Convergence of Iterative Multi-Issue Voting under Uncertainty

Joshua Kavner
Rensselaer Polytechnic Institute
Troy, New York, United States of America
kavnej@rpi.edu

Reshef Meir
Technion Israel Institute of Technology
Haifa, Israel
reshefm@ie.technion.ac.il

Francesca Rossi
IBM T.J. Watson Research Center
Yorktown Heights, New York, United States of America
Francesca.Rossi2@ibm.com

Lirong Xia
Rensselaer Polytechnic Institute
Troy, New York, United States of America
xial@cs.rpi.edu

Abstract

We study the effect of strategic behavior in iterative voting for multiple issues under uncertainty. We synthesize binary multi-issue voting with Meir et al.'s local dominance theory for iterative voting and determine its convergence properties. While restricted local dominance improvement dynamics may fail to converge, with or without uncertainty, we how that cycles are an borderline-case that require exact conditions. We prove that if voters are restricted to O-legal preferences, or are part of a very large (nonatomic) population, then the game is guaranteed to converge for any level of uncertainty.

Keywords

Computational Social Choice; Equilibrium Computation; Uncertainty in AI

1 Introduction

Consider a wedding planner that is determining a wedding's main meal and wants to incorporate the party's invitees' preferences into their decision. Suppose they need to decide on three issues that have two options each: the main course (chicken or beef), the paired wine (red or white), and the dessert cake's flavor (chocolate or vanilla). How should the wedding planner incorporate the large party's preferences? This question lies in the realm of voting in multi-issue domains [17] which focuses on studying *social choice rules* that efficiently aggregate agents' preferences when multiple group decisions have to be coordinated. On the one hand, the planner could ask for a lot of information from each attendee (agent) and request their full preference order over the 2^p possibilities (alternatives), for p binary issues. This, however, has an exponential computational cost and a high cognitive cost for each agent, as it requires them to report their full preference order over many alternatives.

On the other hand, the planner could only ask for each agent's top ranked alternative and decide each of the p issues independently of one another. Although this option is computationally and cognitively simple, it may result in a *multiple election paradox*, where agents can select each of their least favored alternatives by voting for each issue independently. For example, consider our decision problem with three binary issues, and suppose three agents prefer $(1, 1, 0)$, $(1, 0, 1)$, and $(0, 1, 1)$ first respectively and all prefer $(1, 1, 1)$ last [15]. By using the majority rule on each issue independently, the agents collectively vote for $(1, 1, 1)$, their least favored outcome.

A middle-of-the-road approach is for the planner to elicit agents' most preferred alternative and, given information about every other agents' vote, allow agents to *deliberate* and exchange their votes during a fixed time period. In the single-issue setting, with multiple options in the issue's domain, this type of *iterative voting* was first explored in depth by Meir et al. [22]. By the celebrated Gibbard-Satterthwaite theorem, most social choice rules are susceptible to strategic manipulation [11, 27]. This means that an agent may have be incentivized to vote against their truthful preferences in order for the social choice rule to pick a more preferred alternative overall. Meir and colleagues studied the convergence properties of the plurality voting rule when agents sequentially and myopically make *best response* improvement steps to their votes, given all other agents' current votes. Meir et al. inspired a line of research that researched the dynamics, equilibrium, and social welfare properties of other social choice rules and assumptions on agents' behavior (see Related Work). The conclusion from previous research is mixed, as Plurality and Veto have strong convergence properties, but almost any other rule that was studied admits cycles.

A related line of work considered *uncertainty in voting*, and in particular strict uncertainty and dominating actions [5, 21, 25].

In contrast to the simple single-issue setting, uncertainty in a multi-issue domain plays a double role. First (as in single-issue voting), it means that the voter is considering herself as potentially pivotal on any issue that is sufficiently close to a tie. Second—and this part is new—the voter may be uncertain regarding her own preferences on a particular issue, as there depend on undecided issues.

Our primary question is then:

Under what conditions does iterative voting for multiple issues under uncertainty converge?

1.1 Our Contribution

On the conceptual side, we extend the Local Dominance model to multiple referenda, where voting is on one issue at a time. The generalized model naturally captures both types of uncertainty discussed above.

On the technical side, we first show that iterative voting on multiple referenda may not converge, with or without uncertainty.

Appears at the 1st Workshop on Learning with Strategic Agents (LSA 2022). Held as part of the Workshops at the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), N. Bishop, M. Han, L. Tran-Thanh, H. Xu, H. Zhang (chairs), May 9–10, 2022, Online.

However, we show two possible restrictions of the model under which voting is guaranteed to converge when starting from the truth: (a) when *preferences* are restricted to O-legal preferences (i.e., there is an order on issues that restrict the possible dependencies); (b) when *dynamics* is restricted by assuming that voters are more certain about the current issue than about other issues. Furthermore, we show a strong convergence result without any restriction, in a nonatomic variation of the model, which can be thought of as the limit case of a large population.

Finally, we empirically demonstrate that adding uncertainty eliminates almost all cycles in iterative voting. We show that the length of an improvement sequence is correlated with an increase in the quality of the resulting outcome over the truthful winning alternative.

1.2 Related Work

The study of iterative voting in the single-issue setting was initiated by Meir et al. [22]. The authors found an initial convergence result for plurality for best response dynamics under deterministic tie-breaking and provided an upper bound for how quickly an improvement sequence could converge. Subsequent work demonstrated that iterative veto converges [18, 26] while any other voting rules do not [14].

Iterative voting in the multiple-issue setting was studied by Bowman et al. [4], Grandi et al. [13]. They utilize computer simulations and human experiments to show that iterative voting can eliminate multiple election paradoxes and raise agents' social welfare for the chosen outcome. However, the former work's simulations presume agents follow a learning procedure rather than best response or local dominant dynamics.

Our work primarily extends the local dominance model in the single-issue social choice setting of Meir [20], Meir et al. [21] to the multi-issue case. These works introduced agents' local dominance improvement heuristic as a risk-averse behavior using a reasonable amount of information that agents have access to. Meir et al. [21] also found broad conditions for voting equilibrium to exist according to the plurality social choice rule. Meir [20] study a nonatomic version of the local dominance model, where agents have negligible impact on the voting outcome and multiple agents update their votes simultaneously. They showed that a nonatomic model greatly simplifies the dynamics and has stronger convergence properties.

Strategic behavior in single-issue voting when agents have limited information about other agents' votes has also been studied by Conitzer et al. [6], Endriss et al. [9], Reijngoud and Endriss [25]. All three works assume that agents have access to limited amount of information about a vote profile. Conitzer et al. [6] assume agents have access to a subset of the preference relations exhibited by the vote profile, whereas Endriss et al. [9], Reijngoud and Endriss [25] both assume agents have access to a summary of the profile, such as the winner or score of each alternative. Rather than consider which score vectors are possible, as in the local dominance framework, agents reason about which vote profiles E could elicit the same information they receive. Agents can then change their vote if the outcome improves for some profile in E, and cannot worsen for all profiles in E. All three works study different voting rules' susceptibility to manipulation for various types of information functions, while the latter two works analyze convergence properties when manipulations are taken iteratively.

Strategic behavior for binary multi-issue voting was studied by Lang and Xia [16], Xia et al. [28]; see Lang and Xia [17] for a survey.

Different models for *Gibbard-Satterthwaite games* analyzing the game-theoretic properties of agent strategic behavior in social choice was studied by [1, 7, 8, 12].

2 Preliminaries

Basic Model Suppose we have n agents making a decision on p issues $\mathcal{P} = \{1, 2, \ldots, p\}$, over the finite domain $\mathcal{D} = \times_{i=1}^{p} D^i$. Each agent $j \le n$ is endowed with a preference ranking $R_j \in \mathcal{L}(\mathcal{D})$, the set of strict linear orders over the $(\times_{i=1}^{p} |D^i|)!$ alternatives. We call the collection of agents' preferences $P = (R_1, \ldots, R_n)$ a *preference profile* and take $top(P) = (top(R_1), \ldots, top(R_n)) \in \mathcal{D}^n$ to denote each agent's *truthful* most preferred outcome.

