

The Complexity of Campaigning

Cory Siler, Luke Harold Miles, and Judy Goldsmith

University of Kentucky, Lexington, Kentucky, United States
jcsi225@g.uky.edu, luke.lambda@uky.edu, goldsmit@cs.uky.edu

Abstract. In “The Logic of Campaigning”, Dean and Parikh consider a candidate making campaign statements to appeal to the voters. They model these statements as Boolean formulas over variables that represent stances on the issues, and study optimal candidate strategies under three proposed models of voter preferences based on the assignments that satisfy these formulas. We prove that voter utility evaluation is computationally hard under these preference models (in one case, $\#P$ -hard), along with certain problems related to candidate strategic reasoning. Our results raise questions about the desirable characteristics of a voter preference model and to what extent a polynomial-time-evaluatable function can capture them.

1 Introduction

In light of some fairly surprising election outcomes around the world, we are very interested in understanding how politicians construct their platforms. For instance, what motivates many candidates to speak in platitudes that reveal little information about their views? On the other hand, what motivates candidates to commit to specific and sometimes audacious policies? The focus of this paper is a logical formalism introduced by Dean and Parikh [7] (and extended by Parikh and Taşdemir [13]) that aims to explain candidates’ choices of campaign statements to make; these statements are modeled as propositional formulas over variables representing stances on issues, and Dean and Parikh consider different definitions of voters’ utility for a candidate as functions of the possible sets of policies that the candidate might implement based on these statements.

Political scientists have also taken interest in candidates’ decisions about what to say when campaigning; Petrocik [15] found empirical evidence that a candidate will try to focus on issues where the candidate has a good record and their opponents have bad records. Game theorists have a shared interest with Dean and Parikh in what might motivate a candidate to be ambiguous [1, 2, 5]. The game-theoretic models take into account the interaction between multiple candidates (and in the case of Bagdasaryan and Manzoni’s model [5], voters’ uncertainty about their most-preferred policies), but often with simplified representations of a platform (e.g., points on a one-dimensional spectrum or probability distributions over a small set of alternatives). In contrast, Dean and Parikh’s framework abstracts away details of the electoral system like multi-agent

interactions and voter strategy to focus on the implications for an individual candidate of communicating using more expressive logic-based statements.

This expressivity, however, brings computational costs. In this paper, we consider the computational complexity of problems related to voter and candidate reasoning. Although Dean and Parikh’s formulations capture many desirable characteristics of possible voters, and go a long way toward explaining the assumptions about voters that both vague and overspecific candidates might be making, our results raise some questions about whether these are the right models of how voters evaluate candidates’ platforms.

In Section 2 we introduce Dean and Parikh’s model and some computational complexity classes we will use later. We then find computational complexity results for problems related to the model — in Section 3, evaluating voters’ utility for a candidate; in Section 4, choosing campaign statements that optimize total voter utility for the candidate; and in Section 5, choosing campaign statements that motivate enough individual voters to vote for the candidate. In Section 6, we conclude with directions for future work in modeling campaigns, including questions about desirable characteristics of voter evaluation.

2 Preliminaries

2.1 Candidates, Voters, and Statements

In Dean and Parikh’s model, political views are expressed in terms of Boolean variables (atomic propositions) $\mathcal{X} = \{x_1, \dots, x_n\}$ (e.g., x_1 = “Every citizen is entitled to a free pony.”, x_2 = “Tooth-brushing should be mandatory.”, x_3 = “We must invest in zombie-based renewable energy sources.”). A candidate makes statements about their platform in the form of propositional formulas over these variables (e.g., $\neg x_1$, or $x_1 \rightarrow x_2$). The candidate’s current *theory*¹ T consists of the statements the candidate has issued so far and their logical closure. In our discussion of complexity, we will assume that T is given as a set of statements and that anything in the logical closure besides the statements themselves must be computed. In general we assume that T is self-consistent.

A voter v has a preference function $p_v : \mathcal{X} \rightarrow [-1, 1]$ indicating which direction and how strongly v stands on each issue x_i :² A negative $p_v(x_i)$ indicates that v prefers x_i to be false and a positive $p_v(x_i)$ indicates that v prefers x_i to be true, with the magnitude reflecting the strength of preference (and 0 being indifference). If, for example, v was against mandatory tooth-brushing and cared greatly about this issue, then we might have $p_v(x_2) = -0.9$. We assume that candidates have complete knowledge of the voters’ preference functions.

