

Plurality Voting under Uncertainty

Reshef Meir
Harvard University

Abstract

Understanding the nature of strategic voting is the holy grail of social choice theory, where game-theory, social science and recently computational approaches are all applied in order to model the incentives and behavior of voters. In a recent paper, Meir et al. (2014) made another step in this direction, by suggesting a behavioral game-theoretic model for voters under uncertainty. For a specific variation of best-response heuristics, they proved initial existence and convergence results in the Plurality voting system.

This paper extends the model in multiple directions, considering voters with different uncertainty levels, simultaneous strategic decisions, and a more permissive notion of best-response. It is proved that a voting equilibrium exists even in the most general case. Further, any society voting in an iterative setting is guaranteed to converge to an equilibrium.

An alternative behavior is analyzed, where voters try to minimize their worst-case regret. As it turns out, the two behaviors coincide in the simple setting of Meir et al. (2014), but not in the general case.

Introduction

Suppose that ALICE, BOB, and CHARLIE run for office. A and B are currently leading the polls with 45% and 40%, respectively. Your favorite candidate C is currently trailing behind with 15% (see Fig. 1, left).

A game-theorist advice in such a case might be to “vote strategically,” by supporting your next-preferred candidate (say, B). Indeed, desertion of candidates that seem to lose height on polls in common in the real world, even if they are still perceived as suitable.

However, there is no consensus whatsoever regarding how to generalize this straight-forward compromise to an arbitrary situation: should you leave C if he has 25% of the votes? and what if there are three or more candidates that are more popular than your favorite? does the source of the poll matter? and so on.

This lack of a conclusive notion for strategic voting is partly due to the fact that there are many ways to describe the information voters have, as well as their beliefs over other voters’ preferences and actions when casting their

vote. Moreover, even for a given belief there may be several different actions that can be justified as rational. Hence any model of voting behavior and equilibrium must state explicitly its epistemic and behavioral assumptions.

A simple starting point would be to apply some “standard” assumptions from game-theory, for example that voters play a Nash equilibrium of the game induced by their preferences. However, such a prediction turns out to be very uninformative: a single voter can rarely affect the outcome, and thus almost any way the voters vote is a Nash equilibrium (including, for example, when all voters vote for their least preferred candidate). It is therefore natural to consider the role of *uncertainty* in voters’ decisions.

Probabilistic and strict uncertainty A classical economic approach is to assume that voters maximize some expected utility, given some prior distribution on other voters’ preferences and actions (see e.g. the MW model in Related Work). However the assumption that voters have access to an accurate, or even approximate, distribution of this sort seems hard to justify. Furthermore, studies in behavioral psychology show, human decision makers often ignore probabilistic information even when it is given, employing various heuristics instead (Tversky and Kahneman 1974). It is thus unlikely that people are able to represent complex distributions, or to compute and optimize their own expected utility, let alone the equilibria of such games.

In a recent effort to reconcile well-founded decision making approaches with a formal game-theoretic analysis, Meir et al. (2014) suggested a model for Plurality voting relying on *strict uncertainty* and *local dominance*. Informally, voters’ beliefs can be described by a single vector of candidates’ scores. This prospective score vector may be the result of a poll, derived from acquaintance with the other voters, from the outcome of a previous round of voting, etc. Each voter has a single uncertainty level, reflecting how sure she is in the correctness of the prospective scores—higher uncertainty means the voter considers a larger range of outcomes (score vectors) as “possible.”

In the example introduced above, a voter with an intermediate uncertainty level will consider a tie between A and B as “possible,” but will reject any outcome where C wins as impossible.

Given this uncertain, non-probabilistic view, it is not a-

priori clear how a voter should act. Meir et al. adopted one of the classic approaches to decision making under strict uncertainty, assuming that a voter i will refrain from voting for a candidate a_i that is *locally dominated* by another candidate a'_i . Intuitively, a'_i *locally-dominates* a_i if it is always at least as good, and sometimes strictly better, to vote for a'_i than for a_i (taking into account all outcomes that voter i considers possible). A *voting equilibrium* is simply a state where no voter votes for a locally-dominated candidate according to her beliefs. Going back to our example, our voter will consider candidate c as locally-dominated by candidate b , and will thus strategize by voting for b .

Informally, Meir et al. (2014) proved that if all voters have the *same uncertainty level*, start by voting truthfully, and play one at a time, then they always converge to a voting equilibrium. In the special case of zero uncertainty, this simply means that the best-response dynamics converges to a Nash equilibrium (which is known from (Meir et al. 2010)).

However, these assumptions are rather restrictive. Some voters may be less informed than others, or simply more “stubborn” (e.g., require even fewer votes to C in order to be convinced he cannot win). Also, there is no reason to believe that voters initially vote for their most preferred candidate, or that only one voter changes her vote after each poll or voting round (see, e.g., (Reijngoud and Endriss 2012)). Meir et al. (2014) showed *empirically* that convergence occurs even without those assumptions, and conjectured that at least some of the assumptions could be relaxed. However they provided no formal guarantee for such convergence.

Research goals The main purpose of this work is to close this gap, and to prove equilibrium existence and convergence under conditions that are as broad as possible.