Agents submit their preferences as *votes* for single alternatives into the *vote profile* $a = (a_1, \ldots, a_n)$, where $a_j = (a_j^1, \ldots, a_j^p) \in \mathcal{D}$. A *resolute voting rule* $f : \mathcal{D}^n \to \mathcal{D}$ maps vote profiles to a single multi-issue outcome. We will refer to $a \in \mathcal{D}$ and $a^i \in D^i$, for $i \in \mathcal{P}$, as an *alternative* when referring to an agent's vote or preference, whereas we will maintain a and a^i as an *outcome* when output by f and f^i.

Simultaneous Binary Voting Voting on multiple issues may either happen *sequentially*, in which the outcome of each issue is revealed to agents prior to voting on the next issue, or *simultaneously*, in which the outcome per issue is revealed at the same time [15]. In this work, we take the latter approach and adopt the multi-issue notation of Xia et al. [28].

We take all issues as binary, so that the domain $D^i = \{0, 1\}$ for all issues $i \in \mathcal{P}$ and denotes whether each issue is *rejected* or *accepted*, respectively. We will use the *issue-by-issue majority* function $f(a) = (f^1(a), \ldots, f^p(a)) \in \mathcal{D}$ that returns whether a strict majority of agents approved each issue. That is, let $s^i(a) = |\{j \le n : a_j^i = 1\}|$ be the number of agents that *approve* (i.e. vote $a_j^i = 1$) of issue i in the vote profile a. Then $f^i(a) = \mathbb{1}\{s^i(a) > \frac{n}{2}\}$. We call $s(a) = (s^1(a), \ldots, s^p(a))$ the induced *score vector* of a and $\frac{n}{2}$ the *decision threshold*, since this value partitions scores of issues that are rejected or approved. For the duration of this work, we will therefore assume n is odd to disallow ties. Since the vote profile a induces the score vector s, we will often use s, $s(a)$, and a interchangeably for ease of notation.

For a given vote profile a, let $a_{-j} = (a_{-j}^1, \ldots, a_{-j}^p)$ denote the profile without an agent j's vote, where $a^i = (a_{-j}^i, a_j^i)$. Then we denote by $s_{-j} = (s_{-j}^1, \ldots, s_{-j}^p)$ the corresponding *adjusted score vector*, where $s_{-j}^i = s^i - a_j^i$ is taken by component-wise subtraction. For an alternative vote $\hat{a}_j$, we will use $s_{-j} + \hat{a}_j$ to denote the score vector by replacing agent j's vote a_j by $\hat{a}_j$.

Preference Domain Restrictions Let $O = \{o_1, \ldots, o_p\}$ be some order over the p issues. A preference ranking R_j is called *O-legal* if j's preference for each issue $o_i \in \{1, \ldots, p\}$ is conditionally independent of $\{o_{i+1}, \ldots, o_p\}$ given $\{o_1, \ldots, o_{i-1}\}$, for all $i \in \mathcal{P}$ [28]. That is, an agent's preference for o_i may only depend on the outcomes of the issue prior to it on the order O.

We call the preference profile P *O-legal* if every ranking is O-legal for the same order O. A preference ranking R_j is called *separable* if it is O-legal for *any* order O. That is, their preference for each issue o_i does not depend on the outcomes of any other issue.

Example 1. *Suppose we have $n = 3$ agents and $p = 2$ binary issues. Let the agents have the preferences $P = (R_1, R_2, R_3)$, where:*

$$R_1 : (1,0) \succ_1 (1,1) \succ_1 (0,1) \succ_1 (0,0)$$
$$R_2 : (1,1) \succ_2 (0,0) \succ_2 (0,1) \succ_2 (1,0)$$
$$R_3 : (0,0) \succ_3 (0,1) \succ_3 (1,0) \succ_3 (1,1)$$

The agents' truthful vote profile is then $a = top(P) = ((1,0),(1,1),(0,0))$. The score vector for this vote is $s(a) = (2,1)$, where we've added the number of approvals for each issue independently. The outcome for this vote is then $f(a) = (1,0)$, where we've used the decision rule $\mathbb{1}\{s^i(a) > \frac{3}{2}\}$ for each issue i.

Notice that agent 1's ranking R_1 is O-legal for the order $O = \{1,2\}$. That is, they prefer $1 > 0$ on the first issue, independent of their preference for the second. Still, the agent's preference for the second issue depends on the outcome for the first issue. Agent 3's ranking R_3 is separable, as they always prefer $0 > 1$ on the each issue, independent of their preference for the other issue. Agent 2's ranking R_2 is neither separable nor O-legal.

Furthermore, notice that agent 2 can improve the outcome for themselves by voting for $\hat{a}_2 = (0,1)$ instead of $a_2 = (1,1)$. The adjusted score vector without their vote is $s_{-2} = (1,0)$, so $s_{-2} + \hat{a}_2 = (1,1)$. This gets mapped by f to $(0,0) \succ_2 (1,0) = f(a)$.

Improvement Dynamics We make use of the iterative voting model that was introduced by Meir et al. [22] and refined for voting under uncertainty by Meir et al. [21]. Given agents' truthful preferences P, we consider an iterative process of vote profiles $a(t) = (a_1(t), \ldots, a_n(t))$ that describe agents' reported votes over time $t \geq 0$. For each round t, a *scheduler* ϕ chooses an agent j to make an *improvement step* over their prior vote $a_j(t)$ by the application of a specified *response function* $g_j : \mathcal{D}^n \to \mathcal{D}$. All other votes besides j are presumed to be unchanged.

A scheduler is broadly defined as a mapping from sequences of vote profiles $\{a(t)\}_{t\geq 0}$ to an agent with an improvement step in the latest vote profile [2]. An improvement step must be selected if one exists from a particular vote profile, and a vote profile where no improvement step exists is called *an equilibrium.* If an improvement sequence lands on an equilibrium, we say it *converges.*

Definition 1. *A vote profile a is an* equilibrium *if $g_j(a) = a_j$ for all agents j.*

The literature on game dynamics considers different types of response functions, schedulers and other assumptions, such as the initial profile (see e.g., [3, 10, 19]). This means that there are multiple levels in which a voting rule may guarantee convergence [23]. In this work, we consider two types of response functions: best responses, following Meir et al. [22] in the setting without uncertainty, and local dominance, adapted from Meir et al. [21] and Meir [20] in the setting with uncertainty. Both response functions depend on agents' implicit preferences.

We first define the *best response* (BR) function as follows.

Definition 2 (best response). *Fix the vote profile a. Then $g_j(a) := \hat{a}_j$ such that agent j could not have achieved a more preferred outcome $f(a_{-j}, \hat{a}_j)$ than voting for $\hat{a}_j$ from a.*

Second, we consider response functions based on the notions of *strict uncertainty* and *local dominance* [6, 21]. Let $S \subseteq \times_{i=1}^{p} \mathbb{N}^{|D^i|}$ be a set of score vectors. Informally, S describes which score vectors an agent believes to be possible given their uncertainty of the real score vector s. A local dominance step is then an improvement step in an agent's vote that is weakly better off than their original for every possible vector in S, and strictly better off for some vector in S. This is formally defined as follows.

Definition 3. *We say that the vote $\hat{a}_j$ S-**beats** a_j if there is at least one score vector $v \in S$ such that $f(v + \hat{a}_j) \succ_j f(v + a_j)$.*

This corresponds to agent j believing that for some score $v \in S$, they can achieve a strictly better outcome by voting for $\hat{a}_j$ than a_j [6, 21].

