A Local-Dominance Theory of Voting Equilibria

Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein

Abstract

We suggest a new model for strategic voting that takes into account voters' bounded rationality, as well as their limited access to reliable information. We introduce a simple behavioral heuristic based on *local dominance*, where voters each consider a set of possible states without assigning probabilities to them. In a voting equilibrium, all voters vote for candidates that are not dominated within the set of possible states. We prove that local dominance-based dynamics, where a voter moves to the best candidate that locally dominates her current selection, quickly converge to an equilibrium if voters start from the truthful state.

Using extensive simulations of strategic voting on generated preference profiles (and some real ones), we show that convergence is quick and robust, that emerging equilibria replicate widely known patterns of human voting behavior such as Duverger's law, and that strategic voting under the Plurality rule generally improves the quality of the winner compared to truthful voting.

1 Introduction

Game-theoretic considerations have been applied to the study and design of voting systems for centuries, but the question of how people vote, or should vote, is still open. Suppose that we put aside the complications involved in political voting,[1] and focus on a simple scenario that fits all the "standard" assumptions: each of n voters has a complete transitive preference $\prec_i$ over a fixed set of alternatives M, and each voter's only purpose is to bring about the election of her most-favorable alternative. We will further restrict ourselves to discussing the common Plurality rule, where the alternative with the maximal number of votes is the winner. This scenario translates naturally to a *game*, in which the actions of each voter are her possible ballots—the set of alternatives, in the case of Plurality. One might expect game theory to give us a definitive answer as to what would be the outcome of such a game.

However, an attempt to apply the most fundamental solution concept, Nash equilibrium, to the scenario above, reveals a disappointing fact: almost *any* profile of actions is a pure Nash equilibrium, and even a highly unpopular alternative may win in some equilibrium. Yet, people do very often vote strategically.

Example 1. As a running example, we consider a profile with 3 candidates $M = \{a, b, c\}$. Suppose that there are 100 voters, and that according to the last poll, votes are divided as 45 for a, 40 for b, and 15 for c. Among the supporters of c are voters v and v'. Voter v has preference $c \succ b \succ a$, whereas voter v' prefers $c \succ a \succ b$.

While if truthful, v would stay with c, it seems that c has no chance of winning, and thus a wise strategic decision for v would be to *change* her vote to b.

By applying the reasoning above to all voters, we would expect to eventually reach an equilibrium where only a and b get votes. The phenomenon that under the Plurality rule almost all votes divide between two candidates is well known in political science, and is called Duverger's Law [12].

[1] For example, social utilities [22, 4], strategic candidates [5, 15], and other considerations as in [32, 13].

This observation triggered a search for more appropriate solution concepts for voting games. These concepts rely on taking into account various additional factors, such as the information available to the voters, collusion and group behavior, and intrinsic preferences towards certain actions. All of these models had to explicitly or implicitly make assumptions on what it means to vote *strategically.* Fisher [17] offers the following definition for strategic, or *tactical* voting:

> *A tactical voter is someone who votes for a party they believe is more likely to win than their preferred party, to best influence who wins in the constituency.*

The goal of this work is to lay sound foundations for voting models that better explain and predict the behavior of human voters. In particular, we wish to extend the above intuition to general voting scenarios, by providing a new interpretation for the two key components of Fisher's definition: *belief* (what voters know) and *influence* (what voters consider as an effective action).

Our contribution We begin by presenting a brief desiderata that in our opinion should guide the search for a proper solution concept, and review some solutions that have been proposed in the literature. We then lay out our framework, which is a non-probabilistic way of capturing uncertainty. Using this framework and simple behavioral assumptions, we present our response function and equilibrium concept. In the remainder of the paper, we will argue, using formal propositions and empirical analysis, that our solution is indeed the appropriate one for Plurality voting. In particular we show that voters who start from the truthful vote will quickly converge to a pure equilibrium, and that convergence is likely to occur even from arbitrary initial states. We show that under various distributions over preference profiles, using the local dominance framework enables "Duverger-like" outcomes. Also, the attained equilibria often yield "better" winners than the truthful outcome. For example, a Condorcet winner, if one exists, is more likely to be elected in equilibrium.

An extended version of the paper was published in EC 2014 [24]. In this manuscript we present our main results and focus on discussion and implications, while omitting some of the proofs and other technical details. The full version is available for download at `https://arxiv.org/abs/1404.4688`.

2 Desiderata for Voting Models

We now present some arguably-desirable criteria for a theory of voting. We will not be too picky about what is considered a voting model, and whether it is described in terms of individual or collective behavior. The key feature of a model is that given a profile of preferences, it can be mapped to a set of outcomes (i.e., of possible or likely voting profiles under the Plurality rule). We classify desirable criteria into the following classes: theoretic (mainly game-theoretic), behavioral, and scientific.

Theoretic Criteria

Rationality Voters are trying to maximize their own utility based on what they know and/or believe.

Equilibrium Predicted outcomes are in equilibrium (for example, a refinement of Nash equilibrium, or of another popular solution concept from the game theory literature).

Discriminative power (sometimes called predictive power) The predicted outcomes are a small but non-empty set of possible outcomes. More specifically, a small set of possible winners will be predicted, as there may be many voting profiles with the same winner.

Broad scope The model applies (or can be easily adapted to) various scenarios such as simultaneous, sequential or iterative voting, and the use of different voting rules.

In addition, we put forward another less formal requirement: we would like our model to be grounded in familiar concepts from decision theory, game theory, and voting theory. It will thus be easier to understand, and to compare with other models.

Behavioral criteria

By behavioral criteria, we refer to the assumptions the model makes about voters' behavior.

Voters' knowledge Voters' behavior should not be based on information that they are unlikely to have, or that is hard to obtain.

Voters' capabilities The decision of the voter should not rely on complex computations, non-trivial probabilistic reasoning, etc.

In addition, we would like the behavioral assumptions, whether implicit or explicit, to be supported by (or at least not to directly contradict) studies in human decision making.