While we adopt the model of Meir et al. at large, we study a non-atomic variation of it, where the effect of any single voter is negligible. This simplifies the model, and allows us to prove stronger results. We modify the distance function used to determine the possible ranges of candidates’ scores, to one that is better justified by psychological studies (for details see Footnote 4).

In addition, we are interested whether the behavioral assumptions can be weakened, or, alternatively, be replaced with a more nuanced way to select among several undominated candidates. To that end, we define a variation of the model where voters are aiming to minimize their *worst-case regret* over all states they believe possible, and compare this behavior to voting under local dominance.

Our contribution

Our main result is proving that voters with local dominance behavior always converge to an equilibrium. This holds for any population of voters with different preferences and uncertainty levels, for any initial voting profile, and for any order of moves (including moves of arbitrary subsets of voters, and suboptimal moves).

We then turn to study voters minimizing worst-case regret. When voters have the same uncertainty level and start from the truthful profile, we show that regret minimization

coincides with local dominance. However if these requirements are relaxed then the behaviors may significantly differ, and even the existence of equilibrium is not guaranteed.

Due to space constraints, some of the proofs were omitted, and are available in the full version of this paper.¹ In the full version we also show that all of our results hold in the finite case, for voters that move one-at-a-time.

The Formal Model

Basic notations Where possible, we follow the notations and definitions of (Meir, Lev, and Rosenschein 2014). We denote $[x] = \{1, 2, \dots, x\}$ for all $x \in \mathbb{N}$. For a finite set A , $\Delta(A)$ denotes the set of all probability distributions over A (all non-negative vectors of size $|A|$ that sum to 1).

The set of candidates is denoted by M , where $m = |M|$. Let $\mathcal{Q} = \pi(M)$ be the set of all strict orders (permutations) over M . We encode all the information on a voter, including her preferences, in her *type*,² where the set of all types is denoted by \mathcal{V} .

Specifically, a voter of type $v \in \mathcal{V}$ has preferences $Q_v \in \mathcal{Q}$, where $Q_v(a) \in [m]$ is the rank of candidate $a \in M$ (lower is better), and $q_v = Q_v^{-1}(1)$ is her most-preferred candidate. We denote $a \succ_v b$ if $Q_v(a) < Q_v(b)$.

There is no finite set of voters. Rather, a *preference profile* $\mathbf{Q} \in \Delta(\mathcal{Q})$ is a distribution over preferences, specifying the fraction of voters with each preference order. A *population* $\mathbf{V} \in \Delta(\mathcal{V})$ is a distribution over *types*, that is, a preference profile aggregated with any additional information specified by voters’ types (e.g., their beliefs and behaviors).

Under the Plurality rule, every voter selects a single candidate. A given *action profile* should specify how many voters of each type vote for every candidate. Formally, an action profile (also called *state*) \mathbf{a} is a refinement of \mathbf{V} , where $a(v, c) \in \mathbb{R}_+$ denotes the fraction of voters of type v who vote for $c \in M$. Intuitively, the winner under state \mathbf{a} is the candidate c getting at most votes, i.e., maximizing $\sum_v a(v, c)$.

Our next definitions are intended to enable easier notations and analysis. We treat the voter set as if it is partitioned to subsets whose size is ϵ , for some arbitrarily small $\epsilon > 0$.³ We denote by I the collection of these $1/\epsilon$ sets, and refer to an arbitrary voter in the set $i \in I$ as “voter i .” W.l.o.g. all voters in set $i \in I$ are indistinguishable: they have the same type v_i and the same action a_i in any profile \mathbf{a} (otherwise we just further split set i to smaller uniform subsets).

Score vectors and tie-breaking The outcome of the Plurality rule in state \mathbf{a} , denoted $f(\mathbf{a}) \in M$, only depends on the total number of votes for each candidate. We thus define the *score vector* $\mathbf{s}_\mathbf{a}$ induced by action profile \mathbf{a} . That is, $s_\mathbf{a}(c) = |\{i \in I : a_i = c\}| \epsilon = \sum_{v \in \mathcal{V}} a(v, c)$. We will use \mathbf{a} and $\mathbf{s}_\mathbf{a}$ interchangeably, sometimes omitting the subscript \mathbf{a} . We denote by $\mathcal{S} = \mathbb{R}_+^m$ the set of all score vectors. We

¹<http://arxiv.org/abs/1411.4949>

²We later extend the definition of a type to also include the belief structure and behavior of the voter.

³This is not a restrictive assumption. We further discuss its implications in the full version.

also sometimes refer to score vectors as states, although \mathbf{s} may not be attained from an actual action profile. E.g., it is possible that $\sum_{c \in M} s(c) \neq 1$. Note that we may only use the aggregate scores \mathbf{s} in a context where voters’ identifiers are not important.

Every score vector \mathbf{s} implicitly contains an arbitrary “tie-breaker” $Q_{\mathbf{s}}^+ \in \mathcal{Q}$. The *winner* $f(\mathbf{s})$ is the candidate $c \in M$ whose score $s(c)$ is maximal. If there is more than one candidate with maximal score, we break ties according to $Q_{\mathbf{s}}^+$. Formally, $f(\mathbf{s}) = \operatorname{argmin}_{c \in \operatorname{argmax}_{s(c')} Q_{\mathbf{s}}^+(c)}$. Note that the outcome is well defined even for score vectors that are not derived from valid action profiles.