Definition 4. *We say that the vote $\hat{a}_j$ S-**dominates** a_j if both (I) $\hat{a}_j$ S-beats a_j; and (II) a_j does not S-beat $\hat{a}_j$.*

That is, it may be the case that j should vote for $\hat{a}_j$ rather than a_j, but it is never the case that j should vote for a_j rather than $\hat{a}_j$. We then define agent j as having a *local dominance improvement* step as follows:

Definition 5 ((Restricted) local dominance improvement). *Let a be a vote profile, and for a given agent j, let LD_j be the set of alternatives that S-dominate a_j and are not themselves S-dominated. Then*

$$g_j(a) := \begin{cases} a_j, & LD_j = \emptyset \\ \hat{a}_j, & \text{otherwise} \end{cases}$$

*where voting for $\hat{a}_j$ yields j's most preferred outcome among the set $\{f(a_{-j}, \tilde{a}_j) : \tilde{a}_j \in LD_j\}$. We call g_j **restricted** to issue i if we only consider alternatives in LD_j that differ from a_j on the i^{th} issue.*

Note that S-dominance is transitive and antisymmetric, but not complete, so agent j may not necessarily have a local dominance improvement step. Moreover, to fully define the model, we need to specify the set of possible scores S in every profile a. For example, if $S = \{s(a)\}$ and agents have no uncertainty about the score vector, then local dominance moves coincide with best response, and an equilibrium coincides with Nash equilibrium. Therefore, BR dynamics is a special case of LDI dynamics without uncertainty.

3 Multi-Issue Uncertainty Model

In the original local dominance papers, which considered single-issue voting, the uncertainty set S was constructed via *distance-based uncertainty*, i.e. by considering all score vectors which are of certain distance from the current profile [20, 21]. We adapt this idea to multi-issue voting and are motivated by a single agent j making a restricted LDI step on issue i on round $t \geq 0$ from vote profile $a = a(t)$.

Suppose we have some distance measure for score vectors $d : \mathbb{N}^p \times \mathbb{N}^p \to \mathbb{R}$. We model agent j's uncertainty about the adjusted score vector s^i_{-j} for issue i by the following *uncertainty score set*:

$$\tilde{S}^i_{-j}(s; \delta^i_j) := \left\{ v^i \in \mathbb{N} : d(v^i, s^i_{-j}) \leq \delta^i_j \right\}$$

where $\delta_j^i \in \{0, 1, \ldots, \frac{n-1}{2}\}$ is the *uncertainty parameter*. That is, given all other agents' votes a_{-j}^i, agent j is uncertain of what the real score is $v^i \in \tilde{S}_{-j}^i(s; \delta_j^i)$. Then we take $\tilde{S}_{-j}(s; \delta_j^i) := \times_{i=1}^p \tilde{S}_{-j}(s; \delta_j^i)$. We drop the parameters for ease of notation.

The local dominance model specified by Meir et al. [21] allowed for an arbitrary distance function, such as the L1 norm or the infinity norm. Here, we fix $d(v^i, s_{-j}^i) = |v^i - s_{-j}^i|$ as the L1-norm for all issues i, agents j, and rounds t.

We next consider j's ability to offset this uncertainty set with a replacement vote $\hat{a}_j$. We take

$$\tilde{S}_{-j} + \hat{a}_j := \{v + \hat{a}_j : v \in \tilde{S}_{-j}\}$$

to be the uncertainty set without j's vote, offset by considering j's replacement vote $\hat{a}_j$. That is, although j is uncertain about the value of s_{-j}, they know that their new vote $\hat{a}_j$ must offset every possible $v \in \tilde{S}_{-j}$ by a fixed amount.

Finally, we consider how uncertainty affects the voting rule f. Given the vote profile a and a replacement vote $\hat{a}_j$, take

$$f^i(\tilde{S}_{-j}^i + \hat{a}_j^i) = \left\{f^i(v^i + \hat{a}_j^i) : v^i \in \tilde{S}_{-j}^i\right\}$$

to be the set of possible outcomes for issue i for any uncertain score $v^i \in \tilde{S}_{-j}^i$. Since j doesn't know which $v^i \in \tilde{S}_{-j}^i$ is correct, j considers all possible outcomes. Then

$$f(\tilde{S}_{-j} + \hat{a}_j) = \times_{i=1}^p f^i(\tilde{S}_{-j}^i + \hat{a}_j^i)$$

is the issue-by-issue cross-product of the uncertainty outcome sets.

Example 2. *Suppose there are $n = 13$ agents and $p = 2$ binary issues. Let the vote profile a be defined such that the score vector $s = (5, 4)$ and agent 1's vote is $a_1 = (0, 1)$. Suppose the uncertainty parameters for agent 1 are $\delta_1^1 = \delta_1^2 = 1$. Then we have the uncertainty score set for a, without agent 1's vote, is:*

$$\tilde{S}_{-1} = \{4, 5, 6\} \times \{2, 3, 4\}$$

For the first issue, we took a bandwidth of $\delta_1^1 = 1$ around the score $s_{-1}^1 = s^1 - a_1^1 = 5 - 0 = 5$. For the second issue, we likewise took a bandwidth of $\delta_1^2 = 1$ around the score $s_{-1}^2 = s^2 - a_1^2 = 4 - 1 = 3$.

Next, consider the uncertainty outcomes according to f with the alternative vote $\hat{a}_1 = (1, 1)$. Then $\tilde{S}_{-1} + \hat{a}_1 = \{5, 6, 7\} \times \{3, 4, 5\}$. Since our decision threshold for each issue is $\frac{13}{2} = 6.5$, this yields

$$f(\tilde{S}_{-1} + \hat{a}_1) = \{0, 1\} \times \{0\}$$

That is, although the second issue is definitely approved, the first issue may be rejected or approved depending on the score $v^1 \in \tilde{S}_{-1}^1 + \hat{a}_1^1$.

Note that in the case of a single issue with a non-binary domain, our uncertainty model coincides with Meir et al. [21]'s definitions.

We make the following assumptions for the duration of this work. First, we focus on restricted LDI dynamics and will, unless stated otherwise, discuss restricted local dominance in terms of an agent j making an improvement step on issue i with respect to their uncertainty set $\tilde{S}_{-j}$. This captures restricted BR dynamics as a special case when $\delta_j^i = 0\ \forall i \leq n, j \in \mathcal{P}$. Second, we will assume that the uncertainty parameter for each issue is consistent across agents and for all rounds t, although not necessarily across issues. We denote $\delta^i := \delta_1^i = \ldots = \delta_p^i\ \forall i \in \mathcal{P}$. Finally, following Meir et al. [21] and Meir [20] we limit our discussion to improvement sequences beginning from the truthful profile $a(0) = top(P)$. This assumption is consistent within the iterative voting literature, as it is rather plausible that agents may prefer to give their truthful votes when they have no information about others' preferences and no incentive to a priori vote otherwise [6, 22, 24].

3.1 Pivotality and Strategic Responses

Given the vote profile a, consider an agent j changing their vote a_j on issue i to the alternative $\hat{a}_j$. Under BR dynamics, without uncertainty, j only changes their vote if they can improve the outcome $f(a)$ to one more favorable, according to their preference ranking R_j. This happens under two necessary conditions. First, j must be a *pivotal* voter on the i^{th} issue, meaning that they can practically change the outcome. It is easy to verify this only happens if $s^i = \frac{n-1}{2}$ and $a_j^i = 0$ or $s^i = \frac{n+1}{2}$ and $a_j^i = 1$. That is, the score for i must be one vote away from crossing over the decision threshold $\frac{n}{2}$ and changing j's vote will change i's outcome. Otherwise, if s^i is too far from $\frac{n}{2}$, then j won't have any incentive to change their vote as doing so unilaterally won't change the outcome.

