Heuristic Voting as Ordinal Dominance Strategies

Omer Lev, Reshef Meir, Svetlana Obraztsova and Maria Polukarov

Abstract

Decision making under uncertainty is a key component in many AI settings, including those that involve strategic decisions whose outcome depends on the actions of other agents. The common solution of expected utility maximization requires both a cardinal utility function and detailed probabilistic information. However, such probabilities cannot be easily obtained, especially in strategic situations. Moreover, people are notoriously bad at using probabilities even when such are available.
We present a framework that allows "shades of gray" of likelihood without probabilities, creating a hierarchy of sets of states of the world, with inner sets reflect higher likelihood. This hierarchy of likelihoods allows us to define ordinally-dominated strategies. We apply this definition to voting settings and show that it justifies various voting heuristics as bounded-rational strategies.

1 Introduction

The question of how an agent (human or artificial) chooses a strategy when faced with a choice has been a key issue in artificial intelligence since its inception. Various approaches have been presented, many of which fundamentally rely on two primary components: the *epistemic state* of the decision maker (her beliefs on how her actions will affect the world), and her *innate preferences*, i.e., the utility or cost involved in any outcome.

In voting settings, which will be our running application throughout this paper, agents' views are aggregated to reach a shared result. Participants have various strategies they can follow once they know the current state of the world—what other agents are voting—and their own utility function (in most voting settings, an ordinal preference is assumed). This voting decision may either be applied once (based on the current beliefs of the voters), or in an iterative setting where voters have several opportunities to observe the state and change their action.

When the exact state of the world is unknown, the epistemic state of the decision maker might depend on some prior knowledge and/or signals from the environment, which provide *partial information*.

The most common way to address this lack of knowledge has been to assign probabilities to each state of the world, and to assume that agents each maximize their expected utility over all possible states, and traditional voting models often take into account voters' uncertainty and expected utility (e.g., Myerson and Weber [19]). However, in many cases agents may not have the ability to find the precise probabilities of each state of the world. Moreover, it has been well established that people are quite bad in assigning probabilities and in acting according to them [21, 12]. Alternative approaches, focusing on decision making in face of *strict uncertainty* (possible/impossible states) have been formulated (e.g., [11]) and applied in various AI and economic settings [7, 3, 14], and more recently in voting [5, 20, 17].

1.1 Contribution

We extend the framework suggested in the latter papers, by allowing gradual levels of uncertainty, without falling back to a probabilistic approach. Instead, we build on the idea of having well-delineated sets of states of the world, each being a superset of the one before it, where innermost sets correspond to the most likely states. We say that an action is justified if it is not dominated by any other action, *at any level of the hierarchy.* This extends the most basic rule of rationality: *avoid dominated actions* (Aumann [1, 2]). Importantly, an agent's belief in our model is only over other agents' *actions* (in the context of voting, their ballots), and not over their preferences or beliefs. This is in contrast to Aumann and mainstream game theory, but in alignment with recent work in social choice and more traditional work in AI (see Section 1.2). In particular, there is no common belief (or even individual belief) of rationality. Equilibrium, if exists, emerges as a result of heuristic reactive reasoning rather than recursive strategic deliberation.

We then suggest an alternative representation for information structures in voting, based on what we term a *pivot-graph.* Given a hierarchy of pivot-graphs (one for every likelihood level), we can determine, for example, which votes dominate others, using only ordinal preferences. We show that many voting heuristics from the literature can be justified as rational decisions for an appropriate epistemic model (a specific hierarchy of pivot-graphs). Moreover, in an iterative voting setting, certain topological assumptions on the pivot-graphs guarantee that voters converge to equilibrium, thereby generalizing some previous convergence results.

1.2 Related Work

For an up-to-date coverage of iterative voting, heuristics and uncertainty-based models, see Meir [16]. In particular, Conitzer et al. [5] considered a voter faced with an arbitrary information set, and Reijngoud and Endriss [20] considered partial information where, for example, only the ranking of candidates' score or only the identity of the leader are known. Closest to our paper is the Local Dominance model in Meir et al. [17] and Meir [15], in which all voters base their belief on the possible states on a shared *prospective state.* It was shown that in an iterative voting setting where voters play actions that dominate (within the possible set) their current action, they are guaranteed to converge to an equilibrium under certain assumptions on the distance metric. *Voting heuristics* do not explicitly define the beliefs of the voter, and instead specify a (typically) simple function that dictates the vote in every given state, and aims to capture realistic behavior [20, 10]. In particular, some models suggested that a voter either votes truthfully [8] or abstains [6] when not pivotal.

These models stand in contrast to expected utility models. Myerson and Weber [19] examine *calculus of voting* for a large number of voters, in which a voter computes the probability that each action (vote) will be pivotal in every pairwise tie. We see our model as a way to capture the same line of reasoning of identifying the influential ties, albeit without using probabilities. A more fundamental difference with the calculus of voting approach is that the latter assumes a common knowledge of rationality and of the distribution of preferences, from which an equilibrium is derived.

2 Model

In this section, we provide general definitions for our framework. Specifically, we introduce the non-probabilistic model of *information structures* and what we term the *ordinal dominance* relation between possible actions at a given structure. We denote vectors by bold letters, e.g. $\boldsymbol{x} = (x_1, x_2, \ldots)$. The notation $\boldsymbol{x}_{-i}$ refers to all entries of $\boldsymbol{x}$ except x_i. We use terminology from social choice, but all definitions in Section 2 can be easily extended to other decision making problems.

Each agent i has a set of possible actions A^i (ballots), and there is a set of m possible outcomes, C. The outcome of an action, and thus the agent's utility, depends on the state of the world denoted by s, and the agent's own action. The state s may refer to a state of nature and/or to the actions of other agents. In the multiplayer voting games in the focus of this paper, the state $s = s_{\boldsymbol{a}_{-i}}$ contains the score of each candidate, as derived from the voting profile $\boldsymbol{a}_{-i}$ of all voters except i.