Dynamics and equilibria We describe the behavior of a voter of type $v \in \mathcal{V}$ by a *response function* $g_v : M \times \mathcal{S} \rightarrow 2^M \setminus \emptyset$. That is, a mapping from the current state (only taking into account aggregate scores) and current action, to a subset of actions. Together, Q_v and g_v completely define the type v . For an identified voter $i \in I$, we write $g_i(\mathbf{s})$ instead of $g_{v_i}(a_i, \mathbf{s})$. We also write $g_i(\mathbf{a})$ as a shorthand for $g_i(\mathbf{s}_{\mathbf{a}})$. Intuitively, this means that a voter $i \in I$ may choose any action in $g_i(\mathbf{a})$.

Definition 1. A *voting equilibrium* for population \mathbf{V} is a state \mathbf{a} , where $g_i(\mathbf{a}) = \{a_i\}$ for all $i \in I$.

For example, if g_i is the *best-response function* ($g_i(\mathbf{a}) = \{c \in M : c \text{ maximizes the utility of } i \text{ in } \mathbf{a}\}$ for all i), then a voting equilibria coincides with the Nash equilibria of the game. We emphasize that classic results on existence of equilibria in nonatomic games (e.g., (Schmeidler 1973)) do not apply, since even if we assign cardinal utilities to voters, those would be highly discontinuous in the action profile. Note that we do not assume that voters with the same type, or even the same identifier, are coordinated.

Even in cases where an equilibrium exists, players may or may not reach one, depending on their initial state and the order in which they play. We are therefore interested in sufficient conditions under which the game is *acyclic*, i.e. there are no finite cycles of moves by groups of any size.

Uncertainty and Strategic Voting

The most important part of the model is of course the way we define the behavior, which is the response function g_i . This behavior depends on voters’ preferences, and also on their beliefs about the current state. We assume voters derive their beliefs using a distance-based strict uncertainty model, following (Meir, Lev, and Rosenschein 2014).

Distance-based uncertainty

Suppose we have some distance measure for score vectors, denoted by $\delta : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_+$. For any \mathbf{s} , let $S(\mathbf{s}, x) = \{\mathbf{s}' : \delta(\mathbf{s}, \mathbf{s}') \leq x\}$ be the set of vectors that are at distance at most x from \mathbf{s} . A distance function δ is *candidate-wise* if it can be written as $\delta(\mathbf{s}, \mathbf{s}') = \max_{c \in M} \delta(s(c), s'(c))$ for some monotone function $\hat{\delta}$ (meaning that for a fixed s , $\hat{\delta}(s, s')$ is nondecreasing in $|s - s'|$). Thus $\mathbf{s}' \in S(\mathbf{s}, x)$ if and only if the score of every candidate in \mathbf{s}' is sufficiently close to its score in \mathbf{s} , see Fig. 1.

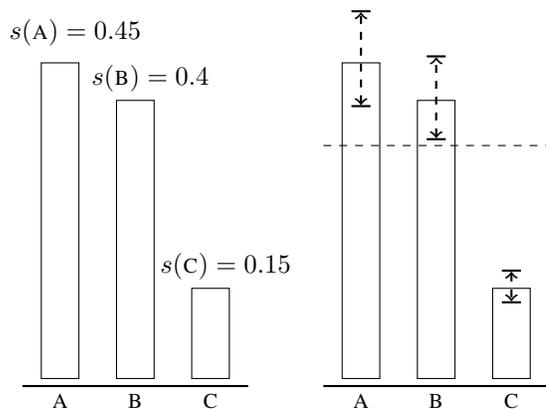


Figure 1: On the left figure we see a given state \mathbf{s} , which can be thought of as the result of a poll, or of the current voting round. On the right we depict the set of possible states $S(\mathbf{s}, r)$ for $r = 0.15$. Any \mathbf{s}' is a possible state as long as the score of each candidate is in the marked range. The dashed line marks the threshold above which a candidate is considered a possible winner.

In (Meir, Lev, and Rosenschein 2014), several metrics have been suggested for the function δ , including various ℓ_d norms. Among those the multiplicative distance (where $\hat{\delta}(s, s') = \max\{s'/s, s/s'\} - 1$) and the ℓ_∞ metric (where $\hat{\delta}(s, s') = |s - s'|$) are candidate-wise.

While all of our results apply for any candidate-wise distance (see full version), for concreteness we assume throughout the paper the multiplicative distance. Note that the multiplicative distance is independent of the amount of voters, i.e. $\delta(\mathbf{s}, \mathbf{s}') = \delta(\alpha\mathbf{s}, \alpha\mathbf{s}')$ for all $\alpha > 0$. This is consistent with findings on how people perceive uncertainty over numerical values (Kahneman and Tversky 1974; Tversky and Kahneman 1974).⁴

The intensity of probabilistic uncertainty is typically captured by some *variance*. The intensity of the agent’s uncertainty in our non-probabilistic model is captured by the *uncertainty parameter* $r_v \in \mathbb{R}_+$, which is part of the agent’s type. Given a profile \mathbf{a} , voter $i \in I$ believes that the actual state may be any $\mathbf{s}' \in S(\mathbf{s}_{\mathbf{a}}, r_i)$, where $r_i = r_{v_i}$.