The second condition is that j must have a preference over the i^{th} issue. That is, they must know whether they prefer $0 \succ 1$ or $1 \succ 0$. This is always the case in the multi-issue setting without uncertainty. Generally, j's preference for i depends on the outcomes of each other issue $k \neq i$. Since j always knows whether each issue is rejected or accepted, they always have a preference for each i.

LDI dynamics, with uncertainty, is similar to BR dynamics with these two conditions, in that j only changes their vote if they are both pivotal and believe they may be able to improve the outcome with respect to R_j. This is formalized as follows:

Definition 6. *For a given vote profile a with induced score vector s, agent j partitions the issues $\mathcal{P}$ into three sets: **certain** issues $C_j(s)$, **borderline** issues $B_j(s)$, and **uncertain** issues $U_j(s)$, as follows:*

- *j perceives i as **certain** if $s^i < \frac{n-1}{2} - \delta_j^i$ or $s^i > \frac{n+1}{2} + \delta_j^i$*
- *j perceives i as **uncertain** if $s^i \in (\frac{n-1}{2} - \delta_j^i, \frac{n+1}{2} + \delta_j^i)$*
- *j perceives i as **borderline** if $s^i \in \{\frac{n-1}{2} - \delta_j^i, \frac{n+1}{2} + \delta_j^i\}$*

Since (δ_j^i) is constant $\forall j \leq n$, we will drop the subscript notation.

These definitions are characterized as follows (see Figure 1): Fix a vote profile a. First, if j perceives i as certain, then $\forall a_j, \hat{a}_j \in \{0, 1\}$ we have all $v^i \in \tilde{S}_{-j} + \hat{a}_j$ is above or below the decision threshold $\frac{n}{2}$. Therefore the outcome $f^i(\tilde{S}_{-j} + \hat{a}_j)$ is either definitively rejected $\{0\}$ or accepted $\{1\}$. No agent can change i's outcome, so no agent is pivotal. This is the case for the second issue in Example 2.

Second, if j perceives i as uncertain, then $\frac{n-1}{2}$ and $\frac{n+1}{2}$ are both in $\tilde{S}_{-j}^i + a_j^i$. In this case, j cannot tell whether the issue is rejected or approved and $f^i(\tilde{S}_{-j}^i + a_j^i) = \{0, 1\}\ \forall a_j \in \{0, 1\}$. Therefore, all agents are pivotal for uncertain issues.

Finally, suppose j perceives i as borderline. Given j's current vote a_j, the outcome $f^i(v^i)$ is either rejected or approved $\forall v^i \in \tilde{S}_{-j}^i + a_j^i$, as though the issue were certain. If j were to improve their vote a_j^i for the issue, changing it to $\hat{a}_j^i = 0$ if $s^i = \frac{n+1}{2} + \delta^i$ or $\hat{a}_j^i = 1$ if $s^i = \frac{n-1}{2} - \delta^i$, then the outcome would become uncertain. Note that j must be voting such that $a_j^i = f^i(a)$ in this case. Otherwise the agent will not have an incentive to change their vote. Therefore, only some agents are pivotal for borderline issues. This is the case for the first issue in Example 2.

Whereas for BR dynamics, agent j always has a preference over issue i, this may not be the case for LDI dynamics, as demonstrated by the following example.

EXAMPLE 3. *Fix $n = 13$, $p = 2$, and the truthful vote profile a such that $s = (6, 6)$. Let agent 1 have ranking*

$$R_1 : (1, 0) \succ_1 (1, 1) \succ_1 (0, 1) \succ_1 (0, 0)$$

which is O-legal for the order $O = \{1, 2\}$. For the first issue, agent 1 always prefers $1 \succ 0$, so they will always maintain their truthful vote. Even though agent 1 is pivotal on the second issue, the alternative vote $\hat{a}_1 = (1, 1)$ does not $(\tilde{S}_{-1})$-beat a_1 because the agent is uncertain about the first issue.

In this example, Agent 1 does not have a restricted LDI step because their preference for the second issue depends on the outcome of the first issue, which is uncertain. This leads to the conclusion that in order for an agent j to have a restricted LDI step on issue i, all other issues $k \neq i$ must either be (1) certain or borderline, or (2) must not affect j's preference for i. These two conditions are formalized as necessary conditions for LDI steps as follows:

LEMMA 1. *Let a be a vote profile and s be its induced score vector. Consider the alternative vote $\hat{a}_j$ for agent j that differs from a_j on issue i. Then j has a restricted LDI step from a_j to $\hat{a}_j$ only if the following two conditions hold:*

(1) Agent j is pivotal on issue i. Formally, $i \in U(s)$, or $s^i = \frac{n-1}{2} - \delta^i$ and $a^i_j = 0$, or $s^i = \frac{n+1}{2} + \delta^i$ and $a^i_j = 1$

(2) Agent j has a preference over the outcome of issue i. Formally, j's preference for issue i is conditionally independent of every $k \in U(s)$ given the outcomes $\{f^h(a) : h \in B(s) \cup C(s)\}$

The proof can be found in Appendix A.

4 Convergence Results

In this section, we discuss sufficient conditions for convergence and non-convergence for restricted BR and LDI dynamics.

4.1 Non-convergence for general preferences

Unlike iterative plurality in the single-choice setting [22], we notice in Example 5 in Appendix A that BR dynamics may fail to converge.

Likewise, when we add uncertainty into our model, we note that uncertainty by itself is not a sufficient condition to guarantee convergence, as demonstrated by the following example.

EXAMPLE 4. *Suppose there are $n = 13$ agents and $p = 2$ binary issues, with $\delta^1 = 1$ and $\delta^2 = 2$. We define four types of agents by their preference rankings as follows:*

(Type 1) Three agents have rankings: $(0, 1) \succ (1, 1) \succ (1, 0) \succ (0, 0)$
(Type 2) Five agents have rankings: $(0, 0) \succ (0, 1) \succ (1, 1) \succ (1, 0)$
(Type 3) Four agents have rankings: $(1, 0) \succ (1, 1) \succ (0, 0) \succ (0, 1)$
(Type 4) One agent has ranking: $(1, 1) \succ (1, 0) \succ (0, 1) \succ (0, 0)$

Figure 2 (right) demonstrates a cycle induced by LDI dynamics across 16 vote profiles. Four of these vote profiles are depicted explicitly by a matrix whose rows represent the vote for every agent of each type. For example, in the truthful profile, every Type 1 agent is voting for $(0, 1)$, while in the upper-right profile, every Type 1 agent is voting for $(1, 1)$. A directed edge denotes which type of agent has an improvement step per round; all agents of that type take their improvement step (in any order). Therefore, between every '1' transition, there are two vote profiles implicitly stated in the figure; between every '2' transition, there are four vote profiles implicitly stated.

Consider, for example, Type 1 agents from the truthful profile. By Lemma 1, every Type 1 agent has an LDI step from $(0, 1)$ to $(1, 1)$ as long as $s^2 = 4 = \frac{n-1}{2} - \delta^2$ is borderline, and $s^1 \in [5, 8] = [\frac{n-1}{2} - \delta, \frac{n+1}{2} + \delta]$ is borderline or uncertain. In the meantime, while $s^1 \in (5, 8)$, no other type of agent has an LDI step. This is because issue 1 is uncertain and each other agent's preferences for issue 2 depends on the outcome for issue 1. The other LDI steps follow from similar reasoning, yielding the cycle shown in the figure. Finally, the figure represents all possible LDI sequences from the truthful vote profile. Neither Type 3 nor 4 agents change their votes as their rankings are fully separable.

This example demonstrates that the uncertainty parameters affect the issues independently. Since agents only change their vote for one issue at a time, each agent only needs to know whether other issues are uncertain or not, but not their scores specifically.