The (finite or infinite) set of all world states is denoted by S. An *outcome function* $f : S \times A^i \to C$ defines the outcome $f(s, a) \in C$ of action a taken by agent i at state s (in Section 3 we show how f captures the effect of a single voter in voting games).

Each agent i has a weak preference order $\succeq^i$ over the set $C \times A_i$.[1] We use the notation $\succ^i$ and $\sim^i$ to specify strict preference or indifference, respectively.

2.1 Information structures

An *information set* is a set of states $S' \subseteq S$. An *information structure* of agent i is a collection of information sets, denoted $\mathcal{S}^i = (S^i_j)^k_{j=1}$. We say that an information structure $\mathcal{S}^i$ is *valid* if $S^i_j \subseteq S^i_{j+1}$ for all j, i.e., they are concentric – each information set contains the sets with a lower index.

An agent does not assign probabilities to states or to information sets, but an intuitive interpretation of the model is that agent i believes any state in S^i_j to be *substantially more likely* than all states outside S^i_j. An information structure can either be shared by all agents, or be agent-specific. Let $\mathscr{S}$ denote the set of all valid information structures on S.

Example 1. *There are 1000 voters and 5 candidates –* a, b, c, d, e *– using the plurality voting system. The voters have access to a poll where votes (in % of total) are* $s = (29, 26, 22, 17, 6)$*. Voters for* a *and* b *have an information structure* $\mathcal{S}^a = (S^a_1, S^a_2)$*. The set* S^a_1 *contains the poll state, along with any states that can result changing the score of any candidate by at most 5% (e.g.,* b *may get up to 31% of the votes). The set* S^b_2 *contains all states within a range of 15% from the poll. In contrast, supporters of* e, f, g *are completely certain that the poll is accurate, and their information structure is* $\mathcal{S}^c = (S^c_1)$*, where* $S^c_1 = \{s\}$*.*

Note that the voters do not reason about the preferences of other individuals—only about their (aggregate) actions. We will return to this example after defining additional components of the model.

2.2 Ordinal dominance

Following [5, 20, 17], for any information set S^i_j and actions $a, b \in A^i$, we say that action a S^i_j*-dominates* action b (denoted $a \succ^i_j b$) if $f(s, a) \succeq^i f(s, b)$ for *all* $s \in S^i_j$ and $f(s, a) \succ^i f(s, b)$ for *at least one* $s \in S^i_j$.[2] Agent i is *indifferent* between actions a, b at S^i_j (denoted $a \sim^i_j b$) if $f(s, a) \sim^i f(s, b)$ for all $s \in S^i_j$. Note that S^i_j-dominance is a *partial order* over actions A^i (transitive, antisymmetric and irreflexive relation).

Definition 1. *Action* a ordinally dominates *action* b *(in structure* $\mathcal{S}^i = (S^i_j)_{j\in[k]}$*) if there is some* $j \in [k]$ *such that action* a S^i_j*-dominates action* b*.*

The next lemma guarantees that it is not possible that $a \succ_j b$ and $b \succ_{j'} a$ for some $j' \neq j$.

Lemma 2. *Ordinal dominance is a partial order.*

[1]We usually ignore the action a_i when looking at preferences. However some preference structures such as truth-bias or lazy-bias (see Section 4.1) require explicit preferences over an agent's own actions.

[2]Since agents care about their actions we should write $(f(s, a), a) \succeq^i (f(s, b), b)$ but we keep the notation simple.

Proof. Transitivity: Suppose action a ordinally dominates action b and b ordinally dominates action c, due to S^i_j and $S^i_{j'}$, respectively. W.l.o.g. $j' \leq j$, then $S^i_{j'} \subseteq S^i_j$. There is a state $s' \in S^i_{j'}$ where $f(s', b) \succ^i f(s', c)$, and since $s' \in S^i_{j'} \subseteq S^i_j$, we also have $f(s', a) \succeq^i f(s', b)$, and so $f(s', a) \succ^i f(s', c)$. Similarly, for any $s \in S^i_{j'}$, $f(s, a) \succeq^i f(s, b) \succeq^i f(s, c)$. Thus $a \succ^i_{j'} c$ which means that a ordinally dominates c.

Antisymmetry: Suppose action a ordinally dominates action b due to S^i_j. For every $j' \leq j$, there cannot be a state $s \in S^i_{j'} \subseteq S^i_j$ where $f(s, b) \succ^i f(s, a)$. Similarly, for any $j' > j$, there is a state $s' \in S^i_j \subseteq S^i_{j'}$ where $f(s', b) \prec^i f(s', a)$. Thus, b does not $S^i_{j'}$-dominates a. Since this is true for any j', b does not ordinally dominate a. □

Dynamics and equilibrium The ordinal dominance (OD) relation does not uniquely define what an agent will do in a given state. An agent playing action a with an information structure $\mathcal{S}$ may have several reasonable responses (just like an agent in a full information setting may play any action that increases her utility). Thus, ordinal dominance defines a mapping $OD : \mathscr{S} \times A \to 2^A$, where $OD_{\succeq^i}(\mathcal{S}, a)$ contains all actions that ordinally dominate a in $\mathcal{S}$ according to preferences $\succeq^i$, and a move from a to such an action $a' \in OD_{\succeq^i}(\mathcal{S}, a)$ is called an *ordinal dominance move*. We omit the subscript $\succeq^i$ when clear from context. We also define a stricter notion of *undominated ordinal dominance move*, where $UOD(\mathcal{S}, a)$ contains all actions a' that ordinally dominate a but are not ordinally dominated themselves. By transitivity, $OD(\mathcal{S}, a) = \emptyset \iff UOD(\mathcal{S}, a) = \emptyset$.

