus \{x, y\}$, and the remaining voter submits $y \succ x \succ \mathcal{X} \setminus \{x, y\}$ in case x lexicographically precedes y and $x \succ y \succ \mathcal{X} \setminus \{x, y\}$ otherwise. Then $F(\succ_i, \mathbf{c}_{-i}) = x$ and $F(\mathbf{c}_i^*, \mathbf{c}_{-i}) = y$. Hence, for any untruthful ballot c_i^* there is a situation where i will do strictly better by voting truthfully. It follows that F is immune to zero-manipulation. \square

Conitzer et al. [5] state a slightly stronger variant of Theorem 4: any resolute voting rule that is (not necessarily strongly) Condorcet-consistent is immune to zero-manipulation. This is true for an *odd* number of voters, as may be seen by revisiting the first part of our proof. For an *even* number of voters, however, Condorcet consistency is not sufficient, as demonstrated by the following example.

EXAMPLE 1. Consider a scenario with 4 voters and 3 alternatives (a, b, c) . Suppose that voting rule F elects the Condorcet winner if one exists, and otherwise the bottom choice of voter 1. Let $\succ_1 = a \succ b \succ c$, and consider ballot $c_1^* = a \succ c \succ b$. Now there is a profile \mathbf{c}_{-1} such that voter 1 benefits from voting untruthfully, namely when the others vote $a \succ b \succ c$, $b \succ a \succ c$, and $b \succ a \succ c$. Then $F(\succ_1, \mathbf{c}_{-1}) = c$ and $F(c_1^*, \mathbf{c}_{-1}) = b$. It is not difficult to check that there is no profile \mathbf{c}_{-1} such that $F(c_1^*, \mathbf{c}_{-1}) \prec_1 F(\succ_1, \mathbf{c}_{-1})$. It follows that F is a resolute voting rule that is Condorcet-consistent and susceptible to zero-manipulation.

We stress that Theorem 4 also cannot be simplified to stating that any voting rule that always elects *some* weak Condorcet winner whenever one exists is immune to zero-manipulation.

The following theorem strengthens another result by Conitzer et al. [5], who use a bound of $n \geq 6m - 12$.

THEOREM 5. When $n \geq 2m - 2$, any positional scoring rule, paired with the lexicographic tie-breaking rule, is immune to zero-manipulation.

PROOF. Let \mathcal{N} , \mathcal{X} and F satisfy above conditions and let π be the zero-PIF. Fix any voter i and any poll profile \mathbf{b} with $b_i = \succ_i$. And fix any untruthful ballot c_i^* such that $F(c_i^*, \mathbf{c}_{-i}) \neq F(\succ_i, \mathbf{c}_{-i})$ for some profile \mathbf{c}_{-i} . Then there exists a pair of alternatives such that $x \succ_i y$ and $y \succ_{c_i^*} x$, and x 's score differs from y 's in \succ_i and c_i^* . Claim: there exists a profile \mathbf{c}_{-i} such that $F(c_i^*, \mathbf{c}_{-i}) \prec_i F(\succ_i, \mathbf{c}_{-i})$. If n is odd, let \mathbf{c}_{-i} be a profile in which $\frac{n-1}{2}$ voters submit $x \succ y \succ \mathcal{X} \setminus \{x, y\}$, and $\frac{n-1}{2}$ voters submit $y \succ x \succ \mathcal{X} \setminus \{x, y\}$, where every alternative $z \in \mathcal{X} \setminus \{x, y\}$ is ranked last by at least one voter. If n is even, let \mathbf{c}_{-i} be a profile in which $\frac{n-2}{2}$ voters submit $x \succ y \succ \mathcal{X} \setminus \{x, y\}$, and $\frac{n-2}{2}$ voters submit $y \succ x \succ \mathcal{X} \setminus \{x, y\}$, where every alternative $z \in \mathcal{X} \setminus \{x, y\}$ is ranked last by at least two voters, and the remaining voter submits the ranking that is like \succ_i but with x and y swapped in case x lexicographically precedes y and c_i^* with x and y swapped otherwise. Then $F(\succ_i, \mathbf{c}_{-i}) = x$ and $F(c_i^*, \mathbf{c}_{-i}) = y$. Hence, F is immune to zero-manipulation. Observe that the bound of $n \geq 2m - 2$ follows from our requirements on the number of voters ranking each $z \in \mathcal{X} \setminus \{x, y\}$ last. (The case with 4 voters and 3 alternatives must be checked separately.) \square

Together, Theorem 4 and Theorem 5 cover a broad range of voting rules. In particular, as is well known [11], the classes of Condorcet-consistent rules and PSRs do not overlap.

