

Convergence of Multi-Issue Iterative Voting under Uncertainty

Joshua Kavner¹, Reshef Meir², Francesca Rossi³ and Lirong Xia¹

¹Rensselaer Polytechnic Institute

²Technion Israel Institute of Technology

³IBM T.J. Watson Research Center

kavnej@rpi.edu, reshefm@ie.technion.ac.il, Francesca.Rossi2@ibm.com, xialirong@gmail.com

Abstract

We study strategic behavior in iterative plurality voting for multiple issues under uncertainty. We introduce a model synthesizing simultaneous multi-issue voting with local dominance theory, in which agents repeatedly update their votes based on sets of vote profiles they deem possible, and determine its convergence properties. After demonstrating that local dominance improvement dynamics may fail to converge, we present two sufficient model refinements that guarantee convergence from any initial vote profile for binary issues: constraining agents to have $\mathcal{O}$ -legal preferences, where issues are ordered by importance, and endowing agents with less uncertainty about issues they are modifying than others. Our empirical studies demonstrate that while cycles are common for agents without uncertainty, introducing uncertainty makes convergence almost guaranteed in practice.

1 Introduction

Consider a wedding planner who is deciding a wedding’s banquet and wants to accommodate the party invitees’ preferences.¹ There are three issues with two candidates each: the main course (chicken or beef), the paired wine (red or white), and the cake flavor (chocolate or vanilla). How should the planner proceed? On the one hand, they could request each attendee’s (agent’s) full preference ranking over the 2^p alternatives, for p binary issues. However, this is computationally prohibitive and imposes a high cognitive cost for agents.

On the other hand, the planner could solicit only agents’ votes and decide each issue simultaneously. Although simpler, this option admits *multiple election paradoxes* whereby agents can collectively select each of their least favored outcomes. For example, suppose three agents prefer $(1, 1, 0)$, $(1, 0, 1)$, and $(0, 1, 1)$ first, respectively on the issues, and all prefer $(1, 1, 1)$ last. Then the agents select $(1, 1, 1)$ by majority rule on each issue independently [Lacy and Niou, 2000].

A third approach is to decide the issues in sequence and have agents vote for their preferred alternative on the *current*

issue given the previously chosen outcomes. Still, the joint outcome may depend on the voting agenda and agents may be uneasy voting on the current issue if their preference depends on the outcomes of later issues [Conitzer *et al.*, 2009].

In this work, we study *iterative voting* (IV) as a different yet natural method for deciding multiple issues [Meir *et al.*, 2010]. We elicit agents’ most preferred alternatives and, given information about others’ votes, allow agents to update their reports before finalizing the group decision. This approach combines the efficiency of simultaneous voting with the dynamics of sequential voting, thus incorporating information about agents’ lower-ranked preferences without directly eliciting them. Like the former approach, agents only report their most preferred alternative. Like the latter approach, agents only update one issue at a time but are unrestricted in the order of improvements.

IV is an effective framework for its adaptability to various information and behavioral schemes. First, we consider agents with full information about the real vote profile, such as in online Doodle polls [Zou *et al.*, 2015], who update their votes to the *best response* of all others. Second, we consider agents with access only to a noisy signal about the real vote profile, such as from imprecise opinion polls [Reijngoud and Endriss, 2012] or latency in a networked system if they can only periodically retrieve accurate vote counts. These agents update their votes to those that *locally dominate* their prior reports – votes that achieve weakly better outcomes for all possible vote profiles and strictly better outcomes for some possible vote profile [Meir *et al.*, 2014].

We ask two primary questions: (1) *Under what conditions does multi-issue IV converge?* (2) *How does introducing and increasing uncertainty affect the rate of convergence?*

Prior work in single-issue IV offers mixed answers, as iterative plurality and veto have strong convergence guarantees but many other rules do not [Meir *et al.*, 2017]. This leaves us with mixed hope in the multi-issue plurality case, and if so, that it can solve other problems like multiple election paradoxes. Furthermore, in contrast to prior work, uncertainty for multiple issues plays a double role. First, like the single-issue case, agents consider themselves as possibly pivotal on any issue that is sufficiently close to a tie. Second—and this part is new—agents may be uncertain whether changing their vote on an issue would improve or degrade the outcome, as this may depend on the outcomes of other uncertain issues.