Scientific criteria

Robustness The model gives similar predictions even if some voters do not exactly follow their prescribed behavior, if we slightly modify the available information, if we change the tie-breaking scheme, etc. Except in a few threshold cases, we would not expect small perturbations to change the identity of the winner.

Few parameters The model should have as few as possible parameters, and we would like each of them to be meaningful (e.g., voters' memory).

Reproduction Simulating generated or real preferences using the model, we would like to see reproduction of common phenomena (e.g., Duverger's law).

Experimental validation The hardest test for a model is to try and predict the behavior of human voters based on their real preferences. By comparing the predicted and real votes (or even just outcomes), we can measure the accuracy of the model.

The behavioral criteria together with the rationality requirement can be thought of as criteria of *bounded rationality*. Lastly, some voting models explain how strategic behavior is *better for society*. For example, equilibrium outcomes in Plurality with a particular voter behavior may have a better Borda score, or coincide more with choosing a Condorcet winner. Although this is not exactly a criterion for a good model (real strategic behavior may not increase welfare), we are interested in the conditions under which the theory predicts increased welfare, as these may be useful for design purposes.

3 Related Work

A broad literature review on various equilibria concepts for voting can be found in the full version of the paper, where we also evaluate some of the models with respect to the desiderata above. We mention here only some models that are closely related to our work.

The Leader Rule Before we start, we would like to highlight a model that fares nicely in almost all of the above criteria, which is Laslier's *leader rule* for Approval voting [21]. This is a simple parameter-free behavioral strategy, where a voter only needs to take into account a prospective ranking of the candidates (which can be available from a poll, a prior belief, or a previous voting round).[2] The leader rule is behaviorally plausible, has attractive theoretical properties, makes minimal informational assumptions, and seems to explain human voting behavior at least in some contexts. Unfortunately, it does not seem to have a natural extension to Plurality voting.

Nash equilibrium The basic notion of a pure Nash equilibrium (PNE) in a normal form voting game with complete information is effectively useless, as almost any outcome (even one where all voters vote for their worst candidate) is a PNE. Consider voter v in Example 1. She is powerless to change the outcome, and therefore has no incentive to change her vote.

Several refinements have been suggested in order to mitigate the equilibria explosion problem, some of which rely on plausible behavioral tendencies.

Truth-bias and Lazy-bias A truth-biased voter gains some negligible additional utility from reporting his true preferences (i.e., his top candidate) [25, 11]. Nash equilibria under Plurality with truth-biased voters have been studied empirically by Thompson et al. [34], and analytically by Obraztsova et al. [27]. Truth-bias eliminates many "weird" equilibria, but whenever the truthful profile happens to be a Nash equilibrium (as in our example above), it also survives truth-bias.

Similarly, in some contexts it is plausible to assume that there is some small cost involved in the voting itself.

The conclusion that a "rational" voter would often rather abstain (as she is rarely pivotal) is typically referred to in the literature as the "no-vote paradox" (see, e.g., [10, 28]). When voting is presented as a normal form game, we can add abstention as an additional allowed action. A "lazy" voter would thus choose to abstain if she cannot affect the outcome. Pure Nash equilibria with lazy voters were studied, for example, in [9]. These are typically highly degenerated voting profiles, where all voters except one abstain.

Voting under uncertainty Uncertainty partly solves the problem of equilibria explosion, since any voter can become pivotal with some probability, and therefore cares about whom to vote for.

Myerson and Weber's [26] *theory of voting equilibria* relaxes the assumption that preferences are public information. Instead, they assume that voters' types are sampled i.i.d. from some known underlying distribution. An *equilibrium* in this model is a distribution over votes, such that every voter is maximizing her expected utility w.r.t. this distribution. Myerson and Weber prove that an equilibrium always exists for a broad class of voting rules including Plurality. Focusing on a few examples with three candidates under Plurality, they show that their model gives reasonable results, and may replicate Duverger's law.

[2] According to the leader rule, the voter approves all candidates that are preferred to the prospective leader, and approves the leader if and only if it is preferred to the runner-up.

The fundamental assumption that voters maximize some expected utility function in prevalent in the political science literature; see, e.g., [33, 29, 2]. However, people are notoriously bad at estimating probabilities, and are known to employ various heuristics instead [35].

An additional disadvantage of the expected utility maximization approach (not just of the Myerson and Weber model), is that voters' preferences must be cardinal and cannot be described as a permutation over candidates.

Strict uncertainty (without probabilities) was considered by Ferejohn and Fiorina [16], who took a minmax regret approach. However their model, like probability-based models, heavily relies on voters having cardinal utilities. Also, they take an extreme approach where voters do not use *any* available information, and thus all states are considered possible.

A different approach to strict uncertainty was taken by Conitzer et al. [7] (although without looking at equilibria). Our notion of local dominance is highly related, and in the full version we explain the connections with [7] and with follow-up work [30, 36].

Iterative and sequential games In sequential voting games voters report their preferences one at a time, where every voter can see all of the previous votes (as in Doodle and Facebook polls). The standard solution concept for sequential games is subgame perfect Nash equilibria [14, 23, 8, 9]. However, subgame perfection is a highly sophisticated behavior that requires a voter to know exactly the preferences of all of her peers. It also requires multiple steps of backward induction, at which human players typically fail [19].

In an *iterative setting*, voters start from some given voting profile, and at each turn one or more voters may change their vote [25, 30]. Meir et al. [25] proved that if voters play one at a time and adopt a myopic best-response strategy they are guaranteed to converge to a Nash equilibrium of the stage game from any initial state. The main problem with this approach is that it does not solve equilibria explosion. More recent papers on the iterative setting suggested other myopic strategies [18], which suffer from similar problems.

When all voters are allowed to change their votes at each step, we essentially have repeated polls [6, 30]. A particular model based on uncertainty with iterated polls was suggested by Reyhani et al. [31], where each voter votes for her most preferred candidates among the leaders.

4 The Formal Model

Basic notations We denote $[x] = \{1, 2, \ldots, x\}$. The sets of candidates and voters are denoted by M and N, respectively, where $m = |M|, n = |N|$. The Plurality voting rule f allows voters to submit their preferences over the candidates by selecting an action from the set M. Then, f chooses the candidate with the highest score, breaking ties lexicographically.