So far, all our immunity results involved either trivial voting rules (dictatorships and constant rules) or the trivial information set (for zero-manipulation). While we should not expect many positive results between these two extremes, they are not impossible to obtain either:

THEOREM 6. When $n \geq 2m - 2$, the antiplurality rule, paired with the lexicographic tie-breaking rule, is immune to winner-manipulation.

PROOF. Let \mathcal{N} and F satisfy above conditions and let π be the winner-PIF wrt. F . W.l.o.g., we may assume that voters only submit an alternative they wish to veto. Fix any voter i , any profile \mathbf{b} with b_i is voter i 's true least favourite alternative, and any ballot $c_i^* \neq b_i$. Claim: voter i never has an incentive to π -manipulate. Suppose $m \geq 3$ and $F(\mathbf{b}) = w$. If $w = b_i$, then i cannot change the outcome. If $w = c_i^*$, let \mathbf{c}_{-i} be a profile in which $n-1$ voters veto some $x \in \mathcal{X} \setminus \{b_i, w\}$, and all alternatives $x \in \mathcal{X} \setminus \{b_i, w\}$ are vetoed by at least one voter. If $w \in \mathcal{X} \setminus \{b_i, c_i^*\}$, let \mathbf{c}_{-i} be a profile in which $n-2$ voters veto some $x \in \mathcal{X} \setminus \{b_i, w\}$, and all alternatives $x \in \mathcal{X} \setminus \{b_i, w\}$ are vetoed by at least two voters, and the remaining voter vetoes w in case w lexicographically precedes b_i and some alternative $x \in \mathcal{X} \setminus \{b_i, w\}$ otherwise. Then $F(b_i, \mathbf{c}_{-i}) = w$ and $F(c_i^*, \mathbf{c}_{-i}) = b_i$. Hence, F is immune to winner-manipulation. (The case with 2 alternatives must be checked separately.) \square

THEOREM 7. When $n \geq 10$, the plurality rule, paired with the lexicographic tie-breaking rule, is immune to MG-manipulation.

PROOF. See Reijngoud [10]. The main insight is that for $n \geq 10$ the majority graph does not give enough information for a voter to infer the identity of the plurality winner. \square

4. REPEATED RESPONSE TO POLLS

In this section we study the case where voters repeatedly react to opinion polls. We assume that all voters will vote truthfully in an initial poll. Then, given the poll information communicated to the voters (using a PIF), one voter is given the opportunity to update her ballot. This process is repeated until no further voter wishes to update her ballot (or until a given maximum number of rounds has been reached). We then apply the voting rule to this final profile to determine the election winner.

An important parameter in this kind of *voting game* is given by the *response policies* that voters use to update their ballots. We shall first formulate several such policies, and then formally introduce the voting games considered here. As we shall see, any such game (the definition of which includes a voting rule F) induces a new voting rule F^t (where t is the number of rounds played), and we will analyse the properties of F^t in view of those of F .

4.1 Response Policies

In each round, a voter i who has the opportunity to update her ballot must decide what to do based on her true preference order \succ_i , her previously submitted ballot b_i , and her current information set \mathcal{W}_i . A *response policy* determines for each voter i a function $\delta_i : \mathcal{L}(\mathcal{X}) \times \mathcal{L}(\mathcal{X}) \times 2^{\mathcal{L}(\mathcal{X})^{\mathcal{N} \setminus \{i\}}} \rightarrow \mathcal{L}(\mathcal{X})$, mapping $(\succ_i, b_i, \mathcal{W}_i)$ to a new ballot b_i' . We shall work with the following policies:

- *Truth-teller*: A truth-teller always votes truthfully, i.e., $\delta_i(\succ_i, b_i, \mathcal{W}_i) = \succ_i$.
- *Strategist*: A strategist computes her best responses to a poll and uses (any) one of them. Only if her current ballot is amongst the best responses, she will always use it. If we restrict attention to the plurality rule and assume that polls give score-information, then this policy is similar to the policy used by Meir et al. [7].
- *Pragmatist*: A pragmatist cannot or does not want to compute her best response to a poll, e.g., because this takes too much effort. A k -pragmatist always moves her favourite amongst the k currently highest ranked alternatives to the first position in her ballot, without changing the relative ranking of the others. This policy is also described by Brams and Fishburn [1].