We let \mathcal{W} be the set of all possible assignments (*worlds*) to the variables \mathcal{X} . Hence, $|\mathcal{W}| = 2^{|\mathcal{X}|} = 2^n$. We denote a specific world as $\omega \in \mathcal{W}$. We say that ω

¹ Some sources, particularly in the belief revision literature, use the term *belief base*.

² Note that we modify our notation from Dean and Parikh’s. In particular, they represent this preference function using two quantities — a weight in $[0, 1]$ and a truth-value preference in $\{-1, 0, 1\}$ — which we combine into the single function p_v .

models a theory T , $\omega \models T$, if ω is consistent with the logical closure of T . For the purpose of defining voter utilities, we treat ω as a function $\omega : \mathcal{X} \rightarrow \{-1, 1\}$ where $\omega(x_i) = 1$ if x_i is true in that world and $\omega(x_i) = -1$ if x_i is false in that world. The voter's utility for some $\omega \in \mathcal{W}$ is

$$u_v(\omega) = \sum_{x_i \in \mathcal{X}} p_v(x_i) \cdot \omega(x_i).$$

The voter's utility for a candidate is a function of the possible worlds modeled by the current theory T of what the candidate has said so far. Dean and Parikh consider three classes of voters:

- *Optimistic* voters evaluate the candidate on the best world modeled by the theory, $ut_v(T) = \max\{u_v(\omega) : \omega \models T\}$.
- *Pessimistic* voters evaluate the candidate on the worst world modeled by the theory, $ut_v(T) = \min\{u_v(\omega) : \omega \models T\}$.
- *Expected-value* voters take the average³ utility over modeled worlds,⁴

$$ut_v(T) = \frac{\sum_{\omega \models T} u_v(\omega)}{|\{\omega : \omega \models T\}|}.$$

Dean and Parikh consider the case of a single candidate who wants to choose statements that maximize the total utility of a population of voters. They prove that with expected-value voters, an optimal strategy involves announcing a stance on every issue, creating a theory that models only one world. (We will refer to such a theory as a *complete* theory; when a theory T' is complete and $T \subseteq T'$, we call T' a *completion* of T .) Furthermore, they observe that with pessimistic voters, it is also advantageous for the candidate to announce a stance on every issue, since eliminating possible worlds can never result in a loss of utility. Only with optimistic voters is it advantageous to remain silent, since eliminating possible worlds can never result in a gain of utility.

2.2 Computational Complexity Classes

We assume familiarity with P and NP . In addition, we will invoke some well-studied but less common complexity notions, which we describe here.

Let C and D be computational complexity classes defined via resource bounds on Turing machines. We denote by C^D those languages or functions computable by a C Turing machine with an *oracle* for D . In other words, we modify a C Turing machine to have an additional tape and state q . If d is a language or function computable by a D Turing machine, then we allow computations of the modified Turing machine to write a string x on the new tape, enter state q , and

³ Dean and Parikh assume that for the purpose of determining “expected value” of the worlds, all worlds are considered equally likely.

⁴ One consequence of this definition is that the utility of an empty or otherwise tautological theory is 0 for any expected-value voter.

in the next step, the new tape contains $d(x)$. We write $A \leq_T^P B$ if $A \in P^B$, meaning “ A is polynomial-time Turing reducible to B ”.

The class NPO is an analogue to NP for optimization problems, i.e., problems that are specified in terms of a definition of valid *instances*, a definition of valid *solutions* with respect to an instance, and a *value function* over the solutions, and ask for a solution with maximum or minimum value. Such a problem is in NPO if and only if it meets the following criteria defined by Ausiello et al. [4]: Instances can be verified as valid in polynomial time, solutions can be verified as valid in time polynomial in the instance size, and the value of a solution can be computed in polynomial time. Note that $NPO \subseteq P^{NP}$.