Let $\pi(M)$ be the set of all orders over M. We denote a preference profile by $\mathbf{Q} \in (\pi(M))^n$. The preferences of voter i are denoted by the total order $Q_i \in \pi(M)$, where $Q_i(a) \in [m]$ is the rank of candidate $a \in M$, and $q_i = Q_i^{-1}(1)$ is the most-preferred candidate. We denote $a \succ_i b$ if $Q_i(a) < Q_i(b)$. Let $\overline{Q}$ be the lexicographic order over candidates. Each voter announces his vote publicly. Thus the action of a voter is $a_i \in M$. The profile of all voters' votes is denoted as $\mathbf{a} \in M^n$, and the profile of all voters except i is denoted by $\mathbf{a}_{-i}$. When abstention is allowed, we have $a_i \in M \cup \{\perp\}$, where $\perp$ denotes abstaining. If $a_i = q_i$ we say that i is *truthful* in $\mathbf{a}$, and voter i is called a *core supporter* of a_i. Otherwise, we say that i is a *strategic supporter* of a_i.

The scoring vector $\mathbf{s_a} \in \mathbb{N}^m$ that corresponds to $\mathbf{a}$ assigns a score to every candidate, taking tie-breaking into account. Formally, we refer to $s_\mathbf{a}(c)$ as equal to the number $|\{i \in N : a_i = c\}|$. When comparing two scores, we write $s_\mathbf{a}(c) >_{\overline{Q}} s_\mathbf{a}(c)$ if either $|\{i \in N : a_i =$

$c\}| > |\{i \in N : a_i = c'\}|$, or c, c' have the same number of votes and $c \succ_{\overline{Q}} c'$. We usually omit the subscript from $>_{\overline{Q}}$ as it is clear from the context.

We will use $\mathbf{a}$ and $\mathbf{s_a}$ interchangeably where possible, sometimes omitting the subscript $\mathbf{a}$ (note that we may only use $\mathbf{s}$ in a context where voters' identities are not important). Given a state $\mathbf{s}$ and an additional vote a_i, in the concatenated state $\mathbf{s}' = (\mathbf{s}, a_i)$ we have $s'(a_i) = s(a_i) + 1$, and $s'(a) = s(a)$ for all $a \neq a_i$.

4.1 An intuitive description of voter response

While the notation we will introduce momentarily is somewhat elaborate and is intended to enable rigorous analysis, the main idea is very simple and intuitive. We lean on two key concepts that are featured in previous models: *dominated strategies*, and *better-response*. From the perspective of voter i, a candidate a *dominates* candidate b if $f(\mathbf{s}, a) \succ_i f(\mathbf{s}, b)$ for *all* $\mathbf{s}$. In contrast, a is a *better-response* for a voter voting for b *in a particular state* $\mathbf{s}^*$, if $f(\mathbf{s}^*, a) \succ_i f(\mathbf{s}^*, b)$.

In our model, we will relax both concepts in a way that takes into account voters' uncertainty over the actual state. We assume that voters have a common *estimated*, uncertain, view of the current state $\mathbf{s}$. In any given state, a voter considers a set of multiple "close" states which might be realized **without assigning probabilities to them**. We say that a *locally dominates* b in $\mathbf{s}$ if $f(\mathbf{s}', a) \succ_i f(\mathbf{s}', b)$ in all $\mathbf{s}'$ that are "close" to $\mathbf{s}$ (by some reasonable metric). Our key behavioral assumption is that a voter will avoid voting for candidates that are locally dominated. The strategic response of a voter in some state would be voting for her most-preferred candidate—among those who locally dominate her current action.

Tying our model back to Fisher's definition of strategic voting, a voter's *belief* is captured by the estimated state $\mathbf{s}$, whereas her *influence* is reflected by local dominance. The different sets of states that voters consider are part of their type, and can account for diverse voter behavior, yet ones that are bounded-rational. Consider Example 1, where the estimated counts are $\mathbf{s} = (45, 40, 15)$. If voter v also considers as possible states where scores vary by ± 10, he will be better off voting for b. Voting for b might influence the outcome, whereas voting for c is futile (unless v considers an even higher variability in scores).

One assumption that requires justification is the existence of the publicly known, estimated state $\mathbf{s}$. In iterative and sequential settings the shared common view is easy to explain, as the estimated state is the actual current voting state. Uncertainty still exists since some voters might change their vote in later rounds. In a simultaneous-vote game, the estimated state might be due to prior acquaintance with the other voters or due to polls. Uncertainty is due to polls' inaccuracy.

4.2 Local dominance

Let $S \subseteq \mathbb{N}^m$ be a set of states.

Definition 1. We say that action a_i S-beats a'_i (w.r.t. voter i) if there is at least one state $\mathbf{s} \in S$ s.t. $f(\mathbf{s}, a_i) \succ_i f(\mathbf{s}, a'_i)$. That is where i strictly prefers $f(\mathbf{s}, a_i)$ over $f(s, a'_i)$.

We can think of S as states that i *believes* to be possible (where these states do not include the action of i himself).

Definition 2 (Local dominance)**.** We say that action a_i S-dominates a'_i (w.r.t. voter i) if (I) a_i S-beats a'_i; and (II) a'_i does *not* S-beat a_i.

Note that S-dominance is a transitive and antisymmetric relation (but not complete).

Our definition coincides with the definition of dominance in [7] and with similar definitions in [30, 36] (see full version for details). The novelty comes from the way S is constructed, which is explained next.

Distance-based dominance Suppose we have some distance metric for states, denoted by $\delta(\mathbf{s}, \mathbf{s}')$. For voter i and $\mathbf{a} \in M^n$, let $S_i(\mathbf{a}, x) \subseteq \mathbb{N}^m$ be the set of states that are at distance at most x from a_{-i}. Formally, $S_i(\mathbf{a}, x) = \{\mathbf{s}' : \delta(\mathbf{s}', \mathbf{s}_{\mathbf{a}_{-i}}) \leq x\}$.