We assume any given voter will use the same policy throughout. Voters also have to decide how to vote in the first round. As mentioned above, we assume that they all choose to vote truthfully then. This is not unreasonable, given the immunity results to zero-manipulation in Section 3.3.

Note that in the framework defined here, voters only take into account information from the latest poll round. However, the framework could be extended to include previous rounds by adding the information sets induced by those rounds to the input arguments of δ_i .

4.2 Induced Voting Rules

Let us now formally define the notion of *voting game*.

DEFINITION 3. A *voting game* is a tuple $G = \langle F, \pi, \delta \rangle$, where F is a resolute voting rule, π is a PIF, and $\delta = (\delta_1, \dots, \delta_n)$ is a vector of response policies.

A voting game proceeds in rounds. Initially, all voters vote truthfully. In each subsequent round, exactly one voter changes her ballot. This voter is selected from the set of voters who wish to change. Whether or not a voter i wishes to do so depends on her response policy δ_i . At the end of each round, $\pi(\mathbf{b})$ is computed for the new poll profile \mathbf{b} and the result is communicated to all voters. For a voting game to be uniquely defined, we need to fix the order in which voters may change their ballot. Any such order is allowed (none of our results will depend on the order chosen).

A voting game G induces a new voting rule F^t when the number of rounds to be played is t .

DEFINITION 4. Let $G = \langle F, \pi, \delta \rangle$ be a voting game, and let $t \in \mathbb{N}$ be the number of rounds to be played. Then a voting rule F^t is **induced** by G by stipulating for any profile \mathbf{b} :

$$F^t(\mathbf{b}) := \left\{ x \in \mathcal{X} \mid \begin{array}{l} x \text{ is an election winner after } t \text{ rounds} \\ \text{when } \mathbf{b} \text{ is the truthful, initial profile} \end{array} \right\}$$

A game *terminates* in round t if no voter wishes to change her ballot according to her response policy at the end of t . If G always terminates after at most t rounds, then we write F^* instead of F^t . Clearly, if all voters are truth-tellers, then any voting game terminates after 0 rounds. Moreover, if F is immune to π -manipulation and all voters are strategists, then G terminates after 0 rounds as well.

Meir et al. [7] show that for any voting game G with F being the plurality rule and π the score-PIF wrt. the plurality rule, if p voters are strategists and $n-p$ voters are truth-tellers, then G terminates after at most $m \cdot p$ rounds. To be

precise, these authors show this for a specific kind of strategist response policy, namely one in which a voter, whenever her current ballot is not amongst the best responses, changes her ballot to the best response in which the next alternative to win is the one she ranks first. We obtain a similar result for an electorate composed of truth-tellers and pragmatists (as opposed to strategists). In fact, under these assumptions the result can be generalised to arbitrary PSRs:

LEMMA 2. Let $G = \langle F, \pi, \delta \rangle$ be any voting game such that F is a PSR, paired with the lexicographic tie-breaking rule, π is the rank-PIF wrt. F , and δ is a vector of p k -pragmatists and $n-p$ truth-tellers. Then G terminates after $\leq p$ rounds.

PROOF. Let G, F, π , and δ satisfy above conditions. Fix any profile \mathbf{b} . Let $H_k^t(\mathbf{b})$ be the set of k highest ranked alternatives in \mathbf{b} according to F^t . Claim: $H_k^t(\mathbf{b}) = H_k^{t+1}(\mathbf{b})$ for any number of rounds $t \in \mathbb{N}$. Suppose that voter i changes her ballot at round t . Since i is a k -pragmatist, and F is a PSR, we have that no alternative $x \in H_k^t(\mathbf{b})$ loses points and no alternative $y \in \mathcal{X} \setminus H_k^t(\mathbf{b})$ wins any. Hence, each voter will update her ballot at most once and G terminates after at most p rounds. \square