Definition 2. *An* ordinal dominance (OD) equilibrium *(for a given information structure $\mathcal{S}$ and preference profile $\succ$) is an action profile $\boldsymbol{a}$ such that for every agent i, $OD_{\succeq^i}(\mathcal{S}^i(s_{\boldsymbol{a}_{-i}}), a_i)$ is empty.*

That is, no agent plays an action that she believes to be ordinally dominated. Note that the definition would not change by using UOD instead of OD.

Observation 3. *For a "full information" epistemic model where $\mathcal{S}^i(s) = (\{s\})$, the set $OD(\mathcal{S}^i(s), a_i)$ coincides with the set of better responses to (s, a_i); the set $UOD(\mathcal{S}^i(s), a_i)$ coincides with best-responses; and OD equilibrium coincides with a pure strategy Nash equilibrium.*

Observation 4. *If $S_k = S$, any globally-dominated action (i.e., dominates any other action, in all possible situations) is also ordinally-dominated.*

3 Voting and Pivotal Actions

We now apply our model to voting scenarios. Specifically, we consider single-winner elections with a large population of voters [19, 15], under scoring voting rules. The set of outcomes is the set of candidates C, and the states of the world are all possible score vectors $S = \mathbb{N}^m$, where $s(c)$ is the number of votes for c. A *score-based voting (SBV) rule* f_A is an outcome function defined by a set $A \subseteq \mathbb{N}^m$ of allowed votes: For example, the set A under Plurality contains all vectors whose sum is 1, Approval allows all binary vectors, and Borda allows all permutations of $(0, \ldots, m-1)$. The action sets are symmetric, i.e., $A^i = A$ for all voters. We denote by $a_i(c) \in \mathbb{N}$ the absolute score given to c by agent i in vote $a_i \in A^i$. The outcome function $f_A(s, a_i)$ selects a candidate $c \in C$ with the maximal total score $s(c) + a_i(c)$, breaking ties (say) lexicographically.

3.1 Pivot graphs

A pair of actions (a', a'') is *pivotal* for a pair of outcomes $c', c'' \in C$ in state $s \in S$, if $f(s, a') = c'$ and $f(s, a'') = c''$. An agent i is pivotal for the pair of outcomes $c', c'' \in C$ in information set S^i_j, if there are $s \in S^i_j$ and actions $a'_i, a''_i \in A^i$ that are pivotal for c', c'' in s.

Information set S^i_j induces a *pivot-graph* $H^i_j = (C, E)$, which contains a vertex for every outcome, and an edge (c', c'') if agent i is pivotal for the pair c', c'' in S^i_j.

Every information structure $\mathcal{S}^i$ induces a *pivot graph structure* $\mathcal{H}^i = (H^i_j)^k_{j=1}$, where each H^i_j is a subgraph of H^i_{j+1}. The set $\mathscr{H}(C)$ contains all pivot graph structures.

Example 5. *Continuing our Example 1 from above, there is a state $s' \in S^a_1$ such that $s' = (28, 28, 19, 20, 5)$ and thus a single voter can change the outcome from $f(s', a) = a$ to $f(s', b) = b$. Thus, for supporters of a and b, H_1 will have an edge between a and b, and similarly between a and c and between b and c. Their H_2 will have these edges and additionally between d and a, b and c. For voters of e, f, g, their pivot-graphs have no edges at all.*

In elections with large populations of voters, each voter is almost insignificant, as she carries a minute effect on the outcome. Therefore, voters do not know the exact score of each candidate, but only have a rough idea of what it is (each candidate's share of the votes). As a result, it is likely that if a voter considers herself pivotal in some possible tie, she will consider *any* change in her vote as possibly pivotal. We capture this property in the following formal definition.

Definition 3 (Sharp Pivot Property (SPP))**.** *An information structure $\mathcal{S}^i$ satisfies the* Sharp Pivot Property *if: for all $c', c'' \in C$, an edge $(c', c'') \in H^i_j$ entails that* any pair of actions *$a'_i, a''_i \in A^i$ such that $a'_i(c') > a''_i(c')$*[3] *and $a'_i(c'') \leq a''_i(c'')$ is pivotal for c', c''.*

Note that structures $\mathcal{S}^i$ and $\mathcal{H}^i$ are two different ways to represent the information of an agent. In general, $\mathcal{H}^i$ contains less information than $\mathcal{S}^i$. We will assume throughout the paper that all information structures have SPP, which means that $\mathcal{H}^i$ contains all the relevant information in $\mathcal{S}^i$.

The SPP is justified when the scale of a single voter's influence is much smaller than the scale of a candidate's score. In the large-population model of Myerson and Weber [19] a similar assumption is made: either c', c'' are not tied (in which case the pivot probability is negligible), or c', c'' are tied in the expectation, in which case a difference of 0,1,2,3, or any other score (up to the maximum allowed by the voting rule to be cast by a single voter) are all equally likely. SPP makes a weaker assumption, that none of these states is substantially more likely than another (i.e., belongs in a different certainty level). For example, if $(c, d) \in H_2$, $a'_i(c) = a_i(c) + 3$ and $a'_i(d) = a_i(d)$ then there is state $s' \in S_2$ where c needs exactly 3 more votes to beat d (so a_i, a'_i are pivotal for d, c in s'), but we do not need to specify exactly what this state s' is (i.e., the exact score of each candidate).

3.2 Computing dominance relations

We show that strategies can be efficiently compared according to ordinal dominance.

Proposition 6. *Given a pivot graph structure $\mathcal{H}^i = (H^i_1, \ldots, H^i_k)$ and any SBV f, voter i can check in time $O(|C|^2 k)$ if vote $a'_i \in A$ dominates vote $a_i \in A$.*

Intuitively, Algorithm 1 checks (for each uncertainty level j), whether the new vote a'_i is "safe" (not worse than a_i in any possible tie), and whether it is "pivotal" (better than a_i

[3] For a vote $a \in \mathbb{R}^m$ and a candidate $c \in C$, we denote as $a(c)$ the value of candidate c's coordinate in a.

in at least one tie). Interestingly, the only dependence on the voting rule is that it defines which actions A are allowed.