A voter facing strict uncertainty may use various heuristics when selecting a strategy. In this work we follow two standard approaches: avoiding dominated strategies (Aumann 1999), and minimizing worst-case regret (Savage 1951; Hyafil and Boutilier 2004).

⁴For example, the studies show that if the average number of girls born daily in a hospital is s , then people believe that the probability that on a given day the number is within $[(1-r)s, (1+r)s]$ is fixed and does not depend on s . Kahneman and Tversky highlight that this reasoning stands in contrast to the scientific truth in this case, where the range r is proportional to $1/\sqrt{s}$. In our case we can think of the score of a candidate as the limit of a Poisson variable (rather than Binomial in the hospital example), and thus the range $[s/(1+r), s(1+r)]$ is more appropriate.

Voter influence In the nonatomic case, every voter has negligible influence, yet they still prefer some actions over others. How does the outcome change when a voter votes for some candidate c ? To solve this discrepancy, we assume that a voter considers herself as pivotal only if she happens to break an exact tie. Formally, we define $f(s, c) = c$ if c has maximal score in s (overriding the default tie-breaker Q_s^+), and $f(s, c) = f(s)$ otherwise.

Local dominance

The first approach we consider follows (Conitzer, Walsh, and Xia 2011; Meir, Lev, and Rosenschein 2014), where voting dynamics is determined by local dominance relations.

Consider a particular voter $i = \langle v_i, a_i \rangle$. Let $S = S(s, r_i)$. We say that action a_i *S-beats* b_i if there is at least one $s' \in S$ s.t. $f(s', a_i) \succ_i f(s', b_i)$. That is where i strictly prefers $f(s', a_i)$ over $f(s', b_i)$. Action a_i *S-dominates* b_i if (I) a_i *S-beats* b_i ; and (II) b_i does not *S-beat* a_i .

We next define the response function that strategic voters apply under the local dominance behavioral model. Consider a voter i , where $v_i = \langle Q_i, r_i \rangle$.

Definition 2 (Strategic move under local dominance (LD)). Let $D \subseteq M$ be the set of candidates that $S(\mathbf{a}, r_i)$ -dominate a_i , and that are not $S(\mathbf{a}, r_i)$ -dominated. If $D = \emptyset$ then $g_i(\mathbf{a}) = \{a_i\}$. Otherwise:

- For a *weak LD voter*, $g_i(\mathbf{a}) = D$.
- For a *strict LD voter*, $g_i(\mathbf{a}) = \operatorname{argmin}_{d \in D} Q_i(d)$.

In other words, a strict LD voter votes the most preferred candidate that locally-dominates her current choice (if such exists). The definition of a weak LD voter is much more permissive, and does not restrict the voter to select the most preferred candidate in D . A higher value of r_i may either indicate that the voter is less informed, or simply that she requires stronger evidence that a move will be beneficial, a tendency we can interpret as stronger *loss aversion* (Kahneman, Knetsch, and Thaler 1991).

The following observation is immediate from Definitions 1 and 2.

Proposition 1. *Let \mathbf{V} be a population of LD voters. A profile \mathbf{a} is a voting equilibrium iff $\forall i \in I, \exists a'_i \in M$, such that a'_i $S(\mathbf{a}, r_i)$ -dominates a_i . I.e., if no voter votes for a locally dominated candidate.*

Convergence with LD Voters

We say that candidate c is a *possible winner* for i in state s if there is a possible state where c wins. Formally, $W_i(s) = \{c \in M : \exists s' \in S(s, r_i) \text{ s.t. } f(s', c) = c\}$. Also denote $W_0(s) = \{c \in M : f(s, c) = c\} = W_i(s)$ for $r_i = 0$, and $W_i(s) = W_i(s_{\mathbf{a}})$.

It is easy to see that under the multiplicative distance $c \in W_i(s)$ iff $s(c) \geq (1 + r_i)^{-2} s(f(s))$, and that similar thresholds exist for other metrics, see (Meir, Lev, and Rosenschein 2014) and the full version of this paper. The next key lemma, however, holds only for candidate-wise distances, which means that the possible scores of the different candidates are calculated independently (it is the only place in the paper where the candidate-wise assumption is used).

Lemma 2. *Every pair of possible winners are tied for victory in some possible state. Formally, for every $b, c \in W_i(s)$, there is $s' \in S(s, r_i)$ s.t. $b, c \in W_0(s')$.*

Proof. Consider some $b, c \in W_i(s)$. Let the score of all candidates except b, c be $s'(a) = s(a)/(1 + r_i)$, and set $s'(b) = s'(c) = \min\{s(b), s(c)\}(1 + r_i)$. Then $s' \in S(s, r_i)$, and $s'(b) = s'(c) \geq s'(a)$ for all $a \in M$. \square

In the example depicted in Fig. 1, A is a possible winner for any r . B is a possible winner for voters for which $0.4(1 + r_i) \geq 0.45/(1 + r_i)$, which means $r_i \geq 0.06$. Only voters whose uncertainty r_i is above 0.73 consider C as a possible winner, as otherwise $0.15(1 + r_i) < 0.45/(1 + r_i)$.