In this paper we assume that δ is the ℓ_1 norm. Thus $\delta(\mathbf{s}', \mathbf{s}) \leq x$ means that we can attain $\mathbf{s}'$ by adding/removing a total of x voters to $\mathbf{a}_{-i}$ (think of x an an integer, and note that the total number of votes in $\mathbf{s}, \mathbf{s}'$ may be different). For a discussion of other natural metrics, see full version.

4.3 Strategic voting and equilibria

Let $g_i : M^n \to M$ be a *response function*, i.e., a mapping from voting profiles to actions (which implicitly depends on the preferences of voter i). Any set of response functions $(g_i)_{i \in N}$ induces a (deterministic) dynamic in the normal form game corresponding to a particular preference profile under Plurality. In particular, it determines all equilibria of this game, which are simply the states $\mathbf{a}$ where no voter has a response that changes the state. We refer to the response function of a voter as her *type*.

Definition 3. Let N be a set of voters with response functions $(g_i)_{i \in N}$. A *voting equilibrium* is a state $\mathbf{a}$, where $a_i = g_i(\mathbf{a})$ for all $i \in N$.

We next define the primary response function that strategic voters in our model apply, which is based on local dominance.

Definition 4. A *strategic voter of type* r (or, in short, an r voter) acts as follows in state $\mathbf{a}$. Let $D \subseteq M$ be the set of candidates that $S_i(\mathbf{a}, r)$-dominate a_i. If D is non-empty, then i votes for his most preferred candidate in D. Formally, $g_i(\mathbf{a}) = \text{argmin}_{d \in D} Q_i(d)$ if $D \neq \emptyset$, and $g_i(\mathbf{a}) = a_i$ otherwise.

We denote such a strategic step by $a_i \xrightarrow{i} a'_i$, where $a'_i = g_i(\mathbf{a})$. We refer to r (or r_i if types differ) as the *uncertainty parameter* of the voter. We observe that:

- For $r = 0$, the voter knows the current voting profile exactly, and thus her response function is simple best-response, as in [25].
- For $r = n$, a voter does not know anything about the current voting profile. Thus an action a'_i locally dominates a_i if and only if it weakly *globally dominates* a_i. A step $a_i \xrightarrow{i} a'_i$ means that a_i is i's least-preferred candidate.

Different definitions of strategic responses (distance metrics, value of r) may induce different sets of voting equilibria. However, the assumption that i votes for the most-preferred candidate in D is irrelevant to the set of induced equilibria. The following is an immediate observation.

Proposition 1. *Let N be a set of voters with preferences* $\mathbf{Q}$ *and following Def. 4 (voters may be of heterogeneous types). A voting profile* $\mathbf{a}$ *is a voting equilibrium, if and only if* $\forall i \in N, !\exists a'_i \in M, a'_i\ S_i(\mathbf{a}, r_i)$*-dominates* a_i*. That is, no voter votes for a locally dominated candidate.*

We assume of course, that the same parameters are used for defining the dominance relation and the strategic response of each voter. For example, under Definition 4 with $r = 0$, a strategic move coincides with best-response, and voting equilibria coincide with pure Nash equilibria.

5 Convergence with Strategic Voters

For any $w \in \mathbb{N}$, let $H_w(\mathbf{s}) \subseteq M$ be the set of candidates that need exactly w more votes to become the winner. Thus $H_0(\mathbf{s}) = \{f(\mathbf{s})\}$, $H_1(\mathbf{s}) = \{c : s(f(\mathbf{s})) > s(c) \geq s(f(\mathbf{s}))-1\}$ (either have the same score as the winner and lose by tie-breaking, or c wins in the tie-breaking but have one vote less), etc. Let $\overline{H}_w(\mathbf{s}) = \bigcup_{w' \leq w} H_{w'}(\mathbf{s}) = \{c : s(c) \geq s(f(\mathbf{s})) - w\}$.

5.1 Strategic responses and a threshold for possible winners

We first show that in every strategic response, a voter always votes for her favorite possible winner.

Lemma 2. *Consider a strategic response $a_i \overset{i}{\rightarrow} a'_i$ s.t. $a_i \notin \overline{H}_{r+1}(\mathbf{a})$. Then $a'_i = \text{argmin}_{c \in \overline{H}_{r+1}(\mathbf{a})} Q_i(c)$.*

The set $\overline{H}_{r+1}(\mathbf{a})$ represents the set of possible winners, however Lemma 2 does not mean that our dynamics coincides with "always vote for the most-preferred possible winner". It only holds when the current vote of i is *not* a possible winner, and when there are at least two possible winners.

Note that Lemma 2 entails that the strategic decision of the voter is greatly simplified from both a behavioral and a computational perspective. There is no need to consider all possible world states, only to check which candidates are sufficiently close to the winner in terms of their prospective score, and to select the one that is most-preferred.

5.2 Convergence to equilibrium from the truthful state

In what follows, we will only consider the ℓ_1 norm for simplicity. However, most results hold for other metrics as well. We also highlight that as in [25], we allow any number of non-strategic voters to participate.

Best-response graphs and schedulers Given a game, any dynamic induces a directed graph whose vertices are the states of the game (M^n in the case of the Plurality game). There is an edge from a state $\mathbf{a}$ to a state $\mathbf{a}'$, if (1) $\mathbf{a}, \mathbf{a}'$ differ only by the action of one player i; and (2) $g_i(\mathbf{a}) = a'_i$. We call this graph the *best-response* graph.

A *scheduler* selects which voter plays at any step of the game (that is, breaks ties when more than one voter has a strategic move). We assume that the order of players is determined by an arbitrary scheduler (see [3]).

We next show that when starting from the truthful state, any scheduler guarantees convergence to an equilibrium. In particular, an equilibrium must exist. Proposition 3 also follows from more general convergence results that we will show later. However, we provide a simple and detailed proof that reveals the natural structure of the equilibrium that is reached. In the full version, we show how to extend the proof to other distance metrics. The voting dynamics and the arguments are quite similar, but rely on a more flexible notion of a "possible winner".