$I[X] \in \{-1, 1\}$ is an indicator variable for statement X.

Algorithm 1: $\textsc{OrdinallyDominates}(a'_i, a_i \in A, \succeq^i, \mathcal{H}^i \in \mathscr{H})$

```
for c, c' ∈ C do
    diff(c, c') ← a'_i(c) + a_i(c') − a_i(c) − a'_i(c');
    effect(c, c') ← sign(diff(c, c') · I[c ≻^i c']);
for j ≤ k do
    safe(j) ← min_{(c,c')∈H^i_j} effect(c, c');
    pivot(j) ← max_{(c,c')∈H^i_j} effect(c, c');
    dom(j) ← I[pivot(j) + safe(j) ≥ 1];
if ∃j ≤ k s.t. dom(j) = 1 then
    return TRUE
else
    return FALSE
```

Proof. Suppose a' ordinally dominates a. Then there is some level $j \leq k$ such that $a' \succ^i_j a$. This means that for any pair of candidates (c, c') that can be tied in H^i_j, either c is preferred to c' and a' weakly reduced c''s score, or c' is preferred to c and a' weakly adds to c''s score (thus $\text{effect}(c, c') \geq 0$). Hence, in particular $\text{safe}(j) \geq 0$. In addition, there must be a pair for which the gain is strict, and $\text{effect}(c, c') = 1$, which means $\text{pivot}(j) = 1$. In total, $\text{dom}(j) \geq 1 + 0 = 1$ so the algorithm returns TRUE.

Otherwise, in every level j, either a'_i, a_i have the same outcome in all states, or there is a pair $(c, c') \in H^i_j$ such that $f(s, a_i) = c$, $f(s, a'_i) = c'$, and $c \succ^i c'$.

In the latter case, since f is a scoring rule this means that $a'_i(c) - a_i(c) < a'_i(c') - a_i(c')$, i.e. that c' gained strictly more score than c when changing from a_i to a'_i. Thus $\text{diff}(c, c') = a'_i(c) + a_i(c') - a_i(c) - a'_i(c') < 0$, and $\text{effect}(c, c') = -1$. The algorithm then computes $\text{safe}(j) = -1$. Therefore $\text{dom}(j) \leq 1 - 1 = 0$.

In the first case, $\text{effect}(c, c') = 0$ for all pairs, and thus $\text{safe}(j) = \text{pivot}(j) = 0$, and $\text{dom}(j) = 0$. □

3.3 Epistemic models

Beliefs are often derived from a current state or signal. An *epistemic model* of agent i maps any state s to an information structure $\mathcal{S}^i(s) = (S^i_1(s), S^i_2(s), \ldots, S^i_k(s))$, or directly to a pivot graph structure $\mathcal{H}^i(s)$. In particular, full information occurs when each S_i contains a single state (the "real" state – s).

An epistemic model is *cliqued* if its mapping to a pivot graph structure $\mathcal{H}(s)$ at every state $s \in S$, $H_j(s)$, is a clique. The epistemic model is *upward closed* if the pivot graph structure $\mathcal{H}(s)$ at every state s has an order L over outcomes such that if $(c, c') \in H_j(s)$ and $c'' >_L c'$ then $(c, c'') \in H_j(s)$. Note that any cliqued epistemic model is upward closed (where L may be an arbitrary order where all candidates in $H_j(s)$ precede all others). More generally, L can be roughly thought of as an order of likelihood of states. For simplicity of notation, we shall be referring to a particular pivot graph structure as the epistemic model which maps into it.

(a) The state s (poll) (b) Information structure

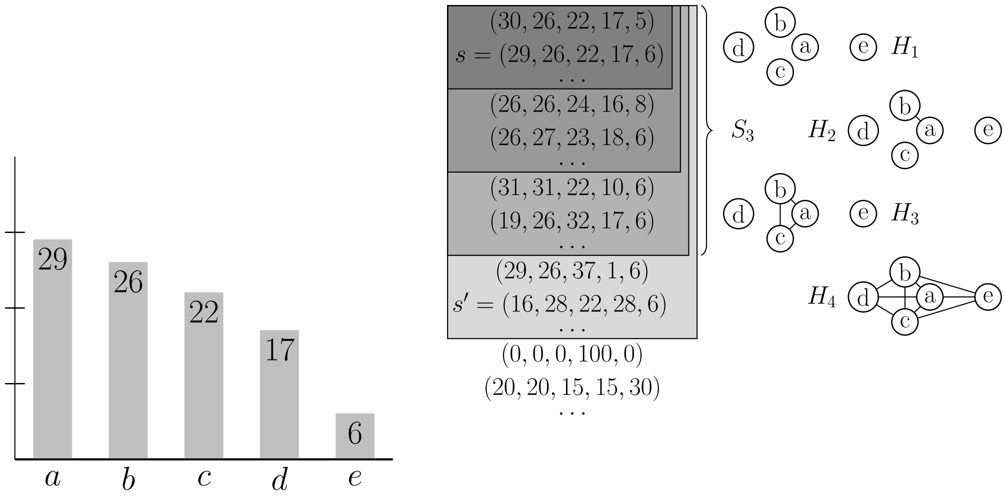

Figure 1: A schematic example of the poll s, the information structure $\mathcal{S}_{\ell_1,\boldsymbol{r}}(s)$, and the induced pivot graph structure $\mathcal{H}_{\ell_1,\boldsymbol{r}}(s)$. E.g., graph H_4 contains the edge (b,d) due to state $s' \in S_4$. States below S_4 are considered impossible. All the pivot graphs are upward closed w.r.t. the lexicographic order on C, but they are not always a clique ($(d,e) \notin H_4$).