We denote by $a_i \xrightarrow{i} a'_i$ valid local dominance steps where $a'_i \in g_i(\mathbf{a})$ and $a'_i \neq a_i$. The next two lemmas characterize such LD moves.

Lemma 3. *Consider an LD move $a_i \xrightarrow{i} a'_i$. Then either (a) $a_i \notin W_i(s)$; or (b) $a_i \prec_i b$ for all $b \in W_i(s)$; or (c) $r_i = 0$, $\{a_i, a'_i\} \subseteq W_0(s)$ and $a'_i \succ_i a_i$.*

Proof. Suppose that $a_i, b \in W_i(\mathbf{a})$, and $a_i \succ_i b$ (i.e., (a) and (b) are violated). Assume first that $a'_i \notin W_0(s)$. By Lemma 2 there is a state $s' \in S(s, r_i)$ where a_i, b have maximal score (possibly with other candidates), strictly above a'_i . W.l.o.g. $Q_{s'}^+(b) < Q_{s'}^+(a_i)$, as the tie-breaker does not affect the distance. Thus $f(s', a_i) = a_i, f(s', a'_i) = b$. Since $a_i \succ_i b$, then we have that a_i $S(\mathbf{a}, r_i)$ -beats a'_i . The remaining case is where $a'_i \in W_0(s)$ and $a_i \succ_i a'_i$. Then in the state s' where a_i, a'_i are tied it is better to vote for a_i . In either case we get that $a'_i \notin D$, which is a contradiction. \square

Lemma 4. *Consider an LD move $a_i \xrightarrow{i} a'_i$. Then*

1. $a'_i \in W_i(s)$, and there is some $c \in W_i(s)$ s.t. $a'_i \succ_i c$.
2. For a *strict LD move*, $a'_i = \operatorname{argmin}_{c \in W_i(\mathbf{a})} Q_i(c)$.
3. If $a_i \notin W_i(s)$, then $|g_i(s)| = 1$.

Part 1 is the only thing we need for our later convergence results, and its proof is immediate: Consider $a_i^* = \operatorname{argmin}_{c \in W_i(\mathbf{a})} Q_i(c)$. Clearly a_i^* locally dominates any candidate not in $W_i(\mathbf{a})$, thus $a'_i \in W_i(s)$. If there is only one possible winner, no action dominates any other action.

Beyond the fact that the above lemmas are required for our main result, they show that a strict LD move boils down to a simple heuristics: vote for the most preferred candidate among those whose score is above the threshold. For a weak LD move, any candidate above the threshold except the least preferred can be selected.

Existence of equilibrium and convergence

Suppose that voters start at some arbitrary state \mathbf{a}^0 , and play repeatedly. We get a sequence of states $\mathbf{a}^0, \mathbf{a}^1, \dots$ where in every iteration some arbitrary subset of voters may change their vote. That is, for any $t \in \mathbb{N}$ and any $i \in I$, either $a_i^{t+1} \in g_i(\mathbf{a}^t)$, or $a_i^{t+1} = a_i^t$. Our primary result states that voters always converge in the nonatomic model.⁵

⁵We assume strict preference orders for consistency with most of the social-choice literature. However we note that the theorem

Theorem 5. *Any sequence of weak LD moves is finite.*

Proof. We only need to show that no valid sequence may contain a cycle. Assume, toward a contradiction, that there is a cyclic path $(\mathbf{a}^t)_{t=0}^T$, and denote by $R \subseteq M$ all candidates that are part of the cycle. Let s^* be the lowest score of any candidate in R during the cycle, w.l.o.g. candidate $a^* \in R$ at time t^* . Thus $s^* = s^{t^*}(a^*) \leq s^t(c)$ for all $c \in R$ for every time $t \leq T$. Consider the next step where some voters join a^* (w.l.o.g. at step t^*), and pick an arbitrary voter $j \in I$ s.t. $a^* = a_j^{t^*+1} \neq a_j^{t^*}$. Thus at step t^* there is a move $a_j \xrightarrow{j} a^*$ where $a_j = a_j^{t^*+1} \in R$.

By Lemma 4, a^* is a possible winner for voter j , i.e., $a^* \in W_j(s^{t^*})$. Since $s^{t^*}(c) \geq s^{t^*}(a^*)$ for all $c \in R$, we have $R \subseteq W_j(\mathbf{a}^{t^*})$, and in particular $a_j \in W_j(\mathbf{a}^{t^*})$.

By Lemma 3, either a_j is the least-preferred candidate for j in $W_j(s^{t^*})$ (Cases I and II below), or the third category of the lemma holds. We treat the latter case separately (Case III), so assume a_j is indeed the least-preferred in $W_j(s^{t^*})$. Since $R \subseteq W_j(\mathbf{a}^{t^*})$, a_j is the least-preferred in R as well.

There must be some step t^{**} in the cycle where a voter of type v_j moves to a_j (w.l.o.g. voter j). So in $s^{t^{**}}$, a_j is preferred by j to some other possible winner z by Lemma 4. Since a_j is the least-preferred in R , and $a_j \succ_j z$, we have that $z \in W_j(s^{t^{**}}) \setminus R$. Denote the (fixed) score of z by $s(z)$.

Case I: $s(z) \geq s^*$. Consider again step t^* . Since $s(z) \geq s^*$, we have $z \in W_j(s^{t^*})$. Since a_j is the least preferred possible winner in t^* , we have that $z \succ_j a_j$, which is a contradiction.