¹The full version of the paper may be found on the archive at: <https://arxiv.org/abs/2301.08873>

1.1 Our Contribution

On the conceptual side, we introduce a novel model that synthesizes prior work in local dominance strategic behavior, iterative plurality voting, and simultaneous voting over multiple issues. This generalized model naturally captures both types of uncertainty discussed above.

On the technical side, we first show that IV with or without uncertainty may not converge. We then present two model refinements that prove sufficient to guarantee convergence for binary issues: restricting agent preferences to have $\mathcal{O}$ -legal preferences and alternating uncertainty, in which agents are more certain about the current issue than others. The former converges because agents' preferences on issues are not interdependent; the latter because fewer preference rankings yield valid improvement steps. These convergence results do not extend to the multi-candidate issues setting, as IV may cycle if agents have partial order preference information.

Our convergence results for binary issues also hold for a nonatomic variant of plurality IV in which agents are part of a large population and arbitrary subsets of agents may change their vote simultaneously, establishing more general convergence results. This is discussed separately in Appendix C.

We conclude with empirical evidence corroborating our findings that introducing uncertainty eliminates almost all cycles in IV for multiple binary issues. Our experiments further suggest IV improves the quality of equilibrium vote profiles relative to their respective truthful profiles, thus reducing multiple election paradoxes. Increasing uncertainty yields faster convergence but degrades this welfare improvement.

1.2 Related Work

Our model pulls insights from research across multi-issue voting, IV, and local dominance strategic behavior. Multi-issue voting is an extensively studied problem in economics and computer science with applications in direct democracy referendums, group planning and committee elections (see e.g., [Lang and Xia, 2016] for a survey). Our work follows research in agent strategic behavior by Lang [2007], Lang and Xia [2009], Conitzer *et al.* [2009], and Xia *et al.* [2011].

Single-issue IV was initially studied by Meir *et al.* [2010] for best response dynamics and the plurality social choice rule, whose authors bounded its guaranteed convergence rate. Subsequent work demonstrated that iterative veto converges [Reyhani and Wilson, 2012; Lev and Rosenschein, 2012], although many other voting rules do not [Koolyk *et al.*, 2017] unless agents are additionally restricted in their behavior [Reijngoud and Endriss, 2012; Grandi *et al.*, 2013; Obraztsova *et al.*, 2015; Endriss *et al.*, 2016].

This review already narrows down the possibility of convergence in the multi-issue setting. We therefore restrict our attention to plurality, extending the models of Meir *et al.* [2014] and Meir [2015]. Their research found broad conditions for voting equilibrium to exist and guaranteed convergence for iterative plurality. In particular, Meir [2015] studied a nonatomic model variation where agents have negligible impact on the outcome but multiple agents update their votes simultaneously, greatly simplifying the dynamics.

Bowman *et al.* [2014] and Grandi *et al.* [2022] empirically demonstrated for multiple binary issues that IV improves the

social welfare of voting outcomes using computer simulations and human experiments respectively. Our work augments this research by characterizing convergence in settings where agents do not have complete information.

Related research studied strategic behavior in epistemic voting games when agents have uncertainty about the preferences or votes of others (see e.g., [Meir, 2018] for a survey). Notably, Chopra *et al.* [2004], Conitzer *et al.* [2011], Reijngoud and Endriss [2012], and Van Ditmarsch *et al.* [2013] focused on the susceptibility and computational complexity of voting rules to local dominance improvement steps. Game-theoretic properties of strategic behavior for Gibbard-Satterthwaite games were studied by Myerson and Weber [1993], Grandi *et al.* [2019], and Elkind *et al.* [2020]. Other forms of IV include work from Airiau and Endriss [2009], Desmedt and Elkind [2010], and Xia and Conitzer [2010].