Distance-based uncertainty Following Meir et al. [17], we consider the following way to derive information sets and information structures. Given a metric $d : S \times S \to \mathbb{R}_+$ and a parameter $r \in \mathbb{R}_+$, every state $s \in S$ explicitly defines an information set $S_{d,r}(s) = \{s' : d(s,s') \leq r\}$. In general, the metric d can be completely arbitrary and the induced set is meaningless.[4] However in many contexts there is a natural metric over states. E.g. if the states represent geographic locations, times, temperatures etc., then it is natural to assume that d is an Euclidean distance metric. Note that "being in the same information set" is not a transitive relation, in contrast to the common partition model. E.g., for two scores $s = 11$ and $s' = 17$, we get $S_{d,4}(s) = [7,15]$ and $S_{d,4}(s) = [13,21]$ which partially overlap.

Distance-based uncertainty is particularly natural in games with many players, where states correspond to action profiles and the distance $d(s,s')$ may reflect how many players change their action between s and s'. Such uncertainty model was applied in Meir et al. [17] and Meir [15] to voting scenarios, where the distance between states (candidate score vectors) were defined by various ℓ-norms and the Earth Mover distance (EMD).[5] Thus, $S_{d,r}(s)$ may reflect a range of possible candidates' scores given a poll or a current state s. It was also applied to routing games in Meir and Parkes [18].

Distance metrics provide us with a simple way to define an information structure: given a metric d and an increasing sequence of distances $\boldsymbol{r} = (r_1, r_2, \ldots, r_k)$, we get an epistemic model where $\mathcal{S}_{d,\boldsymbol{r}}(s) = (S_{d,r_1}(s), S_{d,r_2}(s), \ldots, S_{d,r_k}(s))$. We will call such an information structure *concentric*.

Example 7. *We expand Example 1 where candidates' scores (in % of total) are* $(29, 26, 22, 17, 6)$. *We do not specify the number of voters in the poll. We consider a voter with a concentric information structure, based on the radii* $\boldsymbol{r} = \{1\%, 3\%, 7\%, 20\%\}$ *and the EMD metric. These information sets induce pivot-graphs as illustrated in Figure 1.*

[4] In fact, *any* set $S' \subseteq S$ can be derived from s for some carefully designed metric d.

[5] Under the ℓ_p norm, $d(s,s') = \left(\sum_{c \in C} |s(c) - s(c')|^p\right)^{1/p}$. The EMD is essentially the ℓ_1 norm, with the additional constraint that the total number of votes remains the same.

In Example 7, consider a Plurality voter whose preferences are $e \succ d \succ c \succ b \succ a$. Then the action "$c$" (a shorthand for $(0,0,1,0,0)$) ordinally dominates action "e" (due to H_3) and "b" ordinally dominates everything else due to H_2.

To recap, our model differs from standard (strict uncertainty) epistemic game theory models in two important ways: it does not (necessarily) assume that beliefs form a partition of the state space; and agents do not attribute explicit beliefs, preferences, or rationality to other agents. Rather, our agents each act as bounded rational individual decision makers, reacting to an uncertain but static world.

4 Justifying Voting Heuristics with Ordinal Dominance

Many heuristics have been suggested to analyze how voters behave and change their vote. Most heuristics are derived from a single "prospective state" s, which is assumed to be the state where the voter is (or some proxy for this state). Formally, a *set heuristics* is a function $h : S \times A \to 2^A$ that maps the prospective state and the current action to a set of new possible actions. We say that h is a *point heuristics* if $|h(s,a)| \leq 1$ for every s, a. To be consistent with previous definitions, we always omit a from the set $h(s,a)$, and assume that when $h(s,a) = \emptyset$ the voter simply keeps her current vote.

Definition 4. *We say that an epistemic model* justifies *heuristic h, if for any state $s \in S$ and current action $a \in A^i$: (I) $h(s,a) = \emptyset$ if and only if $UOD(\mathcal{S}(s), a) = \emptyset$; and (II) $h(s,a) \subseteq UOD(\mathcal{S}(s), a)$.*

This means that the heuristic only recommend undominated ordinal-dominance moves under the epistemic model, and only keeps the current action if no such move exists (we could also think of a weaker justification with OD moves).

As a simple example, consider the Plurality rule f (where the action set A coincides with C), and the heuristic $h^{not-last}(s,a)$ that is empty except when action a is the least preferred candidate $\hat{a}_i$, and then it moves to an arbitrary other candidate (regardless of s). The epistemic model which maps to the pivot graph structure $\mathcal{H}^{all}(s) = (H_1^{all})$ where H_1^{all} is the complete graph justifies $h^{not-last}$ as follows: (I) suppose that $a \neq \hat{a}_i$. Then no candidate ordinally dominates a and thus $UOD(\mathcal{H}^{all}(s), a) = \emptyset = h^{not-last}(s,a)$. (II) when $a = \hat{a}_i$, any other candidate is undominated but globally dominates a (since there is a possible state where i is pivotal for c against a), in which case $UOD(\mathcal{H}^{all}(s), a) = C \setminus \{\hat{a}_i\} = h^{not-last}(s,a)$.

4.1 Local dominance

Local dominance [17] heuristic with metric d and parameter r explicitly define a set $S_{d,r}(s) = \{s' : d(s,s') \leq r\}$. The heuristic action $h^{LD}_{d,r}(s, a_i)$ is defined for the Plurality rule as follows: Let $D \subseteq C$ be the set of candidates that $S_{d,r}(s)$-dominate a_i; If D is non-empty, then vote for the most preferred candidate in D.

