

# A Local-Dominance Theory of Voting Equilibria

RESHEF MEIR, Harvard University

OMER LEV, Hebrew University of Jerusalem

JEFFREY S. ROSENSCHEIN, Hebrew University of Jerusalem

We suggest a new model for strategic voting based on *local dominance*, where voters consider a set of possible outcomes without assigning probabilities to them. We prove that voting equilibria under the Plurality rule exist for a broad class of local dominance relations. Furthermore, we show that local dominance-based dynamics quickly converge to an equilibrium if voters start from the truthful state, and we provide weaker convergence guarantees in more general settings. Using extensive simulations of strategic voting on generated and real profiles, we show that emerging equilibria replicate known patterns of human voting behavior such as Duverger's law, and generally improve the quality of the winner compared to truthful voting.

Categories and Subject Descriptors: I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent Systems

Additional Key Words and Phrases: Strict uncertainty; Local dominance; Strategic voting; Equilibrium

## 1. INTRODUCTION

It is often argued that people vote “strategically”, by trying to promote the election of preferable candidates. 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,<sup>1</sup> and focus on a simple scenario that fits all the “standard” assumptions: each of  $n$  voters has complete transitive preferences  $\prec_i$  over a fixed set of alternatives  $M$ , and each voter's only purpose is to bring about the election of her most-favorable alternative possible. 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—voting for one of the alternatives, in 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.

---

<sup>1</sup>For example, social utilities [Manski 1993; Brock and Durlauf 2001], strategic candidates [Calvert 1985; Feddersen et al. 1990], and other considerations (see e.g. [Riker and Ordeshook 1968; Edlin et al. 2007]).

---

This research has been co-financed by the Israel Science Foundation grant #1227/12, Israel Ministry of Science and Technology grant #3-6797, the Google Inter-University Center for Electronic Markets and Auctions, the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI), by the Center for Research on Computation and Society at Harvard (CRCS) and by Microsoft Research through its PhD Scholarship Programme. We thank Joe Halpern, David Parkes, Yaron Singer, Greg Stoddard, and several anonymous reviewers for their valuable feedback.

Authors' addresses: Reshef Meir, School of Engineering and Applied Sciences, Harvard University, Cambridge, MA (rmeir@seas.harvard.edu); Omer Lev, School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel (omer.lev@mail.huji.ac.il); Jeffrey S. Rosenschein, School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel (jeff@cs.huji.ac.il).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

EC'14, June 8–12, 2014, Stanford, CA, USA. Copyright © 2014 ACM 978-1-4503-2565-3/14/06 ...\$15.00.

<http://dx.doi.org/10.1145/2600057.2602860>

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 1.1.* 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 [Duverger 1954].

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 [2004] offers the following definition of 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.* After enumerating the desiderata we believe should guide the search for a proper solution concept, we review some solutions that have been proposed in the literature, and explain where they fall short of meeting these basic requirements. We then lay out our epistemic framework, which is a non-probabilistic way of capturing uncertainty based on distance of voting profiles. Using this framework and simple behavioral assumptions, we define a notion of best-response and the corresponding equilibrium concept. In the remainder of the paper, we will argue, using formal propositions and empirical analysis, that our solution is indeed an 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 in various voter distribution models, using the local-dominance framework enables equilibria with desirable properties — with “better” winners and a “Duverger-like” stable states. Moreover, in complex scenarios for which there is no clear theoretical result on voter behavior, applying our framework resulted in stable equilibria with “desirable” winners and real-world-resembling voter distribution.

A large part of the paper is devoted to discussion of our model and results. All of our technical results with their proofs and the details of our simulations can be found in the full version online.<sup>2</sup>

---

<sup>2</sup>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 to the following classes: Theoretic (mainly game-theoretic), behavioral, and scientific.

### 2.1. 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. E.g., a refinement of Nash equilibrium, or of another popular solution concept from the game theory literature.

**Discriminative power** The predicted outcomes are a small but non-empty set of possible outcomes (sometimes called predictive power). 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.

### 2.2. Behavioral criteria

The assumptions the model makes on 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.

### 2.3. 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 every small perturbation 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*. Although this is not exactly a criterion for a good model (a real strategic behavior may not increase welfare), we are interested in the conditions under which the theory predicts an increased welfare, as these may be useful for design purposes.

### 3. RELATED WORK

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

For example, the *Leader rule* of Laslier [2009] is a strategic model that fares well in all of the above criteria. Unfortunately, it was defined only for Approval voting, and does not have a natural extension to Plurality.

*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 1.1. 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) [Meir et al. 2010; Dutta and Sen 2012]. Nash equilibria under Plurality with truth-biased voters have been studied empirically by Thompson et al. [2013], and analytically by Obraztsova et al. [2013].

Similarly, in some context 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. [Downs 1957; Owen and Grofman 1984]. 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 [Desmedt and Elkind 2010]. 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.<sup>3</sup>

Myerson and Weber’s [1993] *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.

While the Myerson and Weber model is highly attractive in many respects, it suffers from some severe shortcomings. One problem is that it is not clear how voters should find the equilibrium, as the proof is non-constructive. In fact, the voters must engage in

---

<sup>3</sup>Uncertainty have also been proposed as a partial solution to the no-vote paradox [Owen and Grofman 1984].

non-trivial probabilistic reasoning, even if just to verify that they are playing a (mixed) equilibrium strategy.