Theorem 3. *Suppose that all voters are of type r; then a voting equilibrium exists. Moreover, in an iterative setting where voters start from the truthful state, for any scheduler, they will converge to an equilibrium in at most $n(m-1)$ steps.*

Proof. If the truthful state $\mathbf{q}$ is stable, then we are done. Thus assume it is not. Let $\mathbf{a}^t$ (and $\mathbf{s}^t$) be the voting profile after t steps from the initial truthful vote $\mathbf{a}^0 = \mathbf{q}$. Let $a_i \overset{i}{\rightarrow} a'_i$ be a move of voter i at state $\mathbf{s} = \mathbf{s}_{\mathbf{a}^t}$ to state $s' = \mathbf{s}_{\mathbf{a}^{t+1}}$. We claim that the following hold

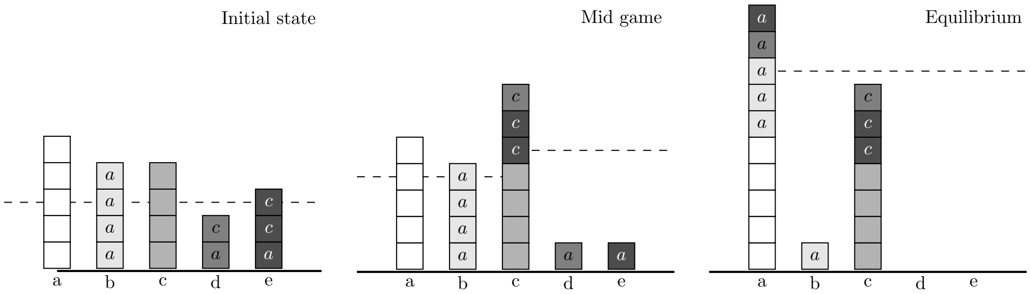

Figure 1: An example of 18 voters voting over candidates $\{a, b, c, d, e\}$. The top left figure shows the initial (truthful) state of the game. The letter inside a voter is his second preference. The dashed line marks the threshold for possible winners (for voters of type $r = 2$). Thus candidates on or above the threshold are the set $\overline{H}_3(\mathbf{a})$. Candidates that are on the dashed line (in $H_3(\mathbf{a})$) are considered possible winners only by voters that do not currently vote for them. Note that due to tie-breaking it is not the same for all candidates. For example, since a beats b in tie-breaking, b needs 2 more votes to win in the initial state.

throughout the game. Recall that by Lemma 2, $\overline{H}_{r+1}$ is the set of possible winners for all voters who are voting for other candidates. We show the following:

1. $a_i \notin \overline{H}_{r+1}(\mathbf{s}')$, i.e., once a candidate is deserted, it is no longer a possible winner.
2. $a'_i \prec_i a_i$, i.e., voters always "compromise" by voting for a less-preferred candidate.
3. $\max_{a\in M} s'(a) \geq \max_{a\in M} s(a)$, i.e., the score of the winner never decreases.
4. $\overline{H}_{r+1}(\mathbf{s}') \subseteq \overline{H}_{r+1}(\mathbf{s})$, i.e., the set of possible winners can only shrink.

We prove this by a complete induction.

1. If this is the first move of i then $a_i = q_i$. Otherwise, by Lemma 2, a_i is the most-preferred candidate of i in $\overline{H}_{r+1}(\mathbf{s}^{t'})$ where t' is the time when i last moved. By induction on (4), $\overline{H}_{r+1}(\mathbf{s}) \subseteq \overline{H}_{r+1}(\mathbf{s}^{t'})$. So if $a_i \in \overline{H}_{r+1}(\mathbf{s})$, it must be the most-preferred candidate in the set. Assume, towards a contradiction, that $a_i \in \overline{H}_{r+1}(\mathbf{s})$; then there is a state $\hat{\mathbf{s}} \in S_i(\mathbf{a}, r)$ where i is pivotal between a_i and $f(\mathbf{a}) \prec_i a_i$. For any $c \neq a_i$, $f(\hat{\mathbf{s}}, c) = f(\mathbf{a})$, and in particular for $c = a'_i$. Therefore a_i $S_i(\mathbf{a}, r)$-beats a'_i, which means a'_i does not $S_i(\mathbf{a}, r)$-dominate a_i. A contradiction.

2. If this is the first move of i then this is immediate. Otherwise, by induction on Lemma 2 and (4), if $a'_i \succ_i a_i$, then i would prefer to vote for a'_i in his previous move, rather than for a_i.

3. As in (1), if i votes for $a_i = f(\mathbf{a}) \in \overline{H}_{r+1}(\mathbf{a})$, then a_i is i's most-preferred possible winner. Thus it cannot be locally dominated by any other candidate.

4. Since by (3) the score of the winner never decreases, the only way to expand $\overline{H}_{r+1}$ is to add a vote to a candidate not in $\overline{H}_{r+1}$. By Lemma 2 this never occurs.

Finally, by property (2), each voter moves at most $m-1$ times before the game converges. □

The proof not only shows that an equilibrium exists, it also describes exactly the way in which such equilibria are reached from the truthful state (see Figure 1). There is always a set of "leaders" ($\overline{H}_{r+1}$ in the case of the ℓ_1 norm). Strategic voters vote for their favorite candidate in this set, if their current candidate is not a possible winner. At some point candidates may "drop out" of the race as their gap from the winner increases, and the set

$\overline{H}_{r+1}$ shrinks. This continues; in the reached equilibrium, all voters are either truthful, or strategically vote for their best possible winner (which is in $\overline{H}_{r+1}$).

5.3 Convergence under broader conditions

Convergence with simultaneous moves and from arbitrary states Unfortunately, if all voters can update their vote at once, convergence is not guaranteed even from the truthful state [25].

However, we can show that under mild restrictions on the scheduler convergence is guaranteed also with simultaneous moves. Moreover, this holds *from any initial state* and not only from the truthful state. In particular, there is a path in the best response graph from any state to an equilibrium. See the full version for details. We conjecture that if voters play one at a time then every path converges to an equilibrium.