2 Preliminaries

Basic model. Let $\mathcal{P} = \{1, 2, \dots, p\}$ be the set of p issues over the joint domain $\mathcal{D} = \times_{i=1}^p D_i$, where D_i is the finite value domain of *candidates* for issue i . We call the issues *binary* if $D_i = \{0, 1\}$ for each $i \in \mathcal{P}$ or *multi-candidate* otherwise. Each of n agents is endowed with a preference ranking $R_j \in \mathcal{L}(\mathcal{D})$, the set of strict linear orders over the $\prod_{i=1}^p |D_i|$ alternatives. We call the collection of agents' preferences $P = (R_1, \dots, R_n)$ a *preference profile* and each agent's most preferred alternative their *truthful* vote. A *vote profile* $a = (a_1, \dots, a_n) \in \mathcal{D}^n$ is a collection of *votes*, where $a_j = (a_j^1, \dots, a_j^p) \in \mathcal{D}$ collects agent j 's single-candidate vote per issue. A *resolute voting rule* $f : \mathcal{D}^n \rightarrow \mathcal{D}$ maps vote profiles to a unique outcome. We call $a \in \mathcal{D}$ and $a^i \in D_i$ for $i \in \mathcal{P}$ an *alternative* or *outcome* synonymously.

Simultaneous plurality voting. A local voting rule, applied to each issue independently, is *simultaneous* if issues' outcomes are revealed to agents at the same time. It is *sequential* according to the order $\mathcal{O} = \{o_1, \dots, o_p\}$ if outcomes of each issue o_i are revealed to agents prior to voting on the next issue o_{i+1} [Lacy and Niou, 2000]. We focus on simultaneous plurality voting and adapt the framework of Xia *et al.* [2011].

The plurality rule $f^i(a)$ applied to vote profile a on issue i only depends on the total number of votes for each of its candidates. We define the *score tuple* $s(a) := (s^i(a))_{i \in \mathcal{P}}$ as a collection of *score vectors* $s^i(a) = (s^i(c; a))_{c \in D_i}$, which compose the score of a candidate $c \in D_i$ as $s^i(c; a) = |\{j \leq n : a_j^i = c\}|$. We use the plurality rule $f(a) = (f^i(a))_{i \in \mathcal{P}} \in \mathcal{D}$, where $f^i(a) = \arg \max_{c \in D_i} s^i(c; a)$, breaking ties lexicographically on each issue.

Let a_{-j} denote the vote profile without agent j and $(a_{-j}, \hat{a}_j)$ the profile a by replacing j 's vote with the prospective vote $\hat{a}_j$. Then s_{-j} and $s_{-j} + \hat{a}_j$ denote corresponding adjusted score tuples without j and upon replacing j 's vote. We may interchange s , $s(a)$, and a for ease of notation.

Preferential dependence. A preference ranking is called *separable* if the relative ordering of candidates in each issue's domain D_i is consistent across all outcomes of the other issues. That is, an agent prefers one outcome over another if they prefer the former's candidates on each issue independently. Such rankings have the advantage that agents

may express their preferences on individual issues and avoid multiple-election paradoxes, but it is a very demanding assumption [Xia *et al.*, 2011; Hodge, 2002]. Relaxing rankings to be $\mathcal{O}$ -legal maintains representation compactness without permitting arbitrary preferential dependencies. Then agents may declare their preferences over an issue's candidates once given the "previous" outcomes.

Formally, for some order $\mathcal{O} = \{o_1, \dots, o_p\}$ over the issues, the preference ranking R is called $\mathcal{O}$ -legal if, given the outcomes of the prior issues $\{f^{o_1}, f^{o_2}, \dots, f^{o_{i-1}}\}$, the relative ordering of candidates in D_{o_i} is constant for any combination of outcomes of the later issues $\{f^{o_{i+1}}, \dots, f^{o_p}\}$ [Lang and Xia, 2009]. Hence, the ranking of candidates D_{o_i} in R depends only, if at all, on outcomes of issues prior to it in $\mathcal{O}$.