To justify $h^{LD}_{d,r}$ with ordinal dominance, we define an epistemic model $\mathcal{H}^{LD}_{d,r}$ where $\mathcal{H}^{LD}(s)_{d,r}$ contains a single pivot graph H_1 which is the pivot graph induced by $S_{d,r}(s)$. Note that our definition applies for *any voting rule*, unlike the one in Meir et al. [17]. In Plurality, $\mathcal{H}^{LD}_{d,r}$ justifies $h^{LD}_{d,r}$ (straightforward proof omitted due to space constraints).

Truth/lazy-bias Denote the top candidate of i by $q_i \in C$, and denote by $\perp$ an "abstain" action that adds no score to candidates. We adopt the suggested variations in Dutta and Laslier [8] and Desmedt and Elkind [6], where the voter prefers the truthful/abstain action if this does no affect the outcome. However, this naïve modification alone may lead to unreasonable behaviors, e.g., where no-one votes [9], even under full information. This issue was handled in Meir et al. [17] by defining an explicit heuristic rule with two distances. In the remainder of this subsection we restrict out attention to Plurality, where $A = C$ for truth-biased agents, and $A = C \cup \{\perp\}$ for lazy-biased agents.[6]

The "truth bias" heuristics $h^{LD+TB}_{d,r_1,r_2}(s, a_i)$ is as follows [17]: (1) perform a local-dominance move at radius r_1, if exists. If such move does not exist, i examines if $f(s', a_i) \succ^i f(s', q_i)$ for some $s' \in S_{d,r_2}(s)$. (2a) If so, agent i keeps the current vote a_i, (2b) otherwise, i moves to q_i.

While the behavior seems to maintain the reason behind truth bias, the definition of h is cumbersome. Instead, we can use r_1, r_2 to define an epistemic model $\mathcal{H}^{LD+TB}_{d,r_1,r_2}$ as follows. We let $H_1(s)$ be as in $\mathcal{H}^{LD}_{d,r_1}$ above. We similarly compute a graph H' from $S_{d,r_2}(s)$. Then, we define H'' to be a *subgraph* of H' containing only edges between a_i and candidates less preferred than a_i, and set $H_2(s) = H_1(s) \cup H''$. Let $\mathcal{H} = \mathcal{H}^{LD+TB}_{d,r_1,r_2}(s) = (H_1(s), H_2(s))$.

Proposition 8. *$\mathcal{H}^{LD+TB}_{d,r_1,r_2}$ justifies h^{LD+TB}_{d,r_1,r_2} in Plurality.*

Proof. First, if $H_1(s)$ is nonempty (at least one tie) then $h^{LD+TB}_{d,r_1,r_2}(s, a_i) = a^*_i \in UOD(\mathcal{H}(s), a_i)$ as in a standard LD move. Otherwise, there are two cases.

If $H_2(s)$ contains some edge (a_i, b), then by SPP for any $a' \neq a_i$ there is a state $s' \in S_{d,r_2}(s)$ where $f(s', a_i) = a_i$ and $f(s', a') = b$ (think of s' as state where a single additional vote for a_i is critical). Since a_i is preferred to b by the definition of $H_2(s)$, we conclude that no candidate a' dominates a_i in $H_2(s)$ (thus $UOD(\mathcal{H}(s), a_i) = \emptyset$); and that $f(s', a_i) = a_i \succ^i b = f(s', q_i)$ (and thus $h^{LD+TB}_{d,r_1,r_2}(s, a_i) = \emptyset$).

In the second case, there is no such edge, then $H_2(s)$ is empty. This means that no action of i can change the outcome whatsoever, and thus by the slight truth-bias q_i is strictly preferred to any other action. In particular, it ordinally dominates a_i and is undominated so $UOD(\mathcal{H}(s), a_i) = \{q_i\}$. Finally, since i is non-pivotal then in particular there is no state $s' \in S_{d,r_2}(s)$ such that $f(s', a_i) \succ^i f(s', q_i)$. Thus $h^{LD+TB}_{d,r_1,r_2}(s, a_i) = q_i \in UOD(\mathcal{H}(s), a_i)$, as required. □

The statement for lazy-bias is similar, and uses the same information structure but with a slight preference to abstain instead of voting truthfully.

4.2 T-pragmatist

The T-pragmatist (point) heuristic [4, 20] considers the leading T candidates in s (denoted $\mathbf{T}$), and sets a new action $a' = h^{T-prag}(s, a_i)$ where a' is identical to a_i except the favorite candidate in $\mathbf{T}$ is moved to the top. E.g., if $a_i = (b_1, b_2, b_3, b_4, b_5)$ is the truthful order, and the voter state s is such that the score order is $(b_3, b_4, b_2, b_1, b_5)$, then $h^{2-prag}(s, a_i) = (b_3, b_1, b_2, b_4, b_5)$ and $h^{3-prag}(s, a_i) = (b_2, b_1, b_3, b_4, b_5)$.

Consider an epistemic model that creates for each agent i $\mathcal{H}^{T,i-star}(s)$, which creates a star graph, in which the center node is the most preferred candidate by voter i in the top T candidates, and it is tied with all other $T-1$ candidates in the top T.

Proposition 9. *$\mathcal{H}^{T,i-star}$ justifies h^{T-prag} for all s and a_i in Plurality (or other rules). That is, $h^{T-prag}(s, a_i) = UOD(\mathcal{H}^{T,i-star}(s), a_i)$.*

Proof. For $T = 1$, the graph has no edges, which means the heuristic advises doing nothing as well. Otherwise, the graph contains $T-1$ edges, and the only OD action is to vote for voter i's favorite candidate in the top T, which is a UOD. This is exactly the recommendation of the heuristic. □

[6] Formally, for truth-biased agents $(f(s, a), a) \succ^i (f(s, b), b)$ either if $f(s, a) \succ^i f(s, b)$, or if $f(s, a) = f(s, b)$ and $a = q_i, b \neq a$. Similarly for lazy-biased agents when $a = \perp$.