Truth-bias and Lazy-bias The notion of local dominance is very flexible, and allows us to define more subtle behaviors. In particular, by adding a negligible utility ϵ to a favorite action, such as truth-telling or abstaining, we get that this action locally dominates any other action where the outcome is the same (that is, the same in all possible world states). We can thus define *truth-biased* or *lazy* voters, who prefer to tell the truth or to abstain whenever they do not see themselves as pivotal.

We highlight that the local neighborhood considered by the voter when deciding whether to cast a strategic vote and when applying truth-bias/laziness, is not necessarily the same neighborhood. Therefore we use an additional parameter k (or k_i if voters' types differ). The voter uses the uncertainty parameter r to decide on her next strategic move (as in Def. 4), and the uncertainty parameter k to decide whether to become truthful/abstain.

For $k = n$ we get that truth-bias or lazy-bias are only applied if a_i is globally dominated, thus the voter almost always votes according to Def. 4. For $r = k = 0$, we get behavior that coincides with [34, 27] (for truth-bias), or with [9] (for lazy-bias).

We argue that it is natural to assume $k > r$, which entails that a voter requires a lower uncertainty level in order to make a new strategic step, rather than to merely keep his current strategic vote. From a behavioral perspective, such an assumption accounts for *default-bias*: decision makers have a higher tendency to stay with their current decision than to adopt a new one [20]. Under this assumption we extend Theorem 3 to societies that include truth-biased and lazy voters (see full version).

6 Simulations of Strategic Voting

We explore via extensive simulations how employing local dominance affects the result of the voting process. These simulations have two primary goals. First, we want to understand better the effect of different parameters on the technical level. More importantly, we use simulations to test the properties of our strategic model with respect to the desiderata listed in Section 2 (for example, what is its discriminative power, is it robust to small changes, and does it replicate common phenomena).

We generate preference profiles from a set of distributions which have been examined in the research literature: the Uniform (or impartial culture) distribution; a uniform Single-peaked distribution; the Polya-Eggenberger Urn model (with 2 urns and with 3 urns); a Riffle distribution; and the Placket-Luce distribution. Urn and Riffle models were particularly designed to resemble preference structures in human societies, whereas in Placket-Luce distributions each voter is assumed to have a noisy signal of some ground truth.

In addition to the generated profiles, we used real datasets of voters' preferences including all 225 currently available full preferences from PrefLib (`http://preflib.org`). See the full version for details of our methods, including the distributions and profiles used.

Methods We generated profiles from all distribution types for various numbers of voters (10, 20, 50) and candidates (between 3 and 8), which resulted in 108 distinct distributions. From each distribution we sampled 200 instances. Then, we simulated strategic voting on each instance, varying the voters' types (basic, truth-biased, lazy) and the uncertainty parameters r and k.

We simulated voting in an iterative setting, where voters start from an initial (truthful) state, and then iteratively make strategic moves until convergence. We repeated each simulation 100 times (as the scheduler may pick a different path each time), and collected multiple statistics on the equilibrium outcomes. All of the collected data can be downloaded from `http://tinyurl.com/k2b775e`.

6.1 Results

We present here the main findings for simulations starting from the truthful state, for the ℓ_1 metric without truth-bias or lazy-bias. The details of our empirical results (including biased agents and random starting states) are available in the full version.

We should note that despite the fact that our convergence proofs do not cover heterogeneous populations, arbitrary simultaneous moves or arbitrary starting points, convergence was just as robust under all of these conditions.

All the patterns we report were robust and clearly visible in all 108 distributions.

Peak r value The most meaningful parameter in the simulations was the uncertainty level r. As we vary the value of r from 0 to 15, there is an increase and then a decrease in the amount of strategic behavior, with a "peak value" for r. We can see the effect of more strategic behavior by looking at the number of steps to convergence, the (lower) agreement with the Plurality winner, and almost any measured property. Intuitively, with low r the voters know the current state exactly, and typically none of them is pivotal. As uncertainty grows a voter considers himself pivotal more often, but beyond the peak r uncertainty is sufficiently large for all voters to believe that their truthful vote is also a possible winner (and then the initial state is stable).

The effect of r, in particular its peak value are determined mainly by the type of the distribution and the number of voters, where the peak r increases with n. The number of candidates may affect the strength of the strategic effect, but not the peak r.

Duverger law As r gets closer to the peak value, over 75% of the voters (all voters in some distributions) end up voting for only two candidates (see Fig. 2).

Winner quality In the Placket-Luce distribution, the quality of a winner can be determined according to its rank in the ground truth used to generate the profile. In the other distributions there is no notion of ground truth, and hence we measured the social welfare of voters, and how often the equilibrium winner agreed with the (truthful) winner of another common voting system or with the Condorcet winner.

According to Borda, Copland, Condorcet consistency, social welfare (see Figure 3) and the ground truth, a clear pattern was observed almost invariably across all distributions. As strategic activity increases, so does the winner quality. In particular, these results are interesting for the Single-Peaked profiles. In such profiles there is always a Condorcet winner, which is the median candidate. As voters strategize more under Plurality, they in fact get closer to the outcome of the strategy-proof median mechanism.

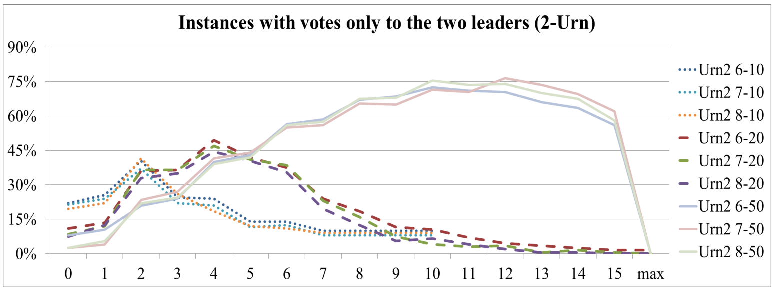

Figure 2: The fraction of simulations in 2-Urn distribution in which *all* voters ended up voting for only two candidates, as a function of r.