The fundamental assumption that voters maximize some expected utility function is prevalent in the political science literature, see e.g. [Silberman and Durden 1975; Palfrey and Rosenthal 1983; Alvarez and Nagler 2000]. However, people are notoriously bad at estimating probabilities, and are known to employ various heuristics instead [Tversky and Kahneman 1974].

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 [1974], 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 (similarly to the dominance-based approach in [Dhillon and Lockwood 2004]), and thus all states are considered possible. See also critique on the minmax approach by Aldrich [1993].

A different approach to strict uncertainty that is closer to ours was taken by Conitzer et al. [2011] and followed in [Reijngoud and Endriss 2012; van Ditmarsch et al. 2013]. See Section 7.2 for details.

*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, and in some internal corporate e-mail surveys). The standard solution concept for sequential games is subgame perfect Nash equilibrium [Farquharson 1969; McKelvey and Niemi 1978; Dekel and Piccione 2000; Desmedt and Elkind 2010]. 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 [Johnson et al. 2002].

Iterative voting sounds like a similar setting, but has generated a very different type of voting models. In an iterative setting, voters start from some given voting profile, and in each turn one or more voters may change their vote [Meir et al. 2010; Reijngoud and Endriss 2012]. Meir et al. [2010] 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 [Grandi et al. 2013], which suffer from similar problems.

When all voters are allowed to change their votes at each step, we essentially have repeated polls [Chopra et al. 2004; Reijngoud and Endriss 2012]. A particular model based on uncertainty with iterated polls was suggested by Reyhani et al. [2012]. According to this model, each voter considers some set of *possible winners* based on the poll and on some internal parameter called *inertia*, and votes for her most-preferred candidate in this set.

#### 4. THE FORMAL MODEL

*Basic notations.* We denote  $[x] = \{1, 2, \dots, x\}$ . 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  $\bar{Q}$  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 profile  $\mathbf{s}_{\mathbf{a}} \in \mathbb{N}^m$  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  $s_{\mathbf{a}}(c) >_Q s_{\mathbf{a}}(c')$  if either  $|\{i \in N : a_i = c\}| > |\{i \in N : a_i = c'\}|$ , or  $c, c'$  have the same number of votes and  $c \succ_Q c'$ . We usually omit the subscript from  $>_Q$  as it is clear from the context.

We will use  $\mathbf{a}$  and  $\mathbf{s}_{\mathbf{a}}$  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 outcome. 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**. These estimated scores may come from a poll, from previous voting round, from prior acquaintance with the other voters, etc.<sup>4</sup>

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 considered “possible” in  $\mathbf{s}$ . Our key behavioral assumption is that a voter will avoid voting for candidates that are locally dominated (a standard assumption in strict uncertainty models, see Section 7.2). In an iterative setting, a voter will vote for her most preferred candidate—among those who locally dominate her current action. Our key epistemic assumption (which is new) is that the possible states are those that are “close” to  $\mathbf{s}$  according to some reasonable metric over vote counts.

Tying our model back to Fisher's definition of tactical 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 1.1, where the estimated counts are  $\mathbf{s} = (45, 40, 15)$ . If voter  $v$  also considers as possible states where scores vary by  $\pm 10$ , 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).

#### 4.2. Local dominance

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

<sup>4</sup>This is similar to the prospective score vector in the Myerson and Weber model, except that they do use it to derive probabilities.

*Definition 4.1.* We say that action  $a_i$   $S$ -beats  $a'_i$  (w.r.t. voter  $i$ ) if there is at least one state  $s \in S$  s.t.  $f(s, a_i) \succ_i f(s, a'_i)$ . That is where  $i$  strictly prefers  $f(s, a_i)$  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). The definition, however, does not depend on this interpretation.

*Definition 4.2 (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). See more on epistemic interpretation in Section 7.2, where we also compare our definition of local dominance with previous work. In particular our definition coincides with the definition of dominance in [Conitzer et al. 2011] and with similar definitions in [Reijngoud and Endriss 2012; van Ditmarsch et al. 2013].

*Distance-based dominance.* Suppose we have some distance metric for states, denoted by  $\delta(s, s')$ . 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,  $S_i(\mathbf{a}, x) = \{s' : \delta(s', \mathbf{s}_{a_{-i}}) \leq x\}$ .

The simplest metric we apply is the  $\ell_1$  norm.  $\delta_{\ell_1}(s', s) \leq x$  means that we can attain  $s'$  by adding/removing a total of  $x$  voters to  $\mathbf{a}_{-i}$ . We similarly define other  $\ell_d$  norms.

Another distance we can consider is the *multiplicative distance*, where  $\delta_M(s', s) \leq x$  if for all  $a \in M$ , both  $s'(a) \leq s(a)(1+x)$  and  $s(a) \leq s'(a)(1+x)$ . Intuitively, this means that the score of each candidate can change (either increase or decrease) by a factor of at most  $(1+x)$ .

A third natural distance is the *Earth Mover distance* (EMD), which is similar to  $\ell_1$ , except that all states must have the same number of voters.