4.3 Leader Rule (Approval voting)

Assume candidates $c_1, \ldots, c_m$ are sorted in decreasing score order in a state s. In Approval voting the allowed actions are $A = 2^C$. The *Leader rule* $a' = h^{LR}(s, a_i)$ is a strategy approving all candidates strictly preferred to the leader of s, and approves the leader of s (candidate c_1) if and only if it is preferred to the runner-up c_2 (i.e., exactly one of c_1, c_2 is being approved in a' [13]).

We consider the epistemic model where $\mathcal{H}^{LR}(s)$ of two nested pivot graphs. The inner graph H_1 contains a single edge between c_1 and c_2. The outer graph H_2 is a star connecting c_1 to all candidates.

Proposition 10. *$a' = h^{LR}(s, a_i)$ ordinally dominates all other actions according to $\mathcal{H}^{LR}$. In particular, $\mathcal{H}^{LR}$ justifies h^{LR}.*

Proof. Let a'' be any alternative vote to a'. We will show that a' dominates a'' in at least one of the tie graphs H_1 or H_2.

Consider a'' that differs from a' on either c_1 or c_2 or both, as well as on any other set of candidates. On the graph H_1, the voter is pivotal for c_1, c_2 and thus there is a state s where $f(s, a'') = c_2 \prec^i c_1 = f(s, a')$, or $f(s, a'') = c_1 \prec^i c_2 = f(s, a')$. Thus a' dominates a'' on H_1.

Next, consider a'' that approves c_1, c_2 iff a' approves them, but differs in (at least) some other candidate c'. If $c_1 \succ c'$, c' is not approved in a' and thus approved in a'' (this is regardless of whether c_1 is approved). Since there is a state s in H_2 where c_1 and c' are tied, $f(s, a'') = c' \prec c_1 = f(s, a')$. If $c_1 \prec c'$, c' is approved in a' but not in a''. Again, since there is a state s where they are tied, $f(s, a') = c' \prec c_1 = f(s, a'')$. Thus $a' \succ^i_2 a''$ and therefore a' ordinally dominates a''.[7] □

5 Ordinal-domination and Iterative Voting

Once the ordinal-domination dynamic is defined on voting settings, it is natural to examine how it behaves when multiple strategic agents apply it, and would it reach a OD equilibrium or cycle. This is iterative voting, in which proceeding from some initial state s^0, and in each iteration an arbitrary subset of voters change their votes. Our convergence results depend on the structure of the pivot graphs in the epistemic model.

We first show that both cliqued and upward-closed epistemic structures are the result distance-based uncertainty with natural assumptions on the distance function.

Proposition 11. 1. *Any neutral distance metric d on scoring vectors induces an upward-closed epistemic model.*

2. *Any candidate-wise distance metric[8] d on scoring vectors induces a cliqued epistemic model.*

Proof sketch. Proof of 1: Assume that there is a state $s = (s_1, \ldots, s_m)$ in which there are $c_1, c_2, c_3 \in C$ such that $s_1 \geq s_2 \geq s_3$; and another state s' within a distance r from s where c_2, c_3 are tied. We construct a (non-normalized) vector s'' where c_1, c_3 are tied, such that $|s''_j - s_j| \leq |s'_j - s_j|$ for all j (hence s'' is closer to s than s') or one where s'' is such that $s''_j = s'_j$ for $j > 2$ and $|s''_1 - s_1| \leq |s'_2 - s_2|$ and $|s''_2 - s_2| \leq |s'_1 - s_1|$.

[7] Note that a' does not dominate the first type of alternative actions on H^i_2, so a single star graph would not have sufficed.

[8] This is a metric on a scoring vector, composed of a singleton metric $D : \mathbb{R}^2 \to \mathbb{R}_+$, where $d(s, s') = \max_{c \in C} D(s(c), s'(c))$. This includes, for example, the ℓ_∞ norm.

W.l.o.g we have $s'_3 \geq s_3$ and $s'_j \leq s_j$ for all $j \neq 2, 3,$. Denote by $w = s'_2 = s'_3$ be the winning score in s'. There are several cases: (I) if $w \geq s_1$, define $s''_1 = s_1, s''_2 = s_2 \leq s''_1, s''_3 = s_1$; (II) if $s_2 < w < s_1$, define $s''_1 = w \in [s'_1 + 1, s_1], s''_2 = s_2, s''_3 = w$. It is easy to check that s'' holds both conditions, thus $d(\frac{s''}{\|s''\|}, s) \leq d(s', s) \leq r$ as required. If (III) $w \leq s_2$, it quite simple to see that by setting $s''_1 = w, s''_2 = w - 1, s''_3 = w$ we are closer to s than s'.

Proof of 2 is similar to Meir [15], Lemma 2. □

Theorem 12. *Suppose agents each have a concentric, cliqued epistemic model (not necessarily the same one). Iterative voting using plurality must converge to OD equilibrium, from any initial state.*

Proof sketch. The proof is a mild modification of the convergence proof from Meir [15]. The key observation is that any ordinal dominance move in $\mathcal{H}^i$ is a local dominance move in *some level* H^i_j. The second observation is that by Prop. 11, $S_{d,r}(s)$ with the ℓ_∞ norm (which is the information set used in Meir [15]) induces the cliqued graph $H_{d,r}(s)$, with edges among all candidates who are possible winners in $S_{d,r}(s)$. Thus $H^i_j(s)$ is equivalent to $H_{d,r}(s)$ for some appropriate $r = r_j$.