4.2 O-legal preferences

Note that in our last example, each ranking depicted was O-legal for different orders O. That is, R_1 and R_3 were O-legal for the order $\{2, 1\}$, while R_2 was O-legal for the order $\{1, 2\}$. In the following theorem, we show that by restricting agents' preferences to O-legal profiles, restricted LDI dynamics do converge.

THEOREM 1. *Suppose $n \geq 3$ odd agents have O-legal preferences abiding by order $\tilde{O} = \{o_1, \ldots, o_p\}$. Let there be $p \geq 2$ issues with uncertainty parameter $\delta \geq 0$. Then, under restricted LDI dynamics, every improvement sequence converges.*

PROOF. Suppose for contradiction there is a cycle among the vote profiles $C = \{a(t_1), \ldots, a(t_T)\}$ from round t_1 to t_T, where $a(t_T + 1) = a(1)$. Let $\hat{o}$ be the highest order issue whose score changed in C according to $\tilde{O}$. Let $t^* \in \{t_1, \ldots, t_{T-1}\}$ be the first round that some agent j changed their vote for issue $\hat{o}$ in the cycle. Likewise, let $t^{**} \in \{t_2, \ldots, t_T\}$ be the last round that j changed their vote for issue $\hat{o}$ in the cycle. Then on round t^{**}, j's preference for $\hat{o}$, conditioned on the scores for all other issues, must be different than that on round t^*. However, since $\hat{o}$ is the highest order issue in $\tilde{O}$ whose score changed in C, the score for any issue before $\hat{o}$ in $\tilde{O}$ mustn't have changed. Moreover, since j's preferences are O-legal, the scores for issues after $\hat{o}$ in $\tilde{O}$ do not affect j's preference for $\hat{o}$. We therefore conclude that j's preference for $\hat{o}$ did not change between round t^* and t^{**}, which is a contradiction. □

As a corollary, we note that the prior result holds for separable preferences as well.

COROLLARY 1. *Suppose $n \geq 3$ odd agents have separable preferences over $p \geq 2$ issues with uncertainty parameter $\delta \geq 0$. Then, under restricted LDI dynamics, every improvement sequence converges.*

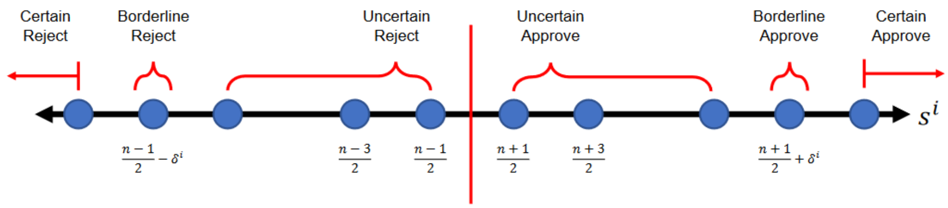

Figure 1: Uncertainty characterization of scores s^i given δ^i

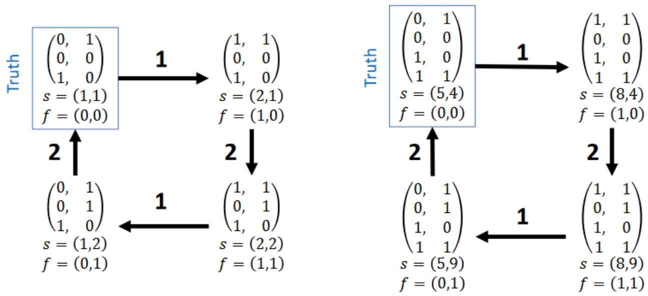

Figure 2: Example of cycles under restricted LDI dynamics without (left) and with (right) uncertainty.

4.3 Dynamic Uncertainty

Suppose an agent j is considering its improvement step on an issue i from some vote profile a. Like Example 4, consider a different issue $k \neq i$ with a different uncertainty parameter $\delta^i \neq \delta^k$ such that $f^k(a)$ is certain. Suppose agent j's preferences for issue i depend on those for k. We know that either $s^k < \frac{n-1}{2} - \delta^k$ or $s^k > \frac{n+1}{2} + \delta^k$. If we were to increase δ^k sufficiently large, then f^k would no longer be certain and agent j may no longer have a preference over issue i. Thus, *increasing* the uncertainty parameter for other issues, other than the one an agent is looking at, eliminates LDI steps.

On the other hand, suppose an agent j is considering its improvement step on issue i for some vote profile a, but that $f^i(a)$ is uncertain. If we were to decrease δ^i sufficiently small, then f^i would gain certainty and agent j would no longer have an incentive to change their vote a^i_j. Thus, *decreasing* the uncertainty parameter for the issue an agent is looking at eliminates LDI steps.

We consider an extension of LDI dynamics with *dynamic uncertainty*, in which the issue an agent is currently deciding on has less uncertainty than any other issue. We show in the following claim that this eliminates cycles.

Theorem 2. *Let there be $n \geq 3$ odd agents, $p \geq 2$ issues, an uncertainty parameter $\delta^c \geq 1$ for the current issue and $\delta^o > \delta^c$ for other issues. Let the scheduler ϕ be defined such that an agent only considers an issue if it's borderline with respect to δ^c. Then any sequence of LDI moves is finite.*

Proof. By Lemma 1, an agent j has an LDI step on issue i only if, for all other issues k that affect j's preference of i, k is either borderline or certain. This means that $s^i \in [\frac{n-1}{2} - \delta^c, \frac{n+1}{2} + \delta^c]$ while either $s^k \leq \frac{n-1}{2} - \delta^o$ or $s^k \geq \frac{n+1}{2} + \delta^o$ for every such k. These ranges do not overlap. If there is only one such i for a truthful vote profile a, then agent j may update their vote for that issue. This issue can only change a finite number of times and the issue's score will stay in the $[\frac{n-1}{2} - \delta^c, \frac{n+1}{2} + \delta^c]$ region.

However, consider the case where there are two issues i_1 and i_2 whose scores are in $[\frac{n-1}{2} - \delta^c, \frac{n+1}{2} + \delta^c]$, according to truthful vote profile a, and suppose j is looking to change issue i_1. Then i_2 is an *other* issue with respect to j and is considered uncertain. If j's preference for i_1 depends on i_2, conditioned on the tertiary outcomes, then j does not have a preference for i_1 and does not have an LDI step.

Similar reasoning holds for more than two issues whose scores are in $[\frac{n-1}{2} - \delta^c, \frac{n+1}{2} + \delta^c]$. If these issues are no dependent on each other, then agents can update these scores like the single-issue case only a finite number of times. ◻

5 Non-atomic Model

We show that cycles are an exception rather than the norm, and their existence hinges on the presence of "borderline profiles" – states in which some voters perceive an issue as certain while others do not due to their own vote. Since borderline profiles become more rare as the number of voters increase, we can think of them as an artifact of the model, rather than an actual barrier to convergence. We therefore turn to study a nonatomic model of voting from the literature (which can be thought of as the limit case of a large population), showing a very general convergence result.

5.1 Model Definition

We follow the basic definitions from [20], using our existing notation where possible.

Basic notations We do not have a finite set of voters. Rather, a *preference profile* $Q \in \Delta(\mathcal{L}(D))$ is a distribution over preferences, specifying the fraction of voters $Q(R)$ with each preference order $R \in \mathcal{L}(D)$. Since a single voter has negligible influence, we only consider moves by subsets of voters whose size is a multiple of ϵ (for some arbitrarily small ϵ). All voters in each set have the same preference and move simultaneously (although in an uncoordinated manner; see Appendix IV of [20]). We denote by J the collection of these $1/\epsilon$ sets. Since all voters in set $j \in J$ are indistinguishable we refer to "voter j" which is an arbitrary voter in the set j. Also $R_j \in \mathcal{L}(D)$ is the preference of an arbitrary voter in the set, and a_j is the vote of an arbitrary voter in the set in profile a.