A similar argument can be used to prove that if F is the Copeland rule and all voters are k -pragmatists, then G terminates after at most n rounds. This also holds in case F is the maximin rule or the Bucklin rule. On the other hand, games defined in terms of other voting rules or other response policies need not always terminate:

EXAMPLE 2. Consider a scenario with 2 voters and 3 alternatives (a, b, c) . Let F be the Copeland rule, paired with the lexicographic tie-breaking rule. Let π be the MG-PIF. Suppose all voters are strategists. Consider the ballot profile $\mathbf{b} = (a \succ b \succ c, c \succ b \succ a)$. Then $F^0(\mathbf{b}) = a$, $F^1(\mathbf{b}) = b$, $F^2(\mathbf{b}) = a$, $F^3(\mathbf{b}) = c$, $F^4(\mathbf{b}) = a, \dots$ (voters 2 and 1 alternate moving alternative b up and down in their ballots).

4.3 Properties of Induced Voting Rules

For a given voting rule F , and under certain assumptions on the PIF used and the response policies of voters, what will be the properties of F^t (and F^* , when it is well-defined)? This is the question we shall investigate next. Specifically, we are interested in properties that *transfer* from F to F^t and F^* . Let us begin with a simple observation: If F is dictatorial, then any induced rule (for any PIF and any response policy defined here) will be dictatorial as well. More interestingly, we also obtain a transfer result for unanimity:

THEOREM 8. Let $G = \langle F, \pi, \delta \rangle$ be any voting game such that F is unanimous, π is a PIF, and δ is a vector of pragmatists, strategists and truth-tellers. Then F^t is unanimous for any $t \in \mathbb{N}$.

PROOF. Let G, F, π , and δ satisfy above conditions. Fix any profile \mathbf{b} such that there is an alternative w that is ranked first by all voters. Claim: $F^t(\mathbf{b}) = w$ for any number of rounds $t \in \mathbb{N}$. Proof by induction. Since F is unanimous, we have that $F^0(\mathbf{b}) = w$. Now, suppose that $F^t(\mathbf{b}) = w$, and that voter i wishes to change her ballot and may do so next. As no truth-teller or pragmatist who already has her favourite alternative winning will ever change her ballot, we only need to consider the case where i is a strategist. Since strategists always switch to a ballot that is at least as good

as their previous ballot for all profiles in their information set (and strictly better for some), we have that $F^{t+1}(\mathbf{b}) = w$. This proves that F^t is unanimous for any $t \in \mathbb{N}$. \square

However, the Pareto condition, which is slightly stronger than unanimity, does not always transfer:

EXAMPLE 3. Consider a scenario with 2 voters and 3 alternatives (a, b, c) . Let F be a voting rule that returns all alternatives that are not Pareto dominated, paired with the lexicographic tie-breaking rule. Let π be the winner-PIF wrt. F . Suppose all voters are strategists and let $t \geq 1$. Consider the ballot profile $\mathbf{b} = (b \succ c \succ a, c \succ a \succ b)$. Then $F^t(\mathbf{b}) = a$ (because the second voter will rank a on top, given that she has no chance to make c win, which is disadvantaged by the tie-breaking rule), even though alternative a is Pareto dominated by alternative c in profile \mathbf{b} .

We also cannot guarantee the transfer of the Pareto condition for voting games in which all voters are pragmatists. Other properties that do not always transfer are surjectivity, anonymity, and Condorcet consistency. For all of these properties there are counterexamples in which all voters are strategists and polls give ballot-information. We omit these examples due to space constraints. On the other hand, under certain conditions, Condorcet consistency does transfer:

THEOREM 9. Let $G = \langle F, \pi, \delta \rangle$ be any voting game such that F is Condorcet-consistent, π is the rank-PIF wrt. F , and δ is a vector of truth-tellers and pragmatists. Then F^t is Condorcet-consistent for any $t \in \mathbb{N}$.