The preference profile P is called $\mathcal{O}$ -legal if every ranking is $\mathcal{O}$ -legal for the same order $\mathcal{O}$; the ranking R is *separable* if it is $\mathcal{O}$ -legal for any order $\mathcal{O}$.

Example 1. Consider $p = 2$ binary issues and $n = 3$ agents with preference profile $P = (R_1, R_2, R_3)$ such that:

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

The truthful vote profile $a = ((1, 0), (1, 1), (0, 0))$ consists of each agent's most preferred alternative. The score tuple is then $s(a) = \{(1, 2), (2, 1)\}$, so with plurality $f(a) = (1, 0)$.

Note that R_1 is $\mathcal{O}$ -legal for $\mathcal{O} = \{2, 1\}$: the agent always prefers $0 \succ 1$ on the second issue, yet their preference for the first issue depends on f^2 . R_3 is separable, as the agent prefers $0 \succ 1$ on each issue independent of the other issue's outcome. R_2 is neither separable nor $\mathcal{O}$ -legal for any $\mathcal{O}$.

Furthermore, 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 tuple is $s_{-2} = \{(1, 1), (2, 0)\}$, so $s_{-2} + \hat{a}_2 = \{(2, 1), (2, 1)\}$ and $f(s_{-2} + \hat{a}_2) = (0, 0) \succ_2 (1, 0) = f(a)$.

Improvement dynamics. We implement iterative voting (IV) as introduced by Meir *et al.* [2010] and refined for uncertainty by Meir *et al.* [2014] and Meir [2015]. Given agents' truthful preferences P and an initial vote profile $a(0)$, we consider an iterative process of vote profiles $a(t) = (a_1(t), \dots, a_n(t))$ 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 applying a specified *response function* $g_j : \mathcal{D}^n \rightarrow \mathcal{D}$. Each agent's response implicitly depends on their preferences and belief about the current *real* vote profile, but they are not aware of others' private preferences. All other votes besides j 's remain unchanged.

A scheduler is broadly defined as a mapping from sequences of vote profiles to an agent with an improvement step in the latest vote profile [Apt and Simon, 2015]. An improvement step must be selected if one exists, and a vote profile where no improvement step exists (i.e., $g_j(a) = a_j \forall j \leq n$) is called an *equilibrium*. The literature on game dynamics considers different types of response functions, schedulers, initial profiles, and other assumptions (see e.g., Fudenberg and Levine [2009], Marden *et al.* [2007], Bowling [2005], Young [1993], and Meir *et al.* [2017]). This means that there are multiple levels in which a voting rule may guarantee con-

vergence to an equilibrium. In this work, we study two response functions: *best response* (BR), without uncertainty, and *local dominance improvements* (LDI), with uncertainty. For both dynamics, we restrict agents to only changing their vote on a single *current* issue $i \in \mathcal{P}$ per round, as determined by the scheduler ϕ . We therefore have the following form of convergence, as described by Kukushkin [2011], Monderer and Shapley [1996], and Milchtaich [1996]:

Definition 1. An IV dynamic has the restricted-finite improvement property if every improvement sequence is finite from any initial vote profile for a given response function.

Under BR dynamics, agents know the real score tuple $s(a)$.

Definition 2 (Best response). Given the vote profile a , $g_j(a) := \hat{a}_j$ which yields agent j 's most preferred outcome of the set $\{f(a_{-j}, \tilde{a}_j) : \tilde{a}_j^i \in D_i, \tilde{a}_j^k = a_j^k \forall k \neq i\}$ unless there is no change in the outcome; then $g_j(a) := a_j$.

LDI dynamics are based on the notions of *strict uncertainty* and *local dominance* [Conitzer *et al.*, 2011; Reijngoud and Endriss, 2012]. Let $S \subseteq \times_{i=1}^p \mathbb{N}^{|D_i|}$ be a set of score tuples that, informally, describes agent j 's uncertainty about the real score tuple $s(a)$. An LDI step to a prospective vote $\hat{a}_j$ is one that is weakly better off than their original a_j for every $v \in S$ and strictly better off for some $v \in S$, as follows.