### 4.3. Strategic voting and equilibrium

Let  $g_i : M^n \rightarrow M$  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 function  $(g_i)_{i \in N}$  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*. We emphasize that the set of voting equilibria depends only on voters' preferences and response functions, and not on whether they vote iteratively or simultaneously.

*Definition 4.3.* 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 4.4.* 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,  $g_i(\mathbf{a}) = \operatorname{argmin}_{d \in D} Q_i(d)$  if  $D \neq \emptyset$ , and  $g_i(\mathbf{a}) = a_i$  otherwise.

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

- If  $a'_i$   $S_i(\mathbf{a}, r)$ -dominates  $a_i$ , then  $a'_i$   $S_i(\mathbf{a}, r')$ -beats  $a_i$ , for any  $r' \geq r$ .
- For  $r = 0$ , the voter knows the current voting profile exactly, and thus his response function is simple best-response, as in [Meir et al. 2010].

- For  $r = m$ , 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*  $a_i$  (which means  $a_i$  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 4.5.** *Let  $N$  be a set of voters with preferences  $\mathbf{Q}$  and following Def. 4.4 (voters may be of heterogeneous types). A voting profile  $\mathbf{a}$  is a voting equilibrium, if and only if no voter votes for a locally dominated candidate. Formally, if  $\forall i \in N, \exists a'_i \in M$ , such that  $a'_i$   $S_i(\mathbf{a}, r_i)$ -dominates  $a_i$ .*

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 4.4 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 with same score as the winner and lose by tie-breaking, or win in tie-breaking but has one vote less), etc. Let  $\bar{H}_w(\mathbf{s}) = \bigcup_{w' \leq w} H_{w'}(\mathbf{s}) = \{c : s(c) \geq s(f(\mathbf{s})) - w\}$ .

### 5.1. Strategic responses and possible winners

We say that candidate  $c$  is a *possible winner* for  $i$  in state  $\mathbf{a}$  if there is a possible state where  $c$  wins. Formally,  $W_i(\mathbf{a}, r) = \{c \in M : \exists \mathbf{s}' \in S_i(\mathbf{a}, r), f(\mathbf{s}', c) = c\}$ .

In contrast with  $\bar{H}_w(\mathbf{s})$ , the definition of  $W_i(\mathbf{a}, r)$  depends on the identity of the voter, and not only on her type.

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

**LEMMA 5.1.** *Consider a strategic response  $a_i \xrightarrow{k} a'_i$  s.t.  $a_i \notin W_i(\mathbf{a}, r)$ . Then  $a'_i = \operatorname{argmin}_{c \in W_i(\mathbf{a}, r)} Q_i(c)$ .*

The lemma holds for any  $\ell_d$  norm,  $d \geq 1$ , the multiplicative distance, and EM distance.

**PROOF.** Consider the set of candidates  $D$  in Def. 4.4. Since  $a'_i \in D$ , it is non-empty. We will show that (I)  $D \subseteq W_i(\mathbf{a}, r)$ ; and (II)  $b = \operatorname{argmin}_{c \in W_i(\mathbf{a}, r)} Q_i(c)$  is in  $D$ . This implies  $b = a'_i$ .

For (I), every candidate in  $D$  in particular  $S_i(\mathbf{a}, r)$ -beats  $a_i$ , and therefore must be a possible winner (transferring  $i$ 's vote from one non-possible winner to another does not change the outcome in any possible state).

For (II), first note that since  $a'_i \in D$ , there is a state  $\mathbf{s}' \in S_i(\mathbf{a}, r)$  where  $i$  prefers  $f(\mathbf{s}', a'_i) = a'_i$  over  $f(\mathbf{s}', a_i) = a_i$ , thus  $a'_i \succ_i a_i$ . Also, by (I)  $a'_i \in W_i(\mathbf{a}, r)$  and thus  $b \succeq_i a'_i \succ_i a_i$ . Thus  $a_i$  cannot  $S_i(\mathbf{a}, r)$ -beat  $b$ . It remains to prove that  $b$   $S_i(\mathbf{a}, r)$ -beats  $a_i$  (and thus locally dominates it).

We consider two cases. Suppose first that for all  $\mathbf{s}' \in S_i(\mathbf{a}, r)$ ,  $f(\mathbf{s}', a_i) = b$ . Then for any  $c \in M$ , and any  $\mathbf{s}'$ ,  $f(\mathbf{s}', c) \in \{b, c\}$ . Moreover, by definition of  $b$ , if  $f(\mathbf{s}', c) = c$  then  $c \preceq_i b$ . Thus no candidate  $S_i(\mathbf{a}, r)$ -beats  $a_i$ , and  $D = \emptyset$ . Thus suppose there is a state  $\mathbf{s}' \in S_i(\mathbf{a}, r)$ ,  $f(\mathbf{s}', a_i) = c \prec_i b$ , and there is also a state (since  $b$  is a possible winner)  $\mathbf{s}'' \in S_i(\mathbf{a}, r)$ ,  $f(\mathbf{s}'', b) = b$ . It can be verified for each of the metrics we consider that there must be a state  $\mathbf{s}^* \in S_i(\mathbf{a}, r)$  where  $f(\mathbf{s}^*, b) = b$  but  $f(\mathbf{s}^*, a_i) = c$ . For the  $\ell_1$  norm,



we get  $s^*$  from  $s'$  by adding votes to  $b$  until  $c, b$  are tied (or until there is a difference of one for  $c$ , if  $b$  beats  $c$  in  $\overline{Q}$ ).<sup>5</sup> Thus  $b$   $S_i(\mathbf{a}, r)$ -beats  $a_i$ .  $\square$

Lemma 5.1 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 outcomes.