The function class $\#P$, introduced by Valiant [17], contains those problems that are equivalent to determining the number of accepting paths in an NP Turing machine. If a problem is in NP , then the problem of counting how many witness strings satisfy the NP machine for a given instance is in $\#P$. Since a nonzero answer for a $\#P$ problem instance entails a positive answer for the corresponding NP problem instance and a zero answer for the $\#P$ problem entails a negative answer for the NP problem, $NP \leq_T^P \#P$.

We have $P \subseteq NP \subseteq P^{NP} \subseteq P^{\#P} \subseteq PSPACE \subseteq EXPTIME$.

3 Complexity of Finding Voter Utility

One important factor in the epistemology of campaigns that Dean and Parikh’s framework (with its implicit assumption of “logical omniscience”) does not explicitly model is cognitive complexity (for which we use *computational* complexity as a proxy) as it pertains to the voters. We argue that even a superficially simple series of campaign statements may induce a complex underlying theory: Though we, and Dean and Parikh, have introduced the set of variables \mathcal{X} in terms of high-level issues (x_1 = “Every citizen is entitled to a free pony”) for explanatory purposes, in reality such issues might more accurately be viewed as complex interplays of finer-granularity subissues (“Every citizen is entitled to a free pony” = $(x'_1 \vee x'_2 \vee x'_3) \wedge (x'_1 \rightarrow x'_4 \vee \neg x'_5) \wedge \dots$, where x'_1, x'_2, x'_3 are potential taxes to fund the pony giveaway, x'_4, x'_5 are about the logistics of pony distribution, and so on).

Furthermore, as Dean and Parikh note, additional information can arise from a statement through *implicature* — that which is suggested by a speaker without directly being part of or entailed by “what is said”. For instance, when Vermin Supreme says, “When I’m president everyone gets a free pony”, we discount the possibility that he plans to give everyone *two* free ponies; if he did, that would not *contradict* his promise, but his omission of information would be *infelicitous*.⁵

⁵ In Grice’s account of implicature [8], participants in a conversation assume each other to be obeying certain maxims of cooperativity (for instance, illustrated here is the maxim of Quantity — roughly, *give as much information as necessary, and do not give more information than necessary*); they interpret each other’s statements in light of this mutual assumption.

A potential source of discrepancies between the framework’s predictions and the reality of campaigns is that, given an elaborate body of information about a candidate’s policies, voters have trouble evaluating the candidate due to the intractability of their utility functions. We will consider the computational complexity of the function problems of determining exact voter utility, but also of decision problems of determining whether the utility meets a given threshold, which are particularly relevant for the “stay-at-home voter” scenario we will discuss in Section 5.

3.1 Optimistic Voter Evaluation

Theorem 1. *Given a theory T , an optimistic voter v , and a value k , the problem of deciding whether $ut_v(T) \geq k$ is NP-complete.*

Proof. For NP membership, observe that given a world modeled by T for which v ’s utility is at least k , we can verify the consistency and utility in polynomial time.

We will show NP-hardness with a polynomial-time reduction from Boolean satisfiability (SAT). Let ϕ be a propositional formula over a set of variables \mathcal{X} . We construct the theory as $T = \{x_* \rightarrow \phi\}$, where x_* is a new variable. We construct an optimistic voter v with preferences set as $p_v(x_*) = 1$ and $p_v(x_i) = 0$ for all $x_i \in \mathcal{X}$. And we let $k = 1$. Let $A = \{\omega : \omega \models T\}$. If ϕ is unsatisfiable, then $A = \{\omega : \omega(x_*) = -1\}$, hence $ut_v(T) = -1$. However, if ϕ is satisfiable then there are some $\omega \in A$ where $\omega(x_*) = 1$ and hence $ut_v(T) = 1$. Finally, we have $ut_v(T) \geq 1 = k$ if and only if ϕ is satisfiable.

Theorem 2. *Given a theory T and an optimistic voter v , the problem of computing $ut_v(T)$ and a corresponding best world modeled by T is NPO-complete.*

Proof. The problem satisfies the criteria for NPO-membership [4]: Instances (i.e., the theory and voter specification) and solutions (i.e., worlds modeled by the theory) are recognizable as such in time polynomial in the instance size, and the value function (i.e., voter utility) is computable in polynomial time.