Winner determination For a given vote profile a we define the *score vector* s induced by the vote profile. In particular $s^i(a) := |\{j : a^i_j = 1\}|\epsilon \in [0, 1]$. Winner determination is exactly as in the atomic model (we assume $|J|$ is odd so there are no ties). $\mathcal{S} = \mathbb{R}_+^b$ denotes the set of all possible score vectors.

Distance-based uncertainty Following [20, 21] we assume voters derive their beliefs using a distance-based strict uncertainty model. That is, for we suppose agent have a fixed uncertainty $\delta^i \in [0, \frac{n}{2}]$ over each issue i. Then let $\tilde{S}(s, \delta) \subseteq \mathcal{S}$ be the *uncertainty set* of score vectors that are an distance at most δ from s. Formally, $\tilde{S}(s, \delta) = \{\tilde{S}^1(s, \delta^1), \ldots, \tilde{S}^p(s, \delta^p)\}$ where:

$$\tilde{S}^i(s, \delta^i) = \{v : |v - s^i| \leq \delta^i\}.$$

Note that in sharp contrast to the atomic model, the uncertainty set $\tilde{S^i}$ does not depend on the playing voter(s), as all voters agree on the possible states.[1]

Given a vote profile a and associated score vector $s = s(a)$, we partition the issues $\mathcal{P}$ into two sets: *uncertain* issues $U(s) := \{i \in \mathcal{P} : s^i \in (\frac{n-1}{2} - \delta^i, \frac{n+1}{2} + \delta^i)\}$ and *certain* issues $C(s) := \{i \in \mathcal{P} : s^i < \frac{n-1}{2} - \delta^i \text{ or } s^i > \frac{n+1}{2} + \delta^i\}$.

Here we have another difference from the atomic model that simplifies matters, namely that there are no borderline profiles. Uncertain issues have their uncertainty sets intersecting with the decision threshold $\frac{n}{2}$, so agents cannot tell whether their outcome is rejected or approved. Certain issues, rather, have their uncertainty sets distinct from the decision threshold, so agents can determine their outcome.

Local dominance An agent of type $j \in J$ has a restricted local dominance improvement step from vote a_j to $\hat{a}_j$ on issue i if $\hat{a}_j$ S-dominates a_j (see Definition 5).

The response function g_j is also as in the atomic model, except that its domain (all possible profiles) is now $\mathcal{D}^{|J|}$ rather than $\mathcal{D}^n$. The definition of equilibrium does not change.

5.2 Convergence Result

This first lemma characterizes local dominance steps.

LEMMA 2. *Suppose an agent j makes a LDI step from a_j to $\hat{a}_j$ from the vote profile a. Then (1) the votes a_j and $\hat{a}_j$ only differ by a single issue i; (2) issue $i \in U(a)$; (3) agent j has a preference over the i^{th} issue, meaning that j's preference for issue i is conditionally independent of the scores of uncertain issues s^k for $k \in U(a)$ given the known outcomes of certain issues $\{f^k(s) : k \in C(a)\}$.*

PROOF. Part 1 is by definition of a restricted LDI step. For Part 2, agent j can only make an improvement on issue i if it is uncertain. Otherwise j cannot change the outcome of the issue if its score is too far from the decision threshold.

Part 3 also follows from the definition of restricted local dominance. If agent j's preference for issue i depends on some uncertain issue $k \in U(a)$, given the certain issues' outcomes, and j does not know the outcome for that issue, then j cannot tell whether they prefer to approve or reject i. □

Our next result states that voters always converge in the non-atomic model.

THEOREM 3. *Any sequence of LDI steps converges from any initial vote profile.*

[1] The same distinction occurs between the atomic and nonatomic model in Plurality [20], and simplifies convergence proofs in that domain as well.

PROOF. Suppose towards a contradiction there is a cycle among the vote profiles $C = \{a(t_1), \ldots, a(t_T)\}$ from round t_1 to t_T, where $a(t_T + 1) = a(t_1)$. By Lemma 2.2 the set of uncertain issues $U(a)$ is constant for all profiles $a \in C$ in the cycle. For some agent $j \in J$, let t^* be the first round that they change their vote in the cycle, and denote the singular issue they change i (by Lemma 2.1). Let $t^{**} > t^*$ be the last round j changes their vote on issue i, back to their original vote in $a_j^i(0)$. By Lemma 2.3, on each round $t \in \{t^*, t^{**}\}$ j's preference for i must be conditionally independent of the score of each issue $k \in U(a(t))$, given the certain issues' outcomes $\{f^k(s(t)) : k \in C(a(t))\}$. Since the certain issues' outcomes don't change due to LDI dynamics and since j's preference for i does not matter on the scores of each issue $k \in U(a), \forall a \in C$, we conclude that j's preference for i is the same in each round $t \in \{t^*, t^{**}\}$. This contradicts our assertion that j made an improvement step on round t^{**}, so such a cycle C does not exist. □

Note that the proof also provides us with a bound on the total number of moves that each voter makes: After a voter moves (say, changes her vote on some issue i), she will not change her vote on i until some other issue becomes certain. Once an issue becomes certain this is irreversible, so the total number of steps by a single voter is at most p^2.

6 Experiments

We generate election data with $n \in \{7, 11, 15, 19\}$ agents, $p \in \{2, 3, 4, 5\}$ issues, and we vary the uncertainty parameter $\delta \in \{0, 1, 2, 3\}$ which is constant for all agents and issues. For each of these combinations, we generate 10,000 (truthful) preference profiles where agents' rankings are uniformly and independently sampled at random – i.e., preferences are sampled from the *impartial culture* (IC) distribution. We simulate LDI dynamics from the truthful vote profile $a(0) = top(P)$ for preference profile P using a uniformly random scheduler ϕ. That is, for each vote profile $a(t)$, ϕ selects $a(t + 1)$ uniformly at random from the set of valid LDI steps across all agents and issues. If there are no such valid LDI steps, we say the LDI sequence has converged in t iterations. Otherwise, we take $t = 10,000$ iterations as a stopping condition, in which we say the sequence has cycled. We use 'd' to denote δ in the following figures.

First, we ask how frequently cycles occur. With uncertainty ($\delta > 0$) the answer is almost never, with less than 5 samples cycling for any parameter combination. Figure 3 demonstrates the number of samples that cycle in the setting without uncertainty ($\delta = 0$). For $p \in \{2, 3\}$ issues, the number of cycles first decreases then increases as n increases, whereas for $p \geq 4$ issues, the number of cycles strictly increases as n increases. This demonstrates that uncertainty can greatly diminish the likelihood of cycles, despite them still existing in the worst case.