*Threshold for possible winners.* We continue with the following lemma, which shows that for some simple metrics, the set of possible winners is exactly all candidates whose score is above a certain threshold. Denote by  $f_i^* = f(\mathbf{a}_{-i})$  the candidate that would win if  $i$  would not vote.

**LEMMA 5.2.** *Each of the metrics  $\delta$  from  $(\ell_1, \ell_\infty, \text{multiplicative})$  induces a function  $\beta = \beta_{\delta, r} : \mathbb{N} \rightarrow \mathbb{N}$ , where*

- *For every  $\mathbf{a}, i \in N$ , if  $a_i \neq f_i^*$ , then  $W_i(\mathbf{a}, r) = \{c : s_{\mathbf{a}_{-i}}(c) \geq \overline{Q} \beta(s(f_i^*))\}$ . That is, possible candidates are all those whose score is above the threshold, which is a function of the score of the winner.<sup>6</sup>*
- *$\beta(s)$  is weakly increasing in  $s$ .*

*In particular, for  $\delta_{\ell_1}$ ,  $W_i(\mathbf{a}, r) = \overline{H}_{r+1}(\mathbf{s}_\mathbf{a})$  for any  $i$  s.t.  $a_i \notin \overline{H}_{r+1}(\mathbf{s}_\mathbf{a})$ .*

A similar result can be proved for EMD and other  $\ell_d$  norms, but the threshold would depend on the score of all candidates and not just the winner.

**PROOF FOR THE  $\ell_1$  METRIC.** Consider first the  $\ell_1$  norm. We set  $\beta(s) = \beta_{\ell_1}(s) = s - r - 1$ . Clearly,  $f_i^*$  is a possible winner and  $s(f_i^*) > \beta(\mathbf{a})$ . Assume  $c \neq f_i^*$ , then  $s(c) \neq \overline{Q} \beta(s(f_i^*))$ . If  $s(c) > \beta(s(f_i^*))$ , consider the state  $s' \in S_i(\mathbf{a}, r)$  where  $c$  has  $r$  additional votes. We have that  $s'(c) + 1 = s(c) + r + 1 > \beta(s(f_i^*)) + r + 1 = s(f_i^*)$ , thus  $f(s', c) = c$ . In contrast, if  $s(c) < \beta(s(f_i^*))$ , then in any  $s' \in S_i(\mathbf{a}, r)$ ,  $s'(f_i^*) - s'(c) > 1$  and thus  $c$  cannot win. Finally, since  $a_i \notin W_i(\mathbf{a}, r)$ , then  $s_\mathbf{a}(c) = s_{\mathbf{a}_{-i}}(c)$  for all  $c \in W_i(\mathbf{a}, r)$ . Thus  $W_i(\mathbf{a}, r) = \{c : s(c) \geq \beta_{\ell_1}(s(f_i^*))\} = \{c : s(c) \geq s(f_i^*) - (r + 1)\} = \overline{H}_{r+1}(\mathbf{s})$ .  $\square$

Note that Lemmas 5.2 and 5.1 together entail 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 select the one that is most preferred.

## 5.2. Existence of equilibrium and convergence from the truthful state


In what follows, we will only consider the  $\ell_1$  norm for simplicity. However most results hold for other metrics as well (see the full version).

A *scheduler* selects which voter or voters play when there is more than one voter for which  $g_i(\mathbf{a}) \neq a_i$ . A *singleton scheduler* always selects a single voter, while in a group scheduler, several voters may change their votes concurrently.<sup>7</sup> We assume that the order of players is determined by an arbitrary singleton scheduler.

<sup>5</sup>For other  $\ell_d$  norms and multiplicative distance, we can consider any path of states between  $s', s''$  that is contained in  $S_i(\mathbf{a}, r)$ , by gradually removing votes from  $c$  and from other candidates, and adding votes to  $b$ . The critical state  $s^*$  will be along this path. For EM distance some paths may fail if we transfer votes directly from  $c$  to  $b$ , but we can use a path where we transfer votes first from winners to  $a_i$ , and then from  $a_i$  to  $b$ .

<sup>6</sup>We break ties with  $\beta(s(f_i^*))$  as if we break ties with  $f_i^*$ .

<sup>7</sup>We emphasize that a group of voters moving at the same time *does not* coincide with a coalitional manipulation. Each player acts as if he is the only one moving.



**Figure 1:** 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  $\beta_{\ell_1}$  of 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. In the next two figures we can see voters leaving their candidates (who are not possible winners for them) to join one of the leaders. The last figure shows an equilibrium that was reached.