Case II: $s(z) < s^*$. Denote $d = a_j^{t^{**}}$, and consider the step $d \xrightarrow{j} a_j$ at time t^{**} . Since $d \in R$ then $s^{t^{**}}(d) \geq s^* > s(z)$, and thus $d \in W_j(s^{t^{**}})$. By Lemma 3 (category (b) or (c)), we have that $a_j \succ_j d$. This is a contradiction since a_j is the least-preferred in R .

Case III: The remaining case is when there is $b \in W_j(s^{t^*})$ s.t. $a_j \succ_j b$. Then by Lemma 3, $a^* \in W_0(s^{t^*})$. However since a^* has minimal score, this means that $R \subseteq W_0(s^{t^*})$, i.e., all candidates in the cycle have the same score s^* at time t^* . Then all of R must have the same score at every time t , since if the score of some candidates goes up, the score of others must go down below the minimum s^* . This means that all of the moves in the cycle fall under categories (b) or (c) of Lemma 3. Thus voters only vote for more preferred candidates, which contradicts a cycle. \square

One can argue that the uncertainty level of a voter may not remain the same throughout the game. For example, there may be less uncertainty as the game advances. We note that Theorem 5 holds even if $W_j(s^{t^*})$ and $W_j(s^{t^{**}})$ are obtained via different values of r_j .

The above theorem proves convergence under very broad conditions, but does not provide much intuition as to what

applies also under weak preferences. Indeed, introducing indifference in the preference relation only eliminate LD moves, and therefore never creates a cycle.

happens along the converging path. Our next result shows that when all voters have the same uncertainty level, convergence is much more structured. Our result extends a similar result in (Meir, Lev, and Rosenschein 2014) for the finite model, but note that in our model convergence is guaranteed even when subsets of voters move simultaneously.

Conveniently, when all voters has the same uncertainty r , at any state \mathbf{a}^t there is just one agreed set of possible winners, denoted by $W^t = W(\mathbf{a}^t)$. We say that a move $a_i \xrightarrow{i} a'_i$ is an *opportunity move* if $a'_i \succ_i a_i$, and otherwise it is a *compromise move*.

Proposition 6. *Consider any non-atomic Plurality voting game with weak LD moves, where all voters have the same uncertainty level r . If \mathbf{a}^0 is the truthful state, then for all t : (A) $W^{t+1} \subseteq W^t$; (B) the score of the winner is non-decreasing; (C) there are only compromise moves; and (D) any a_i^t is either the most preferred candidate for i in W^t , or it is not in W^t .*

The proposition still holds if the uncertainty level r decreases over time.

Regret Minimization

The local dominance approach is appropriate to describe voters who are reluctant to change their vote, unless they know it cannot hurt them, a behavior that is consistent with loss aversion. A different approach to decision making under strict uncertainty is minimization of *worst case regret*, explained by *risk aversion* (Bell 1982). Intuitively, such a voter wants to avoid situations where she could have gained much, but failed to vote for the right candidate.

On one hand, regret minimization is simpler than local dominance since it does not depend on the current vote. Thus we can write $g_v(\mathbf{s})$, $W_v(\mathbf{s})$ rather than $g_i(\mathbf{s})$, $W_i(\mathbf{s})$. On the other hand, regret may depend on *cardinal utilities*, rather than ordinal preferences.

A cardinal utility scale is a *generic* function $u : M \rightarrow \mathbb{R}$, meaning that $u(a) \neq u(b)$ for all $a \neq b$. A cardinal utility scale u fits order $Q \in \pi(M)$, if $u(a) > u(b)$ whenever $Q(a) < Q(b)$. We thus augment the definition of a type v with a cardinal utility scale u_v that fits Q_v .

Formally, the *regret* of a type v voter for voting b in state s' is $REG_v(s', b) = \max_{c \in M} u(f(s', c)) - u(f(s', b))$. Note that $REG_v(s', b) \geq 0$.

The *worst case regret* (WCR) of a type v voter for voting b in \mathbf{s} is $WCR_v(\mathbf{s}, b) = \max_{s' \in S(\mathbf{s}, r_v)} REG_v(s', b)$.

Definition 3 (Strategic move under regret minimization). A WCR voter of type v in profile \mathbf{a} votes for a candidate b minimizing $WCR_v(\mathbf{s}_a, b)$.

We first characterize WCR moves. Under a candidate-wise distance, we can effectively ignore the cardinal utility scale, as only the preference order affects the vote.

Lemma 7. *Either $|W_v(\mathbf{s})| = 1$ (in which case all regrets are 0); or the unique candidate minimizing $WCR_v(\mathbf{s}, c)$ is $a^* = \operatorname{argmin}_{c \in W_v(\mathbf{s})} Q_i(c)$.*

Local dominance vs. Regret minimization

From Lemmas 4 and 7, we get that regret provides a partial justification for strict LD moves.