We will show NPO-hardness with a polynomial-time reduction from the maximum weighted satisfiability problem (MAX-WSAT),⁶ for which Ausiello et al. [4] prove NPO-completeness. A MAX-WSAT instance consists of a propositional formula ϕ and a positive weight r_i for each variable x_i ; the problem is to find a satisfying assignment that maximizes the total weight of the variables assigned

⁶ The name “weighted satisfiability” (WSAT) has been used by different sources to refer to two different groups of problems — one where an instance consists only of a propositional formula and the value of a solution is the number of true variables (the Hamming weight), and the generalization we use here where the instance includes weights for the variables. The maximization/minimization versions of the former are sometimes called “maximum number of ones” (MAX-ONES) / “minimum number of ones” (MIN-ONES), and are complete for NPO-PB [12], a subclass of NPO where the magnitude of a solution’s value is polynomially bounded by the size of the input.

to be true. Let $R = \max\{r_i : 1 \leq i \leq n\}$. We construct the theory as $T = \{\phi\}$ and the voter preferences as $p_v(x_i) = r_i/R$ for each x_i .⁷ Then the voter's best world ω is the optimal assignment for the MAX-WSAT instance, and given $u_v(\omega) \in [-\sum_i r_i/R, \sum_i r_i/R]$ we can retrieve the corresponding total weight for the MAX-WSAT assignment by mapping this range onto $[0, \sum_i r_i]$.

3.2 Pessimistic Voter Evaluation

Theorem 3. *Given a theory T , a pessimistic voter v , and a value k , the problem of deciding whether $ut_v(T) \geq k$ is coNP-complete.*

Proof. For coNP membership, observe that given a world modeled by T for which v 's utility is less than k , we can verify the consistency and utility in polynomial time.

We can show coNP-hardness by polynomial-time reduction to this problem from Boolean unsatisfiability (UNSAT). Given a formula ϕ , we construct the theory T and the voter v 's preferences in the same way as in the proof of Theorem 1, except that v prefers the new variable x_* to be false, $p_v(x_*) = -1$. Then $ut_v(T) \geq 1 = k$ if and only if ϕ is unsatisfiable.

Theorem 4. *Given a theory T and a pessimistic voter v , the problem of computing $ut_v(T)$ and a corresponding worst world modeled by T is NPO-complete.*

Proof. NPO membership applies by the same argument as in the proof of Theorem 2. We can show NPO-hardness by polynomial-time reduction from the minimum weighted satisfiability problem (MIN-WSAT), the minimization counterpart to MAX-WSAT; the mapping is constructed in the same manner as in the proof of Theorem 2.

3.3 Expected-Value Voter Evaluation

Lemma 1. *Given a theory T and an expected-value voter v , the problem of computing $ut_v(T)$ is \leq_T^P -hard for #P.*

Proof. Let ϕ be a Boolean formula over $\{x_1, \dots, x_n\}$ and let $S = \#\text{SAT}(\phi)$ be the number of satisfying assignments of ϕ . We define a formula ψ' over $\{x_1, \dots, x_n\} \cup \{y, z\}$ as follows. First we define $\psi = (\phi \wedge y \wedge z)$, and set

$$\psi' = \psi \vee (y \wedge \neg z \wedge \bigwedge_{i=1}^n x_i).$$

Define a voter D with preferences $p_D(y) = p_D(z) = 1$ and $p_D(x_i) = 0$ for each $x_i \in \{x_1, \dots, x_n\}$. (So this voter is defined over $n + 2$ variables.)

Let $A = \{-1, 1\}^n$ and $B = \{-1, 1\}^{n+2}$ be the possible worlds for ϕ and ψ , respectively. Then we have

$$S = |\{a \in A : a \models \phi\}| = |\{b \in B : b \models \psi\}| = |\{b \in B : b \models \psi'\}| - 1.$$

⁷ We divide through by the maximum weight so that the $p_v(x_i)$'s are in $[0, 1]$.