We next show that when starting from the truthful state, singleton scheduler guarantees convergence to an equilibrium. In particular, an equilibrium must exist. Proposition 5.3 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.

**THEOREM 5.3.** *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 singleton 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  $a_i \xrightarrow{i} a'_i$  be a move of voter  $i$  at state  $\mathbf{s} = \mathbf{s}_{\mathbf{a}^t}$  to state  $\mathbf{s}' = \mathbf{s}_{\mathbf{a}^{t+1}}$ . We claim that the following hold throughout the game. Recall that by Lemma 5.2,  $\overline{H}_{r+1}$  is the set of possible winners for all voters who are voting for other candidates.

- (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 5.1,  $a_i$  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),  $\overline{H}_{r+1}(s) \subseteq \overline{H}_{r+1}(s')$ . So if  $a_i \in \overline{H}_{r+1}(s)$ , it must be the most-preferred candidate in the set. Assume, towards a contradiction, that  $a_i \in \overline{H}_{r+1}(s)$ ; then there is a state  $\hat{s} \in S_i(\mathbf{a}, r)$  where  $i$  is pivotal between  $a_i$  and  $f(\mathbf{a}) \prec_i a_i$ . For any  $c \neq a_i$ ,  $f(\hat{s}, c) = f(\mathbf{a})$ , 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 5.1 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 5.1 this never occurs.

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

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. There is always a set of “leaders” ( $\overline{H}_{r+1}$  in the case of the  $\ell_1$  norm). Strategic voters vote for their favorite candidate in this set, if their current candidate is not a possible winner. At some point candidates can “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 strategic voters vote for their best possible winners (which is in  $\overline{H}_{r+1}$ ).

The next lemma characterizes these equilibria more precisely, in particular showing that there are always at least two leaders (see Figure 1).

**LEMMA 5.4.** *Under the conditions of Theorem 5.3, either  $\mathbf{q}$  is stable, or in every state  $s^t$  we have  $|\overline{H}_{r+1}(s^t)| > 1$ . Also, in the last state  $\mathbf{a}$  either  $|\overline{H}_r(s_{\mathbf{a}})| = 1$ , or all voters vote for possible winners. In the former case, the current winner is the most preferred possible winner for all remaining voters (those not voting for a possible winner).*

### 5.3. Convergence under broader conditions

To show robust convergence results, there seem to be two main extensions. First, we would like the game to converge from any initial state, and not just from the truthful one. Second, we would like convergence to occur even if more than one voter moves between states. In other words, we would like to see convergence under any group scheduler.

Unfortunately, under arbitrary group schedulers, convergence is not guaranteed even from the truthful state. A simple example for  $r = 0$  appears in [Meir et al. 2010], where there are two runner-ups that are preferred to the winner by all of their supporters, but the supporters fail to coordinate on a runner-up to promote. We conjecture that as in the case of best-responses ( $r = 0$ ), a *singleton scheduler* would guarantee convergence from any initial state.

We show that if we make two mild restrictions on the scheduler, then convergence from any state is guaranteed even for groups. We say that a step  $a \xrightarrow{i} a'$  is of type 1 if  $a' \prec_i a$ , and type 2 if  $a' \succ_i a$ . We call type 1 steps *compromise steps*, and type 2 *opportunity steps*.

**PROPOSITION 5.5.** *Suppose that all voters are of type  $r$ . Consider any group scheduler such that (1) any voter has some chance of playing as a singleton (i.e., this will occur eventually); (2) the scheduler always selects (an arbitrary subset of) voters with type 2 moves, if such exist. Then convergence is guaranteed from any initial state after at most  $O(nm)$  singleton steps occur.*

The proof shows that in particular, for singleton schedulers there is a path of best-responses from any state to an equilibrium. The assumption that type 2 moves are played first can be justified to some extent, since type 1 moves are “compromises” and thus voters may be more reluctant to carry them out.

#### 5.4. 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  $\epsilon$  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 strategize or to apply truth-bias/laziness, is not necessarily the same neighborhood.

*Definition 5.6.* A strategic truth-biased voter of type  $(r, k)$  (or, in short, a  $T(t, k)$  voter) acts as follows in state  $\mathbf{a}$ .

- (1) (strategic move) 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  $D$ .
- (2) if  $a_i$   $S_i(\mathbf{a}, k)$ -beats  $q_i$ , then  $i$  keeps current vote  $a_i$ .
- (3) (truth bias move) otherwise,  $i$  moves to  $q_i$ .

A *strategic lazy voter of type  $(r, k)$*  (an  $L(r, k)$  voter) can be similarly defined, replacing  $q_i$  with the action  $\perp$  (abstain).

We make the following observations.

- A  $T(r, n)$  or  $L(r, n)$  voter is very similar to an  $r$  voter, except never voting for a globally dominated candidate.
- $T(0, 0)$  voters follow the truth-bias model in [Thompson et al. 2013; Obratzsova et al. 2013].
- $L(0, 0)$  voters follow the (simultaneous) “lazy” model of Desmedt and Elkind [2010].

We argue that it is natural to assume  $k > r$ , which entails that a voter requires a higher certainty level in order to make a new strategic step, 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 [Kahneman et al. 1991].