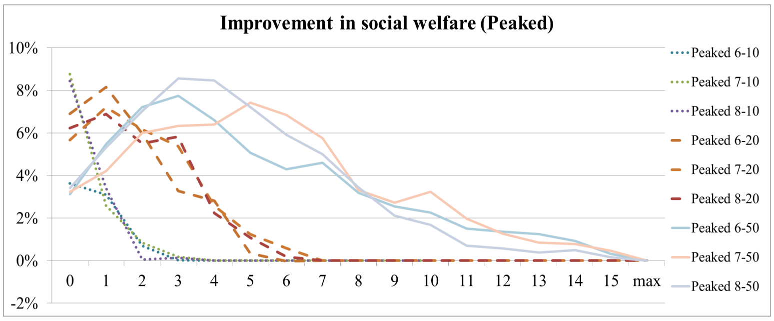

Figure 3: The increase in voters' social welfare, compared to the truthful outcome.

Diverse population Finally, we repeated our simulations with heterogeneous voters, where r_i are sampled uniformly i.i.d. from $\{0, 1, \ldots, \lfloor n/m \rfloor\}$.

The diverse simulations replicated nearly all the patterns of strategic voting across all distributions. Notably, although we used the same simple distribution of r_i values in all simulations, effects of strategic behavior were always approximately as strong as in the peak r value of every profile distribution and across most measured properties.

In particular, winner quality was much higher than the baseline under diverse populations, and generally comparable to quality at peak r.

Regarding Duverger's law, while the number of votes to the top 2 candidates was generally quite similar to the one in peak r (sometimes even higher), with diverse population there were many *fewer* instances where *only* two candidates received votes. See Figure 4. Looking at a typical equilibrium profile reveals that it has a much more "natural" dispersion, with many voters voting for the two leaders, but also some voters (with either very high or very low uncertainty values) voting for other candidates.

7 Discussion

In [1], sophisticated (strategic) voting based on expected utility maximization is defended on the grounds that it "...is a simplification of reality that seeks to capture the most salient features of actual situations. Many voters may see some candidates as having real chances of winning and others as likely losers, and they may weigh these perceptions against the relative attractiveness of the candidates." Our theory is also a simplification of reality, and applies similar logic to explain and justify strategic voting. However, the local-dominance approach

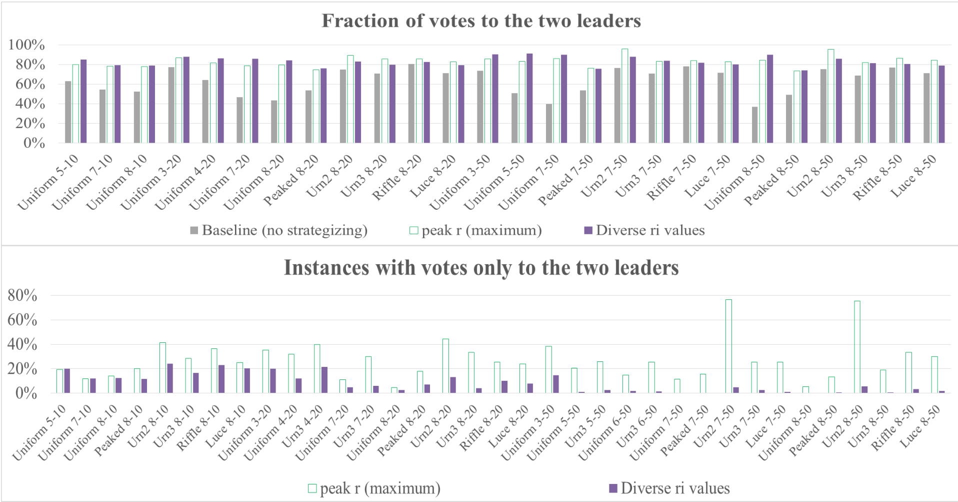

Figure 4: In the top figure we can see that with a diverse population, votes were just as concentrated as with a fixed population with peak r, across all distributions. The bottom figure shows that with fixed r, this concentration is due to many instances where only two candidates get votes, while this is not the case with a diverse population.

allows voters to take into account both "chances of winning" and "relative attractiveness", without regressing to probabilistic calculations and expected utility maximization.

7.1 The model and the desiderata

We summarize by showing how the model of local dominance answers the desiderata we presented in Section 2.

- Looking at *theoretical criteria*, our model is grounded in traditional game-theoretic concepts: voters are trying to maximize their utility, and results are in equilibrium. Further links to decision theory and classical notions of rationality are detailed in the full version.

  When all voters are of the same type, an equilibrium always exists, and convergence of local-dominance dynamics is guaranteed under rather weak conditions. Our simulations show existence and convergence even without these conditions, and demonstrate high discriminative power. The model is broad enough to encompass different scenarios such as simultaneous, sequential and iterative voting, and to account for behaviors such as truth-bias and lazy-bias. Our definitions could be easily extended to other positional scoring rules, although it is an open question whether our results would still hold.

- As argued above, voters in our model fit the *behavioral criteria* we posed, as they avoid complex computations. Moreover, as Lemma 2 shows, voters do not even need to consider the entire space of possible states, but merely to check which candidates have a sufficient score to become possible winners. Our informational assumptions are rather weak and plausible, as we argue at the end of Section 4.1.

- Regarding the *scientific criteria*, once we set the distance metric, every voter can be described by a single parameter (two in the case of lazy or truth-biased voters), which has a clear interpretation as her certainty level. Our extensive simulations demonstrate robustness to the order in which voters play (including whether they act simultaneously or not), and that changing the parameters results in a rather smooth transition. Simulations also show that the model replicates patterns that are common in the real world such as Duverger's Law, and the resulting equilibria, especially with a diverse population, seem reasonable. Experimental validation was outside the scope of this work.

Finally, it is shown that strategic behavior yields a better winner for society according to various measures of quality (compared to the truthful Plurality winner).