That is, ϕ has as many satisfying assignments over n variables as ψ has over $n+2$ variables, and ψ' has one more satisfying assignment than ψ . Then we get the critical equalities:

$$\begin{aligned} \frac{ut_D(\{\psi\})}{ut_D(\{\psi'\})} &= \frac{(\sum_{b\models\psi} u_D(b))/|\{b \in B : b \models \psi\}|}{(\sum_{b\models\psi'} u_D(b))/|\{b \in B : b \models \psi'\}|} \\ &= \frac{(\sum_{b\models\psi} u_D(b))/S}{(\sum_{b\models\psi'} u_D(b))/(S+1)} \\ &= \frac{(\sum_{b\models\psi} u_D(b))/S}{(0 + \sum_{b\models\psi} u_D(b))/(S+1)} = \frac{S+1}{S}. \end{aligned}$$

(The third equality is valid because the only world in $\{b \in B : b \models \psi'\} \setminus \{b \in B : b \models \psi\}$ has utility 0 for voter D .) This allows us to derive an equation to get $\#\text{SAT}$ from ut_D :

$$\frac{1}{ut_D(\psi)/ut_D(\psi') - 1} = \frac{1}{(S+1)/S - S/S} = \frac{1}{1/S} = S = \#\text{SAT}(\phi).$$

Thus, we can use two calls to an oracle for expected-value utility to compute $\#\text{SAT}(\phi)$ in polynomial time.

Lemma 2. *Given a theory T and an expected-value voter v , the problem of computing $ut_v(T)$ is in $P^{\#P}$.*

Proof. Let T be some theory and let v be some expected-value voter. We assume p_v is represented as a vector of rational binary numbers and we set b as the number of bits in the ‘longest’ number. We describe an NP Turing machine M whose number of witness strings is proportional to $u_v(T)$.

The machine M takes in a Boolean formula ϕ and a voter’s preference function p_v . Then M guesses an assignment ω and a binary integer k where $1 \leq k \leq 2^b$. If both $\omega \models \phi$ and $k \leq u_v(\omega) \cdot 2^b$, then the machine accepts. Otherwise, the machine rejects. Note that both checks take polynomial time.

Given ϕ and b and p_v , how many ways can M accept? If ω does not satisfy, then M cannot accept. If ω does satisfy, then M accepts in exactly $u_v(\omega) \cdot 2^b$ different ways. Hence, $\#M(\phi, p_v, b) = 2^b \sum_{w\models\phi} u_v(\omega)$. Finally, we can compute the utility:

$$ut_v(T) = \frac{\#M(\bigwedge T, p_v, b)}{\#\text{SAT}(\bigwedge T) \cdot 2^b}.$$

Theorem 5. *Given a theory T and an expected-value voter v , the problem of computing $ut_v(T)$ is \leq_T^P -complete for $\#P$.*

Proof. This follows from Lemmas 1 and 2.

4 Complexity of Making an Optimal Theory

Dean and Parikh observe that when all voters are optimistic, a candidate looking to increase total voter utility is best off simply saying nothing; as such, this situation does not raise any nontrivial computational issues from the candidate’s perspective. On the other hand, when appealing to an expected-value or pessimistic voter population, the candidate is best off taking an explicit stance on every issue. The candidate’s ability to do so, of course, depends on their knowing *which* stance to take. This has two aspects: Firstly, the candidate must know the voters’ stances on each issue; given the increasing availability of tools like mass surveys and data analytics that let politicians gauge the attitudes of their constituents, this is a reasonable assumption (though models of candidate uncertainty about voter stances are of interest for future study). Secondly, the candidate must be able to compute the best announcements for appealing to the overall voter population, given the individual voter preferences; this is the family of problems we examine here. When we say “optimal theory” in the following results, we mean a theory T that maximizes $\sum_{v \in V} ut_v(T)$ for the voter population V .

In general, we assume that the candidate starts with an empty theory, and that the candidate is willing to craft whatever platform is most advantageous (being what Dean and Parikh call a “Machiavellian” candidate) rather than being committed to personal beliefs. However, in Section 4.3, we show that having to remain consistent with an existing theory raises the complexity of some relevant problems.

4.1 Appealing to Expected-Value Voters

Theorem 6. *Given n variables and a set $V = \{v_1, \dots, v_m\}$ of expected-value voters, a candidate can construct an optimal theory in time $O(n \cdot m)$.*

Proof. In particular, we will describe a procedure for finding a *complete* optimal theory (since Dean and Parikh have established that with expected-value voters there always exists a complete theory that is optimal).