**PROPOSITION 5.7.** *Suppose that each voter  $i$  is of type  $L(r, k_i)$  or  $T(r, k_i)$ , where  $k_i > r$ . Then a voting equilibrium exists. Moreover, in an iterative setting where voters start from the truthful state, they will always converge to an equilibrium in at most  $3nm$  steps.*

We conclude with some properties that must hold for any voting equilibrium, whether it was obtained by an iterative process or not.

**PROPOSITION 5.8.** *In any voting equilibrium it holds that:*

- (1) Every voter is either truthful, or votes for a candidate in  $W_i(\mathbf{a}, k_i)$ .
- (2) No untruthful voter votes for his least-preferred candidate in  $W_i(\mathbf{a}, k_i)$ .
- (3) If  $|W_i(\mathbf{a}, r_i)| > 1$ , then  $a_i$  is not  $i$ 's least-preferred candidate in  $W_i(\mathbf{a}, r_i)$ .

It is possible, though, that all voters vote for their second-least-preferred candidate (consider  $a, b$  ranked last by all, where roughly half rank  $a$  above  $b$ ).

## 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 (for example, how long does it take to reach an equilibrium if we vary the uncertainty level?). More importantly, we use simulations to test the properties of our strategic model with respect to the desiderata listed on Section 2. For example, what is its discriminative power, is it robust to small changes, and whether it replicates common phenomena. See Section 7.1 for a summary of our findings in light of the desiderata.

We generate preference profiles from a set of distributions which have been examined in the research literature, with a focus on distributions that are claimed to resemble preferences of human societies: The Uniform (or impartial culture) distribution; a uniform Single-peaked distribution; Polya-Eggenberger Urn model (with 2 urns and with 3 urns); a Riffle distribution; and Plackett-Luce distribution. Urn models were particularly designed to resemble preference structures in human societies, whereas in Plackett-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 on voters' preferences including all 225 currently available full preferences from PrefLib (<http://preflib.org>). See the full version for details on our methods, including the distributions and profiles we used.

*Methods.* We generated profiles from all distribution types for various numbers of voters and candidates, which resulted in 108 distinct distributions. From each distribution we sampled 200 instances. Then, we simulated strategic voting on each instance varying the distance metric ( $\ell_1$ , multiplicative), 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, and collected multiple statistics on the equilibrium outcomes. All of the collected data can be downloaded from <http://tinyurl.com/lm9axkq>.

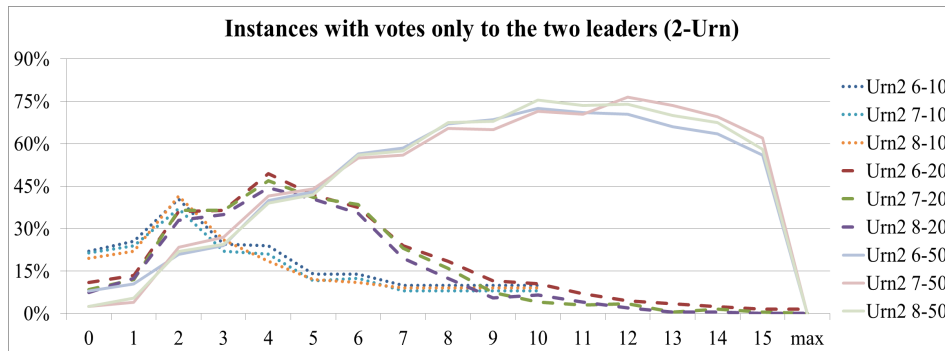
### 6.1. Results

We bring here the main findings for simulations starting from the truthful state, for the  $\ell_1$  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.

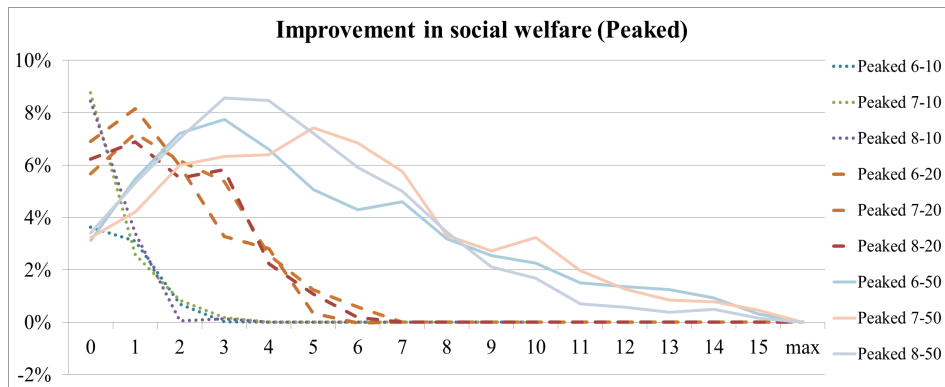
We should note that despite the fact that our convergence proofs do not cover heterogeneous populations, arbitrary simultaneous moves or arbitrary starting points, convergence was just as robust under all of these conditions.

**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  $r$ .



**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.

**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 Plackett-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.

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

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



**Figure 4:** In the top figure we can see that with diverse population, votes were just as concentrated as with 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 diverse population.

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.