Corollary 8. *Let $a_i \xrightarrow{i} a'_i$ be a strict LD move in state s , then a'_i is the WCR response of type v_i , i.e., the unique candidate c minimizing $WCR_{v_i}(s, c)$.*

Thus this mean that the WCR dynamics and the strict LD dynamics coincide? The answer is *no*, since the entailment is only in one direction: any strict LD move is also a WCR move, but there may be an WCR move $a_i \xrightarrow{i} a'_i$ even when a_i is not locally dominated. This may occur for example when both of a_i, a'_i are possible winners. Thus under the WCR dynamics voters have less tendency to stay where they are, and therefore cycles are more likely to form.

It is an open question whether cycles exist when all voters have the same r , yet if we add a restriction on the initial state then convergence is guaranteed.

Proposition 9. *Consider any non-atomic Plurality voting game, where all voters have the same uncertainty level r . If \mathbf{a}^0 is the truthful state (and $|W_r(\mathbf{a}^0)| \geq 2$), then for every time t and any $i \in I$, the WCR move and the strict LD move of i coincide. In particular, the WCR dynamics converges to an equilibrium.*

Proof. Consider any voter $i \in I$. By property (D) of Proposition 6, we have that either $a_i^t \notin W_r(s^t)$, or a_i^t is the best possible winner. In the first case, by Lemma 4, there is a strict LD move, and by Corollary 8, this move coincides with the WCR move. In the latter case, there is no LD move for i (so $a_i^{t+1} = a_i^t$ under LD), and by Lemma 7, a_i^t minimizes worst case regret (so $a_i^{t+1} = a_i^t$ under WCR).

Finally, since the LD dynamics converges, we get that under the conditions of the proposition (same r , truthful initial state), WCR converges as well. \square

Diverse population For the local dominance dynamics, Theorem 5 shows that any game is acyclic. Since under regret minimization voters are more likely to have a strategic move, it is also more likely that cycles emerge, and even the existence of an equilibrium is not guaranteed.

Proposition 10. *There is a voting game where no voting equilibrium exists under WCR dynamics.*

Discussion and Related Work

We highlight some of the ideas and results in the paper, and how they compare with existing literature.

Uncertainty and regret Voting behavior based on regret minimization was considered by Ferejohn and Fiorina (1974). 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. Another regret-based model was suggested in (Merrill 1982), which also ignores any available information. Merrill shows that under the Plurality rule uncertain

voters should be *truthful*, which stands in sharp contrast to behavior observed in the real world.

One of the most prominent models for voting under (probabilistic) uncertainty was suggested by Myerson and Weber (1993), where voters preferences are sampled from a known prior distribution. An equilibrium according to the MW model is a mapping from preferences to a distribution over votes, such that each voter maximizes her expected utility w.r.t. this distribution. Myerson and Weber prove via a fixed-point argument that an equilibrium always exists for every positional scoring rule, and in particular for Plurality.

We see our regret minimization model as a non-probabilistic variation of the MW model. Specifically, in the MW model the voter considers the probability of each tie to conclude her expected utility. In our WCR model the voter focuses on the *most significant* possible tie, which greatly facilitates the decision making process.

Some voting experiments suggest that human voting behavior is consistent with regret minimization (Blais et al. 1995; Krueger and Acevedo 2008), though voters may not see themselves as such. We should note that these studies have limited relevance to our work since they concern the decision *whether to vote* (voter turnout), rather than the strategic decision *what to vote*. Thus more experimentation is required to test the validity of such models.

A related approach is *iterated regret minimization* (Halpern and Pass 2012), which was recently applied to voting (Emery and Wilson 2014). This approach assumes that voters know exactly both the preferences and the decision making process of the other voters, whereas both of these assumptions are avoided in the models we studied.

Local Dominance Concepts that are similar or equivalent to local dominance have been recently introduced in the voting literature, notably *II-manipulation* (Conitzer, Walsh, and Xia 2011; Reijngoud and Endriss 2012) and *'de re' manipulation* (van Ditmarsch, Lang, and Saffidine 2012). However these papers did not consider distance metrics and did not present any results on equilibrium existence or convergence.

For further background, we refer the reader to (Meir, Lev, and Rosenschein 2014), which surveyed many game-theoretic models of voting equilibrium, in particular w.r.t. various approaches to uncertainty, dominance, and iterative voting procedures. In addition, Meir et al. showed via extensive simulations that the equilibria reached by LD voters (especially with diverse uncertainty levels) reproduce patterns observed in the real world. In particular, in equilibrium most voters vote for only two prominent candidates, a phenomenon known as Duverger's Law (Duverger 1954).

Consider the assumption made in (Meir, Lev, and Rosenschein 2014), that among all candidates dominating her current action, a voter will always select the one that is most preferred. Our paper tackles this assumption in two ways. First, we show that it is not required for convergence, and can thus be relaxed (among the other restrictions we relax). Second, we show that this assumption can be justified on the grounds that it minimizes the worst case regret of the voter.

Implications for general games A recent result on memoryless dynamics in atomic games, suggests that in *any game* with more than one pure equilibrium, cycles must occur if arbitrary groups of agents can move simultaneously (Jaggard, Schapira, and Wright 2011). Our result shows that this is no longer true in nonatomic games, and in fact there is a large class of games where convergence is guaranteed despite the existence of multiple equilibria.

Conclusion and Future Directions

We showed that in the Plurality voting system, voters who avoid locally-dominated candidates will always converge to an equilibrium, and that this result is robust to the uncertainty levels in the populations, the initial state, and the order in which voters or groups of voters play.