Second, we identify both the percentage of truthful vote profiles that are not themselves in equilibrium, and how many steps it takes for LDI dynamics to converge, conditioned on profiles whose dynamics do not cycle. Figure 4 demonstrates that the percentage of profiles with LDI dynamics is inversely proportional to $\frac{\delta}{\sqrt{n}}$. That is, a truthful profile is more likely to be in equilibrium if the amount of uncertainty agents have over the issues' score vector is large relative to the number of agents. On the other hand, if the number

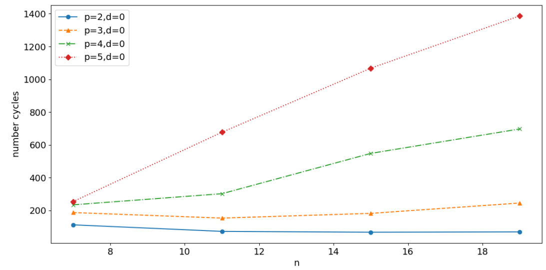

Figure 3: Number of profiles that cycle as n increases; $\delta = 0$

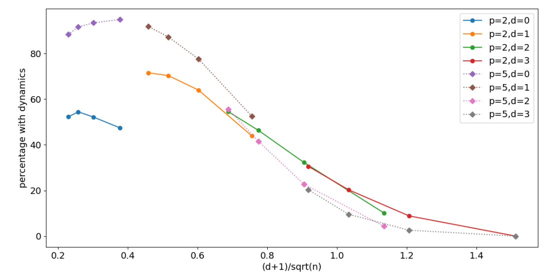

Figure 4: Percentage of profiles with dynamics as $\frac{\delta}{\sqrt{n}}$ increases; $p \in \{2, 5\}$

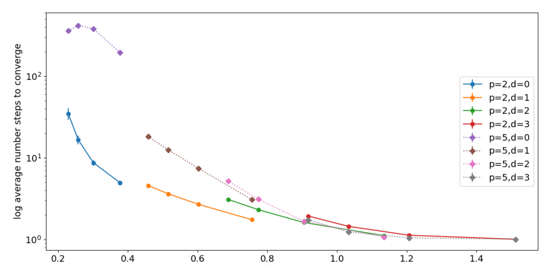

Figure 5: Log average number steps to converge as $\frac{\delta}{\sqrt{n}}$ increases; $p \in \{2, 5\}$; 95% CI (too small to see)

of agents is relatively large, then some agent is more likely to have an improvement step. Figure 5 demonstrates a similar pattern about the length of LDI sequences: an LDI sequence will converge much quicker if δ is relatively larger compared to n.

Intuitively, recall that the range of uncertain scores is proportional to δ and not related to n. Also, many preference rankings have dependencies among their issues, especially as p increases. When $\frac{\delta}{\sqrt{n}}$ is high, more issues are likely to be uncertain, so fewer agents will be decisive about their preference at all and have S-dominant steps. On the other hand, without uncertainty, we note that the percentage of profiles with dynamics increases with the number of issues. Likewise, the length of LDI sequences is significantly higher than among profiles with uncertainty.

Finally, we study how LDI dynamics affects the quality of the equilibrium winning alternative (i.e., $f(a(T))$ if $a(T)$ is an equilibrium) relative to the truthful winning alternative. For every profile that did not cycle, we found the Borda welfare of the truthful and

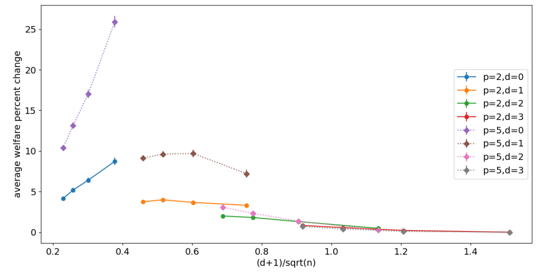

Figure 6: Average percent change in welfare as $\frac{\delta}{\sqrt{n}}$ increases; $p \in \{2, 5\}$; 95% CI (too small to see)

equilibrium winning vote profiles and recorded the percentage change. The Borda utility for a single agent with ranking R for an outcome a is 2^p minus the index of a's position in R; the Borda welfare for a vote profile is the sum total utilities across all agents. Figure 6 plots the average percent change in Borda welfare among profiles that did not cycle, as $\frac{\delta}{\sqrt{n}}$ increases. We immediately notice that all averages are positive, indicating that equilibrium winning alternatives have a higher social welfare on average than the truthful winning alternative. Furthermore, we note a similar trend to the proportion of profiles with dynamics and length of LDI sequences. That is, average change in welfare generally is correlated with n and inversely correlated with δ.

7 Conclusion

In this work we introduced iterative voting for multiple-issue binary-outcome elections. We synthesized this binary multi-issue iterative voting model with Meir et al.'s local dominance theory [21]. We demonstrated that under general agent preferences, restricted local dominance improvement dynamics may fail to converge, with or without uncertainty. On the other hand, if we restrict agents to only have O-legal preferences, then our model converges. In addition, we introduce dynamic uncertainty and non-atomic models as sufficient conditions for convergence.

Finally, we empirically demonstrate that cycles are prevalent for BR dynamics without uncertainty, but that uncertainty eliminates almost all cycles. When LDI dynamics do converge, the outcome they yield is on average a better-quality solution that the truthful winning alternative. We find that the likelihood an agent has an improvement step from the truthful profile, the length of LDI sequence, and the improvement in welfare of the equilibrium solution over the truthful winning alternative are all inversely correlated with $\frac{\delta}{\sqrt{n}}$ for uncertainty parameter δ and n agents. Further study will be necessary to determine the cause or degree of this correlation.

References

[1] Stéphane Airiau and Ulle Endriss. 2009. Iterated majority voting. In *International Conference on Algorithmic Decision Theory*. Springer, 38–49.
[2] Krzysztof R Apt and Sunil Simon. 2015. A classification of weakly acyclic games. *Theory and Decision* 78, 4 (2015), 501–524.
[3] Michael Bowling. 2005. Convergence and no-regret in multiagent learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*. Vancouver, Canada, 209–216.
[4] Clark Bowman, Jonathan K Hodge, and Ada Yu. 2014. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and*

decision 77, 1 (2014), 111–124.
[5] Vincent Conitzer, Toby Walsh, and Lirong Xia. 2011. Dominating manipulations in voting with partial information. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.
[6] Vincent Conitzer, Toby Walsh, and Lirong Xia. 2011. Dominating Manipulations in Voting with Partial Information. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*. San Francisco, CA, USA, 638–643.
[7] Yvo Desmedt and Edith Elkind. 2010. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce*. 347–356.
[8] Edith Elkind, Umberto Grandi, Francesca Rossi, and Arkadii Slinko. 2020. Cognitive hierarchy and voting manipulation in k-approval voting. *Mathematical Social Sciences* 108 (2020), 193–205.
[9] Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. 2016. Strategic Voting with Incomplete Information. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence* (New York, New York, USA) *(IJCAI'16)*. AAAI Press, 236–242.
[10] Drew Fudenberg and David K Levine. 2009. Learning and equilibrium. *Annu. Rev. Econ.* 1, 1 (2009), 385–420.
[11] Allan Gibbard. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41 (1973), 587–601.
[12] Umberto Grandi, Daniel Hughes, Francesca Rossi, and Arkadii Slinko. 2019. Gibbard–Satterthwaite games for k-approval voting rules. *Mathematical Social Sciences* 99 (2019), 24–35.
[13] Umberto Grandi, Jérôme Lang, Ali Ozkes, and Stéphane Airiau. 2020. Voting behavior in one-shot and iterative multiple referenda. *Available at SSRN* (2020).
[14] Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S. Rosenschein. 2017. Convergence and Quality of Iterative Voting Under Non-Scoring Rules. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 273–279.
[15] D. Lacy and E. Niou. 2000. A problem with referenda. *Journal of Theoretical Politics* 12, 1 (2000), 5–31.
[16] Jérôme Lang and Lirong Xia. 2009. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences* 57, 3 (2009), 304–324.
[17] Jérôme Lang and Lirong Xia. 2016. Voting in Combinatorial Domains. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia (Eds.). Cambridge University Press, Chapter 9.
[18] Omer Lev and Jeffrey S Rosenschein. 2012. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. 611–618.
[19] Jason R. Marden, Gürdal Arslan, and Jeff S. Shamma. 2007. Regret Based Dynamics: Convergence in Weakly Acyclic Games. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems* (Honolulu, Hawaii) *(AAMAS '07)*. Association for Computing Machinery, New York, NY, USA, Article 42, 8 pages. https://doi.org/10.1145/1329125.1329175
[20] Reshef Meir. 2015. Plurality Voting under Uncertainty. In *Proceedings of the 29th AAAI on Artificial Intelligence*.
[21] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. 2014. A Local-Dominance Theory of Voting Equilibria. In *Proceedings of the 15th ACM Conference on Electronic Commerce*. Palo Alto, CA, USA, 313–330.
[22] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. 2010. Convergence to Equilibria of Plurality Voting. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*. Atlanta, GA, USA, 823–828.
[23] Reshef Meir, Maria Polukarov, Jeffrey S Rosenschein, and Nicholas R Jennings. 2017. Iterative voting and acyclic games. *Artificial Intelligence* 252 (2017), 100–122.
[24] Svetlana Obraztsova, Evangelos Markakis, and David R. M. Thompson. 2013. Plurality Voting with Truth-Biased Agents. In *Algorithmic Game Theory*, Berthold Vöcking (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 26–37.
[25] Annemieke Reijngoud and Ulle Endriss. 2012. Voter Response to Iterated Poll Information. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Vol. 1. 635–644.
[26] Reyhaneh Reyhani and Mark C. Wilson. 2012. Best Reply Dynamics for Scoring Rules. In *Proceedings of the 20th European Conference on Artificial Intelligence* (Montpellier, France) *(ECAI'12)*. IOS Press, NLD, 672–677.
[27] Mark Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10 (1975), 187–217.
[28] Lirong Xia, Vincent Conitzer, and Jérôme Lang. 2011. Strategic Sequential Voting in Multi-Issue Domains and Multiple-Election Paradoxes. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. San Jose, CA, USA, 179–188.