Winner quality was also generally comparable to peak  $r$ . So was the reproduction of Duverger’s law, except that there were almost no instances where *all* voters vote for the top 2 candidate. Rather, we have a more natural dispersion where votes are highly concentrated, yet all candidates get some votes. See Figure 4.

**7. DISCUSSION**

In [Abramson et al. 1992], 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 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 model of local dominance answers to 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 Section 7.2.

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.<sup>8</sup> Furthermore, as Proposition 5.8 shows, our model rules out many unnatural equilibria.

As argued above, voters in our model fit the *behavioral criteria* we posed, as they avoid complex computations. Moreover, as Lemma 5.2 shows, voters do not even need to consider the entire space of possible states, but merely to check which candidates have sufficient score to become possible winners. Our informational assumptions are rather weak and plausible, as we argue in 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 the Duverger Law, and resulting equilibria, especially with 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 the society according to various measures of quality (compared to the truthful Plurality winner).

## 7.2. Epistemic foundations and rationality

We can phrase dominance relations in terms of modal logic. Consider a Kripke structure over states where  $S_i(\mathbf{a}, r)$  are the states *accessible* from  $s = s_{\mathbf{a}}$ . We note that  $S_i(\mathbf{a}, r)$  defines a Kripke structure that is reflexive and symmetric but non-transitive.

Given a proposition  $P$  (say, “candidate  $b$  wins”) we say that  $P$  is *necessary* in  $s$  if  $P$  holds in all worlds accessible from  $s$ . A common interpretation is to say that “[agent  $i$ ] knows  $P$  in state  $s$ ”, denoted by  $s \models K_i P$  (see, for example, [Aumann 1999]). According to this, we can write “ $b_i$   $S$ -dominates  $a_i$  in state  $s$ ” as the epistemic statement  $s \models K_i(f(b_i, s) \succeq_i f(a_i, s)) \wedge \neg K_i(f(b_i, s) \prec_i f(a_i, s))$  (“ $i$  knows that [voting for]  $b_i$  is at least as good as  $a_i$ , but does not know that  $a_i$  is at least as good as  $b_i$ ”).

Since in our model  $S_i(\mathbf{a}, r) \subseteq S_i(\mathbf{a}, k)$ , a straight-forward extension of the epistemic interpretation is to add *certainty levels*, where a larger (in terms of containment) set of accessible states indicates lower certainty.

*Local dominance and rationality.* According to the standard non-Bayesian incomplete information model (due to Aumann [1995; 1999]), a player  $i$  playing strategy  $a_i$  at some world state  $\pi$  is *rational*, if there is no other strategy  $a'_i$  that yields a same or better outcome in all states accessible from  $\pi$  (and in some states strictly better).

In other words, rationality under strict uncertainty according to Aumann simply means that players avoid locally dominated strategies. Voting equilibria in our model are therefore rational (Prop. 4.5). Our model is more specific in that it specifies a particular dynamic of how voters act when their current strategy is dominated.

*Local dominance and voting.* Dominance within a restricted set of states was considered by several recent papers. In [Reijngoud and Endriss 2012; van Ditmarsch et al. 2013] the assumption is that voters information sets can be described as a partition  $\Pi$ —a common assumption about knowledge also made by Aumann. Reijngoud and Endriss say that a voter *has an incentive to  $\Pi$ -manipulate using ballot  $P'_i$*  (under profile

<sup>8</sup>Extensions to other voting rules are simple with the EM distance, but may be ill-defined with other metrics.



P), if she weakly gains by voting  $P'_i$  in every state that is “equivalent” to P according to  $\Pi$ .<sup>9</sup> In the special case of Plurality, the definition coincides local dominance: Consider Def. 4.2, where we set  $S$  to be all states equivalent to a under  $\Pi$ . Then  $a'_i$   $S$ -dominates  $a_i$  iff  $i$  has an incentive to  $\Pi$ -manipulate using ballot  $a'_i$ . In the terminology of [van Ditmarsch et al. 2013], voter  $i$  *knows ‘de re’ that she can weakly successfully manipulate*. Our definition of local dominance also coincides with the definition of dominance in [Conitzer et al. 2011], which do not make any assumption on the “information set”  $S$ .

In our work the accessibility relation is defined by a distance metric and is *not* a partition. Still, many of the definitions in [Reijngoud and Endriss 2012; van Ditmarsch et al. 2013] can be applied just the same in our case. In particular, a combination of these works can be used to extend the notion of local dominance to other voting rules.

### 7.3. Conclusion and future work

We see a unifying theory 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, more specific, 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.

Finally, we 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.