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

Definition 4 (Local dominance improvement). Given the vote profile a and agent j , let LD_j^i be the set of votes that S -dominate a_j , only differ from a_j on the i^{th} issue, and are not themselves S -dominated by any other vote differing from a_j only on the i^{th} issue. Then $g_j(a) = a_j$ if $LD_j^i = \emptyset$ and $\hat{a}_j \in LD_j^i$ otherwise.

This definition distinguishes from (weak) LDI in Meir [2015] in that agents may change their votes consecutively but only on different issues. Note that S -dominance is transitive and antisymmetric, but not complete, so an agent j may not have an improvement step. To fully define the model, we need to specify S for every a . For example, if $S = \{s(a_{-j})\}$ and each j has no uncertainty about the real score tuple, then LDI coincides with BR and an equilibrium coincides with *Nash equilibrium*. Therefore, LDI broadens BR dynamics.

Distance-based uncertainty. Agents in the single-issue model construed their uncertainty sets using *distance-based uncertainty*, in which all score vectors close enough to the current profile were believed possible [Meir *et al.*, 2014; Meir, 2015]. We adapt this to the multi-issue setting by assuming agents uphold candidate-wise distance-based uncertainty over score vectors for each issue independently.

For any issue $i \in \mathcal{P}$, let $\delta(s^i(a), \tilde{s}^i(a))$ be a distance measure for score vectors for any vote profile a . This measure is *candidate-wise* if it can be written as $\delta(s^i(a), \tilde{s}^i(a)) = \max_{c \in D_i} \hat{\delta}(s^i(c; a), \tilde{s}^i(c; a))$ for some monotone function $\hat{\delta}$. For example, the ℓ_∞ metric, where $\hat{\delta}(s, \tilde{s}) = |s - \tilde{s}|$, is candidate-wise.

Given the vote profile a and issue $i \in \mathcal{P}$, we model agent j 's uncertainty about the adjusted score vector s_{-j}^i by the *uncertainty score set* $\tilde{S}_{-j}^i(s; r_j^i) := \{v^i : \delta(v^i, s_{-j}^i) \leq r_j^i\}$ with an *uncertainty parameter* r_j^i . That is, given other votes a_{-j}^i , agent j is not sure what the real score vector is within $\tilde{S}_{-j}^i(s; r_j^i)$. We define $\tilde{S}_{-j}(s, r_j) := \times_{i=1}^p \tilde{S}_{-j}^i(s; r_j^i)$ for $r_j = (r_j^i)_{i \in \mathcal{P}}$, and drop the parameters if the context is clear.

Example 2. Consider $p = 2$ binary issues and $n = 13$ agents with the vote profile a defined such that seven agents vote $(0, 0)$, three agents vote $(1, 1)$, two agents vote $(1, 0)$, and the last agent, which we label j , votes $a_j = (0, 1)$. The score tuple is then $s(a) = \{(8, 5), (9, 4)\}$, so $f(a) = (0, 0)$.

Under BR dynamics, j has complete information about $s(a)$ and can compute $s_{-j}(a) = \{(7, 5), (9, 3)\}$. Clearly, no prospective vote $\hat{a}_j$ can change the outcome $f(a_{-j}, \hat{a}_j)$.

Under LDI dynamics, agent j has incomplete information about $s(a)$. Suppose that j uses the ℓ_∞ uncertainty metric with uncertainty parameters $(r_j^1, r_j^2) = (1, 1)$. By the above definitions, the uncertainty score set for issue $i \in \{1, 2\}$ is

$$\begin{aligned} \tilde{S}_{-j}^i(s; r_j^i) &= \{v^i : |v^i - s_{-j}^i| \leq r_j^i\} \\ &= \{(6, 7, 8) \times (4, 5, 6)\} \times \{(8, 9, 10) \times (2, 3, 4)\} \end{aligned}$$

which is a bandwidth of $r_j^i = 1$ around each real score s_{-j}^i . Finally, consider the prospective vote $\hat{a}_j = (1, 1)$. Then

$$\tilde{S}_{-j} + \hat{a}_j = \{(6, 7, 8) \times (5, 6, 7)\} \times \{(8, 9, 10) \times (3, 4, 5)\}$$

so that $\{f(v + \hat{a}_j) : v \in \tilde{S}_{-j}\} = \{(0, 0), (1, 0)\}$.