We can reformulate the set V of m voters with possibly many different preference functions into a new set V' of m voters with all the same preference function such that the candidate receives the same total utility, i.e., $\sum_{v \in V} ut_v(T) = \sum_{v' \in V'} ut_{v'}(T)$ for any theory T . We define

$$p_{v'}(x_i) = \sum_{v \in V} \frac{p_v(x_i)}{|V|}$$

for each variable x_i , for each $v' \in V'$. This reformulation takes time $O(m)$ to compute for each of the n variables. The candidate can then construct a theory that models only the world with the preferred assignment to each variable according to the new preference function.

4.2 Appealing to Pessimistic Voters

Theorem 7. *Given n variables and a set $V = \{v_1, \dots, v_m\}$ of pessimistic voters, a candidate can construct an optimal theory in time $O(n \cdot m)$.*

Proof. The procedure from the proof of Theorem 6 can also be used to construct a complete optimal theory for pessimistic voters; while the equality between the total utilities of the original and reformulated sets of voters no longer holds in general, it still holds for theories like the constructed one that model only a single world (since the utility for this world is both the pessimistic and expected value for the theory).

4.3 Extending an Existing Theory

Until now, we have assumed that a candidate starts with a “blank slate”, able and willing to shape the voters’ beliefs with no restrictions. However, there are many reasons why the candidate may instead need to stay consistent with particular set of formulas — the candidate may be an experienced politician who has revealed platform information in prior elections and incumbencies, may be a member of a political party with established doctrine, or may be “tactically honest” [13] — willing to make strategic statements only insofar as they do not contradict certain deeply-held opinions. The strategy of choosing the most informative theory possible to appeal to expected-value or pessimistic voters becomes harder when the candidate must also remain consistent with an existing theory:

Theorem 8. *Given an expected-value or pessimistic voter v and current theory T , the problem of computing an optimal completion of T is NPO-complete.*

Proof. Observe that this problem is equivalent to the optimistic voter utility problem from Theorem 2 (except that instead of yielding a best world ω modeled by a theory, we are yielding a *theory* that models only ω , which can be accomplished by inserting into T a conjunction of literals with their assignments in ω); thus, the proof of Theorem 2 applies here as well.

5 Complexity of Motivating Enough Voters to Vote

While having an enthusiastic constituency is no doubt correlated with a candidate’s success, the more direct measure is whether enough supporters actually turn up to vote. A 2006 Pew Research Center study addressed the question of when people vote — since so many people do not, at least in the US. Their findings included the following, which addresses the questions of showing up, rather than the decision about how to vote:

The Pew analysis identifies basic attitudes and lifestyles that keep these intermittent voters less engaged in politics and the political process. Political knowledge is key: Six-in-ten intermittent voters say they

sometimes don't know enough about candidates to vote compared with 44% of regular voters the single most important attitudinal difference between intermittent and regular voters identified in the survey. [...] One other key difference: Regular voters are more likely than intermittent voters to say they have been contacted by a candidate or political group encouraging them to vote, underscoring the value of get-out-to-vote campaigns and other forms of party outreach for encouraging political participation.

Pew Research Center [16]

Refusal to vote does not necessarily indicate irrationality on the voter's part; under decision-theoretic models of *expressive voting* [3, 9], where a voter's foremost goal is expressing their views rather than bringing about an outcome, abstinence from voting is a rational choice under certain circumstances. Parikh and Taşdemir [13] suggest the presence of "stay-at-home voters", whose utility for a candidate must meet a certain threshold before they will vote, to explain why a candidate might remain silent in situations where Dean and Parikh's model would otherwise suggest a strategy of explicitness.

5.1 Appealing to Optimistic Voters

Theorem 9. *Given an integer h and a set of optimistic voters $V = \{v_1, \dots, v_m\}$ with thresholds $\{k_1, \dots, k_m\}$, the problem of deciding the existence of a theory T such that $ut_{v_i}(T) \geq k_i$ for at least h voters is in P .*

Proof. Since the empty theory $T = \emptyset$ has the maximum utility for any optimistic voter, it suffices to compute the utility of each voter v_i 's best world, $\sum_{x_j \in \mathcal{X}} |p_{v_i}(x_j)|$, and check whether at least h of these utilities meet their respective voters' thresholds.