## REFERENCES

- Paul R Abramson, John H Aldrich, Phil Paolino, and David W Rohde. 1992. “Sophisticated” Voting in the 1988 Presidential Primaries. *American Political Science Review* 86, 01 (1992), 55–69.
- John H Aldrich. 1993. Rational choice and turnout. *American Journal of Political Science* (1993), 246–278.
- R Michael Alvarez and Jonathan Nagler. 2000. A new approach for modelling strategic voting in multiparty elections. *British Journal of Political Science* 30, 1 (2000), 57–75.
- Robert J Aumann. 1995. Backward induction and common knowledge of rationality. *Games and Economic Behavior* 8, 1 (1995), 6–19.
- Robert J Aumann. 1999. Interactive epistemology I: knowledge. *International Journal of Game Theory* 28, 3 (1999), 263–300.
- William A Brock and Steven N Durlauf. 2001. Discrete choice with social interactions. *The Review of Economic Studies* 68, 2 (2001), 235–260.
- Randall L Calvert. 1985. Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science* (1985), 69–95.
- Samir Chopra, Eric Pacuit, and Rohit Parikh. 2004. Knowledge-Theoretic Properties of Strategic Voting. Presented in JELIA-04, Lisbon, Portugal. (2004).
- Vincent Conitzer, Toby Walsh, and Lirong Xia. 2011. Dominating Manipulations in Voting with Partial Information.. In *AAAI*, Vol. 11. 638–643.
- Eddie Dekel and Michele Piccione. 2000. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy* 108, 1 (2000), 34–55.
- Yvo Desmedt and Edith Elkind. 2010. Equilibria of plurality voting with abstentions. In *ACM-EC’10*. 347–356.

<sup>9</sup>The definition is for arbitrary voting rules.

- Amrita Dhillon and Ben Lockwood. 2004. When are Plurality Rule Voting Games Dominance-solvable? *Games and Economic Behavior* 46 (2004), 55–75.
- Anthony Downs. 1957. An economic theory of democracy. (1957).
- Bhaskar Dutta and Arunava Sen. 2012. Nash implementation with partially honest individuals. *Games and Economic Behavior* 74, 1 (2012), 154–169.
- Maurice Duverger. 1954. *Political Parties: Their Organization and Activity in the Modern State*. New York: John Wiley. Y.
- Aaron S Edlin, Andrew Gelman, and Noah Kaplan. 2007. Voting as a rational choice: Why and how people vote to improve the well-being of others. *Rationality and society* 1 (2007).
- Robin Farquharson. 1969. *Theory of Voting*. Yale Uni. Press.
- Timothy J. Feddersen, Itai Sened, and Stephen G. Wright. 1990. Rational Voting and Candidate Entry under Plurality Rule. *American Journal of Political Science* 34, 4 (1990), 1005–1016.
- John A Ferejohn and Morris P Fiorina. 1974. The paradox of not voting: A decision theoretic analysis. *The American political science review* (1974), 525–536.
- Stephen D Fisher. 2004. Definition and measurement of tactical voting: the role of rational choice. *British Journal of Political Science* 34, 1 (2004), 152–166.
- Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. 2013. Restricted Manipulation in Iterative Voting: Condorcet Efficiency and Borda Score. In *ADT'13*. 181–192.
- Eric J Johnson, Colin Camerer, Sankar Sen, and Talia Rymon. 2002. Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory* 104, 1 (2002), 16–47.
- Daniel Kahneman, Jack L Knetsch, and Richard H Thaler. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives* 5, 1 (1991), 193–206.
- Jean-François Laslier. 2009. The Leader Rule: A Model of Strategic Approval Voting in a Large Electorate. *Journal of Theoretical Politics* 21, 1 (2009), 113–136.
- Charles F Manski. 1993. Identification of endogenous social effects: The reflection problem. *The review of economic studies* 60, 3 (1993), 531–542.
- Richard D. McKelvey and Richard Niemi. 1978. A multistage representation of sophisticated voting for binary procedures. *Journal of Economic Theory* 18 (1978), 1–22.
- Reshef Meir, Maria Polukarov, Jeffrey S Rosenschein, and Nicholas R Jennings. 2010. Convergence to Equilibria in Plurality Voting. In *AAAI'10*.
- Roger B. Myerson and Robert J. Weber. 1993. A Theory of Voting Equilibria. *The American Political Science Review* 87, 1 (1993), 102–114.
- Svetlana Obraztsova, Evangelos Markakis, and David RM Thompson. 2013. Plurality Voting with Truth-Biased Agents. In *SAGT'13*. 26–37.
- Guillermo Owen and Bernard Grofman. 1984. To vote or not to vote: The paradox of nonvoting. *Public Choice* 42, 3 (1984), 311–325.
- Thomas R Palfrey and Howard Rosenthal. 1983. A strategic calculus of voting. *Public Choice* 41, 1 (1983), 7–53.
- Annemieke Reijngoud and Ulle Endriss. 2012. Voter response to iterated poll information. In *AAMAS'12*. 635–644.
- Reyhaneh Reyhani, Mark C Wilson, and Javad Khazaei. 2012. Coordination via polling in plurality voting games under inertia. In *COMSOC'12*.
- William H Riker and Peter C Ordeshook. 1968. A Theory of the Calculus of Voting. *The American Political Science Review* (1968), 25–42.
- Jonathan Silberman and Garey Durden. 1975. The rational behavior theory of voter participation. *Public Choice* 23, 1 (1975), 101–108.
- David RM Thompson, Omer Lev, Kevin Leyton-Brown, and Jeffrey Rosenschein. 2013. Empirical analysis of plurality election equilibria. In *AAMAS'13*. 391–398.
- Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science* 185, 4157 (1974), 1124–1131.
- Hans van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. 2013. Strategic voting and the logic of knowledge. In *TARK'13*. To appear.