A Supplementary Material

Example 5. *Fix $p = 2$ issues without uncertainty ($\delta^1 = \delta^2 = 0$) and consider $n = 3$ agents with preferences:*

$$R_1 : (0,1) \succ_1 (1,1) \succ_1 (1,0) \succ_1 (0,0)$$
$$R_2 : (0,0) \succ_2 (0,1) \succ_2 (1,1) \succ_2 (1,0)$$
$$R_3 : (1,0) \succ_3 (1,1) \succ_3 (0,0) \succ_3 (0,1)$$

Figure 2 (left) demonstrates a cycle induced by BR dynamics across four vote profiles. Each vote profile a, represented by a matrix containing votes as row-vectors, is paired with its induced score vector s and outcome f. Each directed edge is labelled with the agent that makes the BR step. For example, agent 1 improves their truthful vote $(0,1)$ to $(1,1)$, thus changing the score vector to $(2,1)$ and outcome to $(1,0)$.

The figure represents all possible BR sequences from the truthful vote profile. Notice that agent 3 never changes their vote since their ranking is fully separable.

Lemma 1. *Let a be a vote profile and consider its induced score vector s. Consider the alternative vote $\hat{a}_j$ for agent j that differs from a_j only on issue i. Then j has a restricted LDI step from a_j to $\hat{a}_j$ only if the following conditions hold:*

(1) $i \in U(s)$, OR $s^i = \frac{n-1}{2} - \delta^i$ and $a^i_j = 0$, OR $s^i = \frac{n+1}{2} + \delta^i$ and $a^i_j = 1$
(2) Agent j's preference for issue i is conditionally independent of every $k \in U(s)$ given the outcomes $\{f^h(a) : h \in B(s) \cup C(s)\}$.

Proof. We prove the lemma by demonstrating that if either of these statements do not hold, then j does not have the stated LDI step.

Suppose first that statement 1 fails. Then exclusively either (a) $i \in C(s)$, (b) $s^i = \frac{n-1}{2} - \delta^i$ and $a^i_j = 1$, or (c) $s^i = \frac{n+1}{2} + \delta^i$ and $a^i_j = 0$. Suppose case (a) and $s^i < \frac{n}{2} - \delta^i$ or $s^i > \frac{n}{2} + \delta^i$. Then regardless of $a^i_j = 1 - \hat{a}^i_j \in \{0,1\}$, we have $\forall v^i \in \tilde{S}^i_{-j} + \hat{a}^i_j$, either $v^i < \frac{n}{2}$ or $v^i > \frac{n}{2}$. Since all scores are on one side of the decision threshold, agent j cannot change i's outcome so they will not have an incentive to change their vote. Suppose case (b). Then $\forall v^i \in \tilde{S}^i_{-j} + a_j$ we have $v^i < \frac{n}{2}$ and j cannot change i's outcome by decreasing their vote. Suppose case (c). Then $\forall v^i \in \tilde{S}^i_{-j} + a_j$ we have $v^i > \frac{n}{2}$ and j cannot change i's outcome by increasing their vote.

Now suppose that statement 2 fails and statement 1 succeeds (i.e., changing j's vote from a^i_j to $\hat{a}^i_j$ can change i's outcome for some $v^i \in \tilde{S}^i_{-j}$; specifically, for $v^i = \frac{n-1}{2}$). We show that $\hat{a}_j$ does not $(\tilde{S}_{-j})$-dominate a_j. This happens if $\exists v, w \in \tilde{S}_{-j}$ such that $f(v + \hat{a}_j) \succ_j f(v + a_j)$ and $f(w + a_j) \succ_j f(w + \hat{a}_j)$. We will now construct these score vectors.

Without statement 2, then there exists a maximal subset of issues $Q \subseteq U(s) \backslash \{i\}$ such that agent j's preference for issue i depends on the joint outcomes of $\{f^k : k \in Q\}$, conditioned on the known outcomes of $\{f^h : h \in B(s) \cup C(s)\}$. However, j's preference for i is conditionally independent of $g \in U(s) \backslash (Q \cup \{i\})$. Then there exists two outcomes o_1, o_2 where $(1, o_1^{-i}) \succ_j (0, o_1^{-i})$ and $(0, o_2^{-i}) \succ_j (1, o_2^{-i})$, such that $o^k_1 = o^k_2 = f^k(a)$ for $k \notin Q$ and o_1, o_2 differ for some issues in Q. We then define v and w as follows:

(1) Set $v^i = w^i = \frac{n-1}{2}$
(2) For $h \in \mathcal{P} \backslash (Q \cup \{i\})$ set $v^h = w^h = s^k$
(3) For $k \in Q$ and $a_j^k = \hat{a}_j^k = 0$, set $v^k = \frac{n+1}{2}$ if $o_1^k = 1$ and $v^k = \frac{n-1}{2}$ otherwise. Likewise, set $w^k = \frac{n+1}{2}$ if $o_2^k = 1$ and $w^k = \frac{n-1}{2}$ otherwise.
(4) For $k \in Q$ and $a_j^k = \hat{a}_j^k = 1$, set $v^k = \frac{n-1}{2}$ if $o_1^k = 1$ and $v^k = \frac{n-3}{2}$ otherwise. Likewise, set $w^k = \frac{n-1}{2}$ if $o_2^k = 1$ and $w^k = \frac{n-3}{2}$ otherwise.

Statement 2 suggests that j knows the outcomes for issues $h \in B(s) \cup C(s)$, and j's preference for i is conditionally independent of issues $g \in U(s) \backslash (Q \cup \{i\})$. Since every issue $k \in U(s)$ is uncertain, we know that $\frac{n-1}{2}, \frac{n+1}{2} \in \tilde{S}^k$. Then setting v^k and w^k as defined in statements 3 and 4 will ensure that $f^k(v^k + a_j^k) = f^k(v^k + \hat{a}_j^k) = o_1^k$ and $f^k(w^k + a_j^k) = f^k(w^k + \hat{a}_j^k) = o_2^k$. Finally, statement 1 ensures that j has an incentive to vote according to their preference for issue i conditioned on all other outcomes. Since j's preference for i differs depending on whether the score vector is v or w, per the other statements, we have $f(v + \hat{a}_j) \succ_j f(v + a_j)$ and $f(w + a_j) \succ_j f(w + \hat{a}_j)$ as desired. Therefore we conclude that $\hat{a}_j$ does not $(\tilde{S}_{-j})$-dominate a_j, so j does not have the stated LDI step. □