7.2 Conclusion and Future Work

We see a unifying theory such as the one we present as a productive step in the quest to understand voting. We hope that future researchers will find our theoretical framework useful for formulating new voting behaviors. Furthermore, our particular distance-based model can serve as a strong baseline for competing theories. Experiments with human voters will be important to settle how close each of these theories comes in adequately describing human voting behavior.

On the technical level, we conjecture that stronger convergence properties can be proved; in particular, that there are no cycles in voting games with voters of the same type, and that a voting equilibrium exists even in games with heterogeneous voters.

We also believe that distance-based local dominance, with the necessary adaptations, can provide a useful non-probabilistic framework for uncertainty in other classes of games where there are natural distance metrics over states, such as congestion games.

Finally, insights based on our theory, for example on how voters' uncertainty level affects quality of the outcome, can be useful in designing better voting mechanisms.

Acknowledgments

This research has been co-financed by Microsoft Research through its PhD Scholarship Programme, Israel Science Foundation grant #1227/12, the Israel Ministry of Science and Technology — Knowledge Center in Machine Learning and Artificial Intelligence grant #3-9243, the Google Inter-University Center for Electronic Markets and Auctions, the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI), and by the Center for Research on Computation and Society at Harvard (CRCS). We thank Joe Halpern, David Parkes, Yaron Singer, Greg Stoddard, and several anonymous reviewers for their valuable feedback.

References

[1] Paul R Abramson, John H Aldrich, Phil Paolino, and David W Rohde. "sophisticated" voting in the 1988 presidential primaries. *American Political Science Review*, 86(01):55–69, 1992.

[2] R Michael Alvarez and Jonathan Nagler. A new approach for modelling strategic voting in multiparty elections. *British Journal of Political Science*, 30(1):57–75, 2000.

[3] Krzysztof R Apt and Sunil Simon. A classification of weakly acyclic games. In *SAGT'12*, pages 1–12. 2012.

[4] William A Brock and Steven N Durlauf. Discrete choice with social interactions. *The Review of Economic Studies*, 68(2):235–260, 2001.

[5] Randall L Calvert. Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science*, pages 69–95, 1985.

[6] Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. Presented in JELIA-04, Lisbon, Portugal, 2004.

[7] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *AAAI*, volume 11, pages 638–643, 2011.

[8] Eddie Dekel and Michele Piccione. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, 108(1):34–55, 2000.

[9] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 347–356. ACM, 2010.

[10] Anthony Downs. An economic theory of democracy. 1957.

[11] Bhaskar Dutta and Arunava Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74(1):154–169, 2012.

[12] Maurice Duverger. *Political Parties: Their Organization and Activity in the Modern State*. New York: John Wiley. Y, 1954.

[13] Aaron S Edlin, Andrew Gelman, and Noah Kaplan. Voting as a rational choice: Why and how people vote to improve the well-being of others. *Rationality and society*, 1, 2007.

[14] Robin Farquharson. *Theory of Voting*. Yale Uni. Press, 1969.

[15] Timothy J. Feddersen, Itai Sened, and Stephen G. Wright. Rational voting and candidate entry under plurality rule. *American Journal of Political Science*, 34(4):1005–1016, 1990.

[16] John A Ferejohn and Morris P Fiorina. The paradox of not voting: A decision theoretic analysis. *The American political science review*, pages 525–536, 1974.

[17] Stephen D Fisher. Definition and measurement of tactical voting: the role of rational choice. *British Journal of Political Science*, 34(1):152–166, 2004.

[18] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and borda score. In *ADT'13*, pages 181–192. 2013.

[19] Eric J Johnson, Colin Camerer, Sankar Sen, and Talia Rymon. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1):16–47, 2002.

[20] Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives*, 5(1):193–206, 1991.

[21] Jean-François Laslier. The leader rule: A model of strategic approval voting in a large electorate. *Journal of Theoretical Politics*, 21(1):113–136, 2009.

[22] Charles F Manski. Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542, 1993.

[23] Richard D. McKelvey and Richard Niemi. A multistage representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, 18:1–22, 1978.

[24] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *ACM-EC'14*, 2014. to appear. For full version see arXiv:1404.4688.

[25] Reshef Meir, Maria Polukarov, Jeffrey S Rosenschein, and Nicholas R Jennings. Convergence to equilibria in plurality voting. In *AAAI'10*, 2010.

[26] Roger B. Myerson and Robert J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.

[27] Svetlana Obraztsova, Evangelos Markakis, and David RM Thompson. Plurality voting with truth-biased agents. In *SAGT'13*, pages 26–37. 2013.

[28] Guillermo Owen and Bernard Grofman. To vote or not to vote: The paradox of nonvoting. *Public Choice*, 42(3):311–325, 1984.

[29] Thomas R Palfrey and Howard Rosenthal. A strategic calculus of voting. *Public Choice*, 41(1):7–53, 1983.

[30] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *AAMAS'12*, pages 635–644, 2012.

[31] Reyhaneh Reyhani, Mark C Wilson, and Javad Khazaei. Coordination via polling in plurality voting games under inertia. In *COMSOC'12*, 2012.

[32] William H Riker and Peter C Ordeshook. A theory of the calculus of voting. *The American Political Science Review*, pages 25–42, 1968.

[33] Jonathan Silberman and Garey Durden. The rational behavior theory of voter participation. *Public Choice*, 23(1):101–108, 1975.

[34] David RM Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey Rosenschein. Empirical analysis of plurality election equilibria. In *AAMAS'13*, pages 391–398, 2013.

[35] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131, 1974.

[36] Hans van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. In *TARK'13*, 2013. To appear.

Reshef Meir
School of Engineering and Applies Sciences
Harvard University
Cambridge, MA
Email: `rmeir@seas.harvard.edu`

Omer Lev
School of Computer Science and Engineering
Hebrew University
Jerusalem, Israel
Email: omerl@cs.huji.ac.il

Jeffrey S. Rosenschein
School of Computer Science and Engineering
Hebrew University
Jerusalem, Israel
Email: jeff@cs.huji.ac.il