3 Best-Response Dynamics

Given the vote profile a , consider agent j changing their vote a_j on issue i to the prospective vote $\hat{a}_j$. Under BR dynamics, without uncertainty, j changes their vote only if they can feasibly improve the outcome $f(a)$ to one more favorable with respect to R_j . This happens under two conditions. First, j must be *pivotal* on the i^{th} issue, meaning that changing their vote will necessarily change the outcome. Second, j must be preferential to change i by voting for $\hat{a}_j^i$ over a_j^i given the outcomes of the other issues $\mathcal{P} \setminus \{i\}$. Agent j 's preferences are always well-defined since they know every issue's real outcome. Thus BR dynamics behave similar to the single-issue setting, which we recall converges [Meir *et al.*, 2010]. However, in the multi-issue setting, agents' preferences on each issue may change as other issues' outcomes change. This entails the possibility of a cycle, as declared in the following proposition and proved with the subsequent example.

Proposition 1. *BR dynamics for multiple issues may not converge, even if issues are binary.*

Example 3. Let there be $p = 2$ binary issues and $n = 3$ agents without uncertainty and the following preferences:

$$\begin{aligned} R_1 &: (0, 1) \succ_1 (1, 1) \succ_1 (1, 0) \succ_1 (0, 0), \\ R_2 &: (0, 0) \succ_2 (0, 1) \succ_2 (1, 1) \succ_2 (1, 0), \text{ and} \\ R_3 &: (1, 0) \succ_3 (1, 1) \succ_3 (0, 0) \succ_3 (0, 1). \end{aligned}$$

Table 1 demonstrates a cycle via BR dynamics from the truthful vote profile $a(0)$. The order of improvement steps is $j = (1, 2, 1, 2)$. No other BR step is possible from any profile in the cycle, so no agent scheduler can lead to convergence.

Agent j	$a_j(0)$	$a_j(1)$	$a_j(2)$	$a_j(3)$
1	(0, 1)	(1, 1)	(1, 1)	(0, 1)
2	(0, 0)	(0, 0)	(0, 1)	(0, 1)
3	(1, 0)	(1, 0)	(1, 0)	(1, 0)
$f(a)$	(0, 0)	(1, 0)	(1, 1)	(0, 1)

Table 1: Agents' votes for $a(0)$ (truthful), $a(1)$, $a(2)$, and $a(3)$.

4 Local Dominance Improvement Dynamics

LDI dynamics broadens best-response since agents' uncertainty score sets contain the true score tuple, by definition, but it is initially unclear how uncertainty affects the possibility of cycles. Seemingly, greater uncertainty over an agent's current issue increases the possibility of having LDI steps over that issue, whereas greater uncertainty over other issues decreases this possibility. We demonstrate in Section 4.1 below that this relationship holds only for binary issues, but it does not eliminate the possibility of cycles, as declared in the following proposition and proved with Example 5 in Appendix B.

Proposition 2. *LDI dynamics with multiple issues may not converge, even if agents have the same constant uncertainty parameters and issues are binary.*

This finding contrasts convergence guaranteed in the single-issue setting with uncertainty [Meir, 2015]. After explaining the effect of uncertainty on LDI steps, we conclude the section with two model refinements that prove sufficient to guarantee convergence for binary issues: $\mathcal{O}$ -legal preferences and a form of dynamic uncertainty.

4.1 Effect of Uncertainty on LDI Steps

Given the vote profile a among binary issues, consider agent j changing their vote a_j on issue i to the prospective vote $\hat{a}_j$. Under LDI dynamics, j changes their vote only if two conditions hold, similar to BR dynamics: if (I) they believe they may be pivotal on issue i and (II) they can improve the outcome with respect to R_j . Notice that if the agent is pivotal on the binary issue i with respect to an uncertainty parameter r_j^i , it is pivotal with respect to all larger parameters $\tilde{r}_j^i : \tilde{r}_j^i > r_j^i$ over i . Furthermore, recall that j 's preference over candidates of issue i may depend on the outcomes of other issues, which j may be uncertain about. It stands to reason that the more uncertainty j has over other issues, the less clarity the agent has over their own preference for issue i 's candidates.