Intuitively, convergence is as follows: Assume towards a contradiction that a cycle exists, and consider state s' when a voter i moves from a_i to the lowest-score candidate in the cycle (say, z). We then look at state s'' in the cycle when i had moved to a_i, and show that there are no r', r'' such that z is a local dominance move in $S_{d,r'}(s')$, *and* a_i is a local dominance move in $S_{d,r''}(s'')$. □

Theorem 13. *Suppose agents each have a concentric, cliqued epistemic model (not necessarily the same one). Iterative voting using veto must converge to OD equilibrium, from any initial state.*

Proof. Assume, for contradiction, that the process does not converge. Let R be the set of candidates whose score changes an infinite number of times, and let $z \in R$ be the candidate which has the lowest score in the cycle (breaking ties using the tie-breaking rule), and let $\boldsymbol{s}^q$ be the state where it reaches this abysmal score. That is, some voter j moves from vetoing candidate a to vetoing candidate z. Candidate a's (and any other $c \in R$) score is above z's, as otherwise its own vetoing before would give it a lower score than z. Since this is a cliqued epistemic model, leaving a means it is the favorite candidate of voter j over all candidates with scores above z, in particular, for any $c \in R$, $a \succ^j c$.

At some point in the future $\boldsymbol{s}^{q'}$, due to the cycle, voter j will move from vetoing some candidate $b \in C$ to veto a, due to an edge in its relevant pivot-graph, indicating a tie between a and some other candidate x. If x's score at $\boldsymbol{s}^q$ was higher than z, then we know a is preferred over it from z's vetoing. If x's score was lower, we know it hasn't changed (as it isn't in R), meaning b is still tied with a as well in the pivot-graph of $\boldsymbol{s}^{q'}$ as it was in $\boldsymbol{s}^q$, hence voter j will not move (since $a \succ^j b$). □

This result also implies the first non-plurality result for local dominance.

Corollary 14. *Using any candidate-wise metric, local-dominance converges to an equilibrium when using veto.*

6 Discussion

This paper presents a framework to model games in which players do not have perfect information of the world. Moreover, they do not even have an exact understanding of their uncertainty of the world's state. Hence, their understanding is modeled in a coarser way –

as "shades of likelihood" of various states of the world around them. Such a model is robust enough to capture many previously suggested heuristics and strategies of voter behavior under uncertainty.

Indeed, the use of the pivot-graph and its topological properties to show convergence (or lack of it), opens the question of whether we can discuss issues of convergence in terms of graph structures (and the metrics or properties that induce them). Perhaps a wider variety of voting rules can converge in iterative voting under different pivot-graph structures. The fact that ordinal dominance in a large population voting scenario can be computed efficiently, stands in contrast to the negative results in Conitzer et al. [5], where verifying whether vote a' dominates a is NP-hard under the Borda rule. This is due to our simplifying assumption on the sharp pivot property that allows us to replace (arbitrarily complicated) information sets with a simple pivot graph representation.

A natural and important use of our model is to reformulate heuristics from various game-theoretic domains – not limited to social choice – as ordinally-dominant strategies. This might offer an insight into the built-in assumptions inherent in these heuristics, and allowing, perhaps, novel formulations of new heuristics and methods, tailored to particular uncertainty structures.

References

[1] Robert J. Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 6(1):6–19, 1995.

[2] Robert J. Aumann. Interactive epistemology I: Knowledge. *International Journal of Game Theory*, 28(3):263–300, August 1999.

[3] Craig Boutilier. Toward a logic for qualitative decision theory. *KR*, 94:75–86, 1994.

[4] Steven J Brams and Peter C Fishburn. Approval voting. *American Political Science Review*, 72(3):831–847, 1978.

[5] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Proceedings of the 25th National Conference on Artificial Intelligence (AAAI)*, pages 638–643, San Francisco, California, August 2011.

[6] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic Commerce (EC)*, pages 347–356, Cambridge, Massachusetts, June 2010.

[7] James Dow and Sérgio Ribeiro da Costa Werlang. Nash equilibrium under knightian uncertainty: breaking down backward induction. *Journal of Economic Theory*, 64(2):305–324, 1994.

[8] Bhaskar Dutta and Jean-François Laslier. Costless honesty in voting. in 10th International Meeting of the Society for Social Choice and Welfare, Moscow, 2010.

[9] Edith Elkind, Evangelos Markakis, Svetlana Obraztsova, and Piotr Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. In *International Symposium on Algorithmic Game Theory*, pages 110–122. Springer, 2015.

[10] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Proceedings of 3rd International Conference of Algorithmic Decision Theory (ADT)*, pages 181–192, Brussels, Belguim, November 2013.

[11] Joseph Y. Halpern. Defining relative likelihood in partially-ordered preferential structures. *Journal of Articial Intelligence Research*, 7:1–24, July 1997.

[12] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, March 1979.

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

[14] Paul-Amaury Matt, Francesca Toni, and Juan R Vaccari. Dominant decisions by argumentation agents. In *International Workshop on Argumentation in Multi-Agent Systems*, pages 42–59. Springer, 2009.

[15] Reshef Meir. Plurality voting under uncertainty. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, pages 2103–2109, Austin, Texas, January 2015.

[16] Reshef Meir. Iterative voting. In Ulle Endriss, editor, *Trends in Computational Social Choice*, chapter 4, pages 69–86. AI Access, 2017.

[17] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *Proceedings of the 15th ACM conference on Economics and Computation (EC)*, pages 313–330, Palo Alto, Califronia, June 2014.

[18] Reshef Meir and David Parkes. Playing the wrong game: Smoothness bounds for congestion games with behavioral biases. In *Proceedings of the 10th Workshop on the Economics of Networks, Systems and Computation (NetEcon)*, pages 67–70, Portland, Oregon, June 2015.

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

[20] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, volume 2, pages 635–644, Valencia, Spain, June 2012.

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

Omer Lev
Ben-Gurion University
Beersheba, Israel
Email: `omerlev@bgu.ac.il`

Reshef Meir
Technion
Haifa, Israel
Email: `reshefm@ie.technion.ac.il`

Svetlana Obraztsova
Nanyang Technological University
Singapore
Email: `lana@ntu.edu.sg`

Maria Polukarov
King's College London
London, United Kingdom
Email: `maria.polukarov@kcl.ac.uk`