PROOF. Let G, F, π , and δ satisfy above conditions. Fix any profile \mathbf{b} with a Condorcet winner w . Claim: $F^t(\mathbf{b}) = w$ for any number of rounds $t \in \mathbb{N}$. Proof by induction. Since F is Condorcet-consistent, we have that $F^0(\mathbf{b}) = w$. Now, suppose that $F^t(\mathbf{b}) = w$, and that voter i wishes to change her ballot and may do so next. Since i is a k -pragmatist for some $k \in \mathbb{N}$, and w is among the k currently highest ranked alternatives, we have that w cannot lose support in any pairwise contest with respect to its original pairwise scores. It follows that $F^{t+1}(\mathbf{b}) = w$. This proves that F^t is Condorcet-consistent for any $t \in \mathbb{N}$. \square

4.4 Condorcet Efficiency: Simulations

It is widely acknowledged that Condorcet consistency is a highly desirable property, but many important voting rules do not satisfy it [11]. The *Condorcet efficiency* of a voting rule is its tendency to elect the Condorcet winner. Theorem 9 identifies conditions under which Condorcet consistency transfers from F to F^t , but it does not say anything about the transfer of Condorcet efficiency. Brams and Fishburn [1] give several examples that show that polls can have a positive and a negative effect on the Condorcet efficiency of a voting rule. To study *how* positive or negative this effect is we shall make use of simulations.

We generated election data with n ranging from 10 to 100 in steps of 5 while keeping m fixed at 5, and m ranging from 3 to 15 in steps of 1 while keeping n fixed at 50 (using a simple program implemented in JAVA 1.6.0). For each of these combinations we generated 10,000 (truthful) ballot profiles with a Condorcet winner using the *impartial culture* (IC) assumption, which states that any permutation of alternatives is equally likely to occur as a voter’s preference order. The limitations of the IC assumption are well known [9]; in




Figure 1: Average probability of electing the Condorcet winner for 50 voters and 5 alternatives over 10,000 trials. Poll effect on Condorcet efficiency is significant ($p < 0.05$) for plurality, STV and Bucklin.

particular, we should not expect the preferences in a real-world electorate to be distributed uniformly. Nevertheless, the IC assumption is still the *de facto* standard used in social choice theory; results based on it provide an important base line and allow for direct comparison with a large number of findings documented in the literature.

Our first experiment was set up to test the effect of polls on the Condorcet efficiency of plurality, Borda, Copeland, STV, and Bucklin when polls provide (at least) rank-information, and all voters are 2-pragmatists or all voters are 3-pragmatists. That is, voters want to have as much electoral influence as they can without having to think too hard about a new ballot, unlike strategists. Note that we included the Copeland rule for comparison, but we already know that the rule itself is Condorcet-consistent, and thus its induced voting rule will be as well (cf. Theorem 9). We fixed the order in which voters may change their ballot to be the *ascending order*: voters were offered a chance to update their ballot according to their index in \mathcal{N} , beginning with the successor of the voter who was the last to change. All induced voting rules were run until termination.

Figure 1 shows the results for 50 voters and 5 alternatives; the other voter-alternative combinations showed a similar pattern, except for plurality, to which we will come back later. We used R to analyse our data [8], and McNemar’s test to determine whether the poll effect was significant. The results for no-polls vs. polls for 2-pragmatists were $p = 0$, $p = 0.13$, $p < 0.001$, and $p = 0$ for plurality, Borda, STV, and Bucklin respectively. For no-polls vs. polls for 3-pragmatists they were $p < 0.001$, $p = 0.13$, $p < 0.001$, and $p < 0.001$ for the same rules. Thus, polls had a significant positive effect on the Condorcet efficiency of plurality, STV and Bucklin, and no significant effect for Borda.

Intuitively, one can think of the pragmatist response policy as offering a Condorcet winner another chance to win if it ended up among the k highest ranked alternatives in the first round. For plurality and Bucklin we can state this intuition as a general rule: if all voters are 2-pragmatists and the Condorcet winner is among the two highest ranked alternatives in the first round, then it will always win under induced vot-




Figure 2: Average probability of electing the Condorcet winner for 50 voters and 5 alternatives over 10,000 trials. Poll effect on Condorcet efficiency is significant ($p < 0.05$) if polls give score-information.

ing rule F^* . This follows from Lemma 2 and properties of F . However, if all voters are 3-pragmatists, then this no longer holds. Generally, an alternative in a runoff between 2 alternatives has a greater chance of winning to start with than an alternative in a runoff between 3. We would therefore expect the 2-pragmatist response policy to have a greater positive effect on the Condorcet efficiency than the 3-pragmatist response policy. Indeed, our data reflect this expectation. On the other hand, as m increases, it becomes less likely that the Condorcet winner ends up amongst the two highest ranked alternatives in the first round. This would explain that for large numbers of alternatives ($m \geq 12$), the 3-pragmatist response policy had a greater positive effect on the Condorcet efficiency of plurality than the 2-pragmatist response policy.