5.2 Appealing to Pessimistic Voters

Theorem 10. *Given an integer h and a set of pessimistic voters $V = \{v_1, \dots, v_m\}$ with thresholds $\{k_1, \dots, k_m\}$, the problem deciding the existence of a theory T such that $ut_{v_i}(T) \geq k_i$ for at least h voters is NP-complete.*

Proof. For NP membership: If there exists a theory for which at least h pessimistic voters meet their thresholds, then these voters also meet their thresholds in any completion of this theory (since eliminating worlds never decreases pessimistic voter utility); given the world modeled by one of these completions, we can verify in polynomial time that the thresholds are met.

We will show NP -hardness with a reduction from conjunctive normal form Boolean satisfiability (CNF-SAT). Let ϕ be a Boolean formula in conjunctive normal form. For each clause containing r literals, we construct a pessimistic voter v_i with preferences as follows: For each variable x_j , $p_{v_i}(x_j) = \frac{1}{r}$ if x_j

appears in the clause with positive polarity, $p_{v_i}(x_j) = -\frac{1}{r}$ if x_j appears in the clause with negative polarity ($\neg x_j$), and $p_{v_i}(x_j) = 0$ if x_j does not appear in the clause. We set v_i 's threshold as $k_i = -\frac{r-1}{r}$ so that $ut_{v_i}(T) \geq k_i$ if and only if all worlds modeled by T have at least one variable assigned to match its polarity in the clause. Finally, we set h equal to the number of clauses to require that all worlds modeled by T have at least one variable assigned to match its polarity for *every* clause; such a T exists if and only if ϕ is satisfiable.

5.3 Appealing to Mixed Voters

Theorem 11. *Given an integer h , a set of voters $V = \{v_1, \dots, v_m\} = V^o \cup V^p$ where V^o consists of optimistic voters and V^p consists of pessimistic voters, and thresholds $\{k_1, \dots, k_m\}$, the problem of deciding the existence of a theory T such that $ut_{v_i}(T) \geq k_i$ for at least h voters is NP-complete.*

Proof. NP-hardness follows from the fact that this is a generalization of the problem with only pessimistic voters from Theorem 10.

For NP-membership: Let $V' \subseteq V$ be a set of h or more voters. We could guess an acceptable world (not necessarily distinct) ω_i for each optimistic voter $v_i \in V' \cap V^o$ (i.e., $u_{v_i}(\omega_i) > k_i$), such that ω_i is also acceptable for each pessimistic voter $v_j \in V' \cap V^p$ (i.e., $u_{v_j}(\omega_i) > k_j$). Then the disjunction $T = \{\vee_i w_i\}$ would satisfy all the voters in V' , and each voter can verify this in polynomial time, since there will only be as many modeled worlds as there are optimistic voters. Furthermore, if there exist theories that satisfy h or more voters, then there exists at least one in the aforementioned form.

6 Conclusions

There are many bodies of research that are, or might be, relevant to whether voters show up to vote, and once there, how they vote. We have explored the computational side of one such theory, and showed that it proposes computationally intensive methods for voter evaluation of platforms and for multi-voter satisficing. Given the basic premise that voter satisfaction or satisficing is a combinatorial problem, the intractability is not surprising.

There are many ways this investigation of the computational complexity of modeling voters' behavior can be extended. They include:

- adding value to informativeness of the candidate(s)' platform when predicting whether a voter will show up to vote, as per the Pew study [16];
- decreasing the complexity of candidates' platforms (conjunctions or disjunctions of atomic propositions; Horn formulas; ...);
- modeling change over time in voter priorities [10] or opinions;
- adding affective variables to the voter models [14];
- investigating social-network models of voter interaction and influence [6];
- using game-theoretic models of candidate-candidate interactions and voter choices [11];

- including group-based identity in the decision to show up as well as the choice of candidate.