In finite games without uncertainty, it is known that convergence must occur in polynomial time in the number of voters and candidates. The speed of convergence in nonatomic games is not easily defined. For example voters can just move in smaller and smaller masses in an infinite (acyclic) path. However one can ask if there is a bound on the maximal number of times a *single* voter can move. For example under the conditions of Prop. 6, any voter can move at most $m - 1$ times, so we can say that convergence is “fast.” Our general convergence proof does not provide a bound of such sort, but we believe it is an interesting open question for future research. We conjecture that no agent should move more than a polynomial (in m) number of times until convergence occurs.

The ultimate test for any scientific theory is in the accuracy of its predictions. An advantage of the theory studied in this paper is that it is possible to test its epistemic assumptions (distance-based uncertainty) and behavioral assumptions (local dominance/WCR minimization) almost independently. We intend to check whether empirical and experimental evidence (e.g. from (Van der Straeten et al. 2010)) support the theory, and if so, what can we say about actual uncertainty levels in real populations.

While the definition of weak LD voters naturally extends to many other voting rules (as do the definitions in (Conitzer, Walsh, and Xia 2011; Reijngoud and Endriss 2012; van Ditmarsch, Lang, and Saffidine 2012)), the action space in most rules contains all permutations of M , and there may be many actions that dominate a particular action. Thus the mere definitions are not very instructive as to how a voter would act. We hope the (local) worst-case regret minimization approach will be useful in defining reasonable voting behaviors under uncertainty for other voting rules. Finally, distance based uncertainty may prove useful in other games where there is a natural metric over action profiles.

Acknowledgments

The author would like to thank Omer Lev, David Parkes and Maria Polukarov for insightful discussions and for commenting on drafts of this paper, and Harvard CRCS for the support. Some of the future directions and clarifications were pointed out by anonymous referees.

References

- Aumann, R. J. 1999. Interactive epistemology i: knowledge. *International Journal of Game Theory* 28(3):263–300.
- Bell, D. E. 1982. Regret in decision making under uncertainty. *Operations research* 30(5):961–981.
- Blais, A.; Young, R.; Fleury, C.; and Lapp, M. 1995. Do people vote on the basis of minimax regret? *Political Research Quarterly* 48(4):827–836.
- Conitzer, V.; Walsh, T.; and Xia, L. 2011. Dominating manipulations in voting with partial information. In *AAAI’11*, 638–643.
- Duverger, M. 1954. *Political Parties: Their Organization and Activity in the Modern State*. New York: John Wiley. Y.
- Emery, M., and Wilson, M. C. 2014. Iterated regret minimization in voting games. In *COMSOC’14*.
- Ferejohn, J. A., and Fiorina, M. P. 1974. The paradox of not voting: A decision theoretic analysis. *The American political science review* 525–536.
- Halpern, J. Y., and Pass, R. 2012. Iterated regret minimization: A new solution concept. *Games and Economic Behavior* 74(1):184–207.
- Hyafil, N., and Boutilier, C. 2004. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *UAI’04*, 268–277.
- Jaggard, A. D.; Schapira, M.; and Wright, R. N. 2011. Distributed computing with adaptive heuristics. In *ITCS’11*.
- Kahneman, D., and Tversky, A. 1974. Subjective probability: A judgment of representativeness. In *The Concept of Probability in Psychological Experiments*. Springer. 25–48.
- Kahneman, D.; Knetsch, J. L.; and Thaler, R. H. 1991. Anomalies: The endowment effect, loss aversion, and status quo bias. *The journal of economic perspectives* 193–206.
- Krueger, J. I., and Acevedo, M. 2008. A game-theoretic view of voting. *Journal of Social Issues* 64(3):467–485.
- Meir, R.; Polukarov, M.; Rosenschein, J. S.; and Jennings, N. R. 2010. Convergence to equilibria in plurality voting. In *AAAI’10*, 823–828.
- Meir, R.; Lev, O.; and Rosenschein, J. S. 2014. A local-dominance theory of voting equilibria. In *ACM-EC’14*, 313–330.
- Merrill, S. 1982. Strategic voting in multicandidate elections under uncertainty and under risk. In *Power, voting, and voting power*. Springer. 179–187.
- Myerson, R. B., and Weber, R. J. 1993. A theory of voting equilibria. *The American Political Science Review* 87(1):102–114.
- Reijngoud, A., and Endriss, U. 2012. Voter response to iterated poll information. In *AAMAS’12*, 635–644.
- Savage, L. J. 1951. The theory of statistical decision. *Journal of the American Statistical association* 46(253):55–67.
- Schmeidler, D. 1973. Equilibrium points of nonatomic games. *Journal of Statistical Physics* 7(4):295–300.
- Tversky, A., and Kahneman, D. 1974. Judgment under uncertainty: Heuristics and biases. *science* 185(4157):1124–1131.
- Van der Straeten, K.; Laslier, J.-F.; Sauger, N.; and Blais, A. 2010. Strategic, sincere, and heuristic voting under four election rules: an experimental study. *Social Choice and Welfare* 35(3):435–472.
- van Ditmarsch, H.; Lang, J.; and Saffidine, A. 2012. Strategic voting and the logic of knowledge. In *AAMAS’12*, 1247–1248.