Our results also show that polls had a greater effect on plurality and Bucklin than on STV and Borda. This might be due to the substantially lower Condorcet efficiency of these rules, leaving more room for improvement.

So, polls improve the Condorcet efficiency of the widely used plurality rule if all voters are pragmatists. What about strategists? In our second experiment we tested the effect of polls on the Condorcet efficiency of plurality under the assumption that polls give rank- or score-information and that all voters are strategists. As under rank-information the induced rule did not always terminate, we ran $t = 10,000$ poll rounds, after which 38% of all elections did terminate. The results of the second experiment are shown in Figure 2. Again, we only show the results for elections with 50 voters and 5 alternatives. McNemar’s test on paired results gave $p = 0.27$ for no-polls vs. rank-polls and $p < 0.001$ for no-polls vs. score-polls. Thus, polls had a significant positive effect on the Condorcet efficiency of plurality when voters received score-information. On the other hand, when polls only gave rank-information, we observed no significant effect on Condorcet efficiency. The latter effect, however, turned significantly positive for large m/n ratios, and significantly negative for small m/n ratios.

How can we explain this pattern? Providing a strategist with less information has two opposing effects: (1) she is more likely to update her ballot, because even if she actually

cannot make a different alternative win, she may think she can; and (2) she is less likely to update her ballot, because she is risk-averse. Which effect is stronger is difficult to predict, but our simulation results suggest that this depends at least in part on the number of voters and alternatives.

5. CONCLUSION

We have developed a framework to study the effects of polls on voting behaviour and election outcomes. We found that when voters do not have any information about the voting intentions of others, then for many voting rules they never have an incentive to vote untruthfully. However, they start having these incentives as soon as they know who is currently winning, according to the poll. This does not necessarily mean that polls have a negative effect on the election outcome. Some favourable properties of voting rules do persist, and may even be strengthened, when a rule is complemented with a series of polls to which the voters can respond.

For properties that do not persist in general, further work on simulations, similar to our study of Condorcet efficiency, will be required. Beyond this, it would be interesting to investigate whether *combinations* of properties of voting rules induce particular properties of elections with polls. Finally, it would be worthwhile to consider additional types of policies according to which voters respond to poll information and to investigate their influence on election outcomes.

6. REFERENCES

- [1] S. J. Brams and P. C. Fishburn. *Approval Voting*. Birkhäuser, 1983.
- [2] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. In *Proc. SOFSEM-2007*, 2007.
- [3] Y. Chevaleyre, J. Lang, N. Maudet, and G. Ravailly-Abadie. Compiling the votes of a subelectorate. In *Proc. IJCAI-2009*, 2009.
- [4] S. Chopra, E. Pacuit, and R. Parikh. Knowledge-theoretic properties of strategic voting. In *Proc. JELIA-2004*, 2004.
- [5] V. Conitzer, T. Walsh, and L. Xia. Dominating manipulations in voting with partial information. In *Proc. AAAI-2011*, 2011.
- [6] G. A. Irwin and J. J. M. van Holsteyn. According to the polls: The influence of opinion polls on expectations. *Pub. Opinion Quart.*, 66:92–104, 2002.
- [7] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria in plurality voting. In *Proc. AAAI-2010*, 2010.
- [8] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2011.
- [9] M. Regenwetter, B. Grofman, A. A. J. Marley, and I. M. Tsetlin. *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*. Cambridge University Press, 2006.
- [10] A. Reijngoud. *Voter Response to Iterated Poll Information*. Master’s thesis, ILLC, University of Amsterdam, 2011.
- [11] A. D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.
- [12] M. Wooldridge. *An Introduction to Multiagent Systems*. Wiley, 2nd edition, 2009.