In addition, we could start from axiomatic characterizations of models of candidate platforms and voter choice. What are good properties to model? (For instance, if a candidate adds to the specificity of their platform in ways that agree with a voter’s preferences, should that increase the likelihood that a voter chooses that candidate, or the likelihood that the voter shows up to vote?) Will we run into Arrow-style impossibility results for achieving all the desiderata we propose?

Acknowledgments The authors thank three anonymous reviewers for their helpful feedback and thank Alec Gilbert for catching errors in a late draft. All remaining errors are the responsibility of the authors. This material is based upon work partially supported by the National Science Foundation under Grants No. IIS-1646887 and No. IIS-1649152. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

I will promise your electorate heart anything you desire, because you are my constituents, you are the informed voting public, and I have no intention of keeping any promise that I make.

Vermin Supreme⁸

⁸ From Revolution PAC’s 2012 interview (<https://www.youtube.com/watch?v=9jKszduiK8E>)

Bibliography

- [1] Enriqueta Aragones and Zvika Neeman. Strategic ambiguity in electoral competition. *Journal of Theoretical Politics*, 12(2):183–204, 2000. doi: 10.1177/0951692800012002003.
- [2] Enriqueta Aragones and Andrew Postlewaite. Ambiguity in election games. *Review of Economic Design*, 7(3):233–255, 2002. doi: 10.1007/s100580200081.
- [3] Enriqueta Aragones, Itzhak Gilboa, and Andrew Weiss. Making statements and approval voting. *Theory and Decision*, 71(4):461–472, 2011. doi: 10.1007/s11238-010-9208-5.
- [4] Giorgio Ausiello, Pierluigi Crescenzi, Giorgio Gambosi, Viggo Kann, Alberto Marchetti-Spaccamela, and Marco Protasi. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer, 1999. doi: 10.1007/978-3-642-58412-1.
- [5] Vardan Baghdasaryan and Elena Manzoni. Set them (almost) free. an analysis of discretion in electoral campaigns. Technical Report 354, University of Milano-Bicocca, DEMS, December 2016.
- [6] Robert M. Bond, Christopher J. Fariss, Jason J. Jones, Adam DI Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298, 2012. doi: 10.1038/nature11421.
- [7] Walter Dean and Rohit Parikh. The logic of campaigning. In *Indian Conference on Logic and its Applications*, pages 38–49. Springer, 2011. doi: 10.1007/978-3-642-18026-2_5.
- [8] H. Paul Grice. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and Semantics 3: Speech Acts*, pages 41–58. New York: Academic Press, 1975.
- [9] Dominik Klein. Expressive voting: Modeling a voter’s decision to vote. In *ESLLI Workshop on Logical Models of Group Decision Making*, 2013.
- [10] Dominik Klein and Eric Pacuit. Focusing on campaigns. In Can Başkent, Lawrence S. Moss, and Ramaswamy Ramanujam, editors, *Rohit Parikh on Logic, Language and Society*, pages 77–89. Springer, 2017. doi: 10.1007/978-3-319-47843-2_5.
- [11] Joachim I. Krueger and Melissa Acevedo. A game-theoretic view of voting. *Journal of Social Issues*, 64(3):467–485, 2008. doi: 10.1111/j.1540-4560.2008.00573.x.
- [12] Alessandro Panconesi and Desh Ranjan. Quantifiers and approximation. *Theoretical Computer Science*, 107(1):145–163, 1993. doi: 10.1016/0304-3975(93)90259-V.
- [13] Rohit Parikh and Çağil Taşdemir. The strategy of campaigning. In *Fields of Logic and Computation II*, pages 253–260. Springer, 2015. doi: 10.1007/978-3-319-23534-9_15.

- [14] Michael T. Parker and Linda M. Isbell. How I vote depends on how I feel: The differential impact of anger and fear on political information processing. *Psychological Science*, 2010. doi: 10.1177/0956797610364006.
- [15] John R. Petrocik. Issue ownership in presidential elections, with a 1980 case study. *American Journal of Political Science*, 40(3):825–850, 1996. doi: 10.2307/2111797.
- [16] Pew Research Center. Who votes, who doesn't, and why: Regular voters, intermittent voters, and those who don't, 2006.
- [17] Leslie G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979. doi: 10.1016/0304-3975(79)90044-6.