We realize the following monotonic relationship between the magnitude of agents' uncertainty parameters and whether they have an LDI step over an issue: increasing uncertainty on issue i may only add LDI steps over issue i but may only eliminate LDI steps over each other issue. This is stated technically in the following proposition. First, we define three uncertainty parameters α_j , r_j , and β_j such that: r_j and α_j only differ on issue $k \neq i$ such that $r_j^k < \alpha_j^k$; r_j and β_j only differ on issue i such that $r_j^i < \beta_j^i$. Let $LD_j^i(r_j)$ denote agent j 's possible LDI steps as in Definition 4 with respect to the uncertainty parameter r_j .

Proposition 3. *Given binary issues, consider agent j changing their vote on issue i in vote profile a with one of three*

uncertainty parameters – α_j , r_j , or β_j – as defined above. Then $LD_j^i(\alpha_j) \subseteq LD_j^i(r_j) \subseteq LD_j^i(\beta_j)$.

The theorem is proved in two parts by demonstrating that if a vote $\hat{a}_j \tilde{S}_{-j}(a; \alpha_j)$ -dominates a_j , then it must hold that $\hat{a}_j \tilde{S}_{-j}(a; r_j)$ -dominates a_j ; likewise, this implies that $\hat{a}_j \tilde{S}_{-j}(a; \beta_j)$ -dominates a_j . Each of these relationships arise as a result of $\tilde{S}_{-j}(a; r_j) \subseteq \tilde{S}_{-j}(a; \alpha_j)$ and $\tilde{S}_{-j}(a; r_j) \subseteq \tilde{S}_{-j}(a; \beta_j)$. This is sufficient to prove since issues are binary. The full proof may be found in Appendix A.

We find that this relationship between agents' uncertainty parameters and their LDI steps is not monotonic in the generalized case of multi-candidate issues, as Example 6 in Appendix B demonstrates that different sets of prospective votes LD_j^i may not be comparable for an agent j with different uncertainty parameters even from the same vote profile a .

4.2 Strategic Responses and $\mathcal{O}$ -legal Preferences

We are motivated by observing in Examples 3 and 5 that cycles appear due to agents' interdependent preferences among the issues. Specifically, in Table 1, a cycle is formed as agents 1 and 2 switch their preferences among candidates for one issue when the other issue changes outcomes, and this holds for opposite issues. It therefore stands to reason that eliminating interdependent preferences by fixing agents with a $\mathcal{O}$ -legal preference profile would guarantee convergence.

We prove that this is the case in Theorem 1. To state this result technically, we first introduce a characterization about agents' strategic responses, extending a lemma from Meir [2015] to the multi-issue setting.

Definition 5. Agent j believes a candidate c on issue i is a possible winner if there is some score vector where c wins:

$W_j^i(a) := \{c \in D_i : \exists v \in \tilde{S}_{-j}(a; r_j) \text{ s.t. } f^i(v + a_j) = c\}$
In contrast, j calls c a potential winner if there is some score vector in which they can vote to make c win:

$H_j^i(a) = \{c \in D_i : \exists v \in \tilde{S}_{-j}(a; r_j) \text{ and } \hat{a}_j \text{ s.t. } \hat{a}_j^i = c, \hat{a}_j^k = a_j^k \forall k \neq i, \text{ s.t. } f^i(v + \hat{a}_j) = c\}$. The set of real potential winners is denoted: $H_0^i(a) = \{c \in D_i : f^i(s_{-j} + \hat{a}_j) = c \text{ where } \hat{a}_j^i = c, \hat{a}_j^k = a_j^k \forall k \neq i\}$.

By this definition, $W_j^i(a) \subseteq H_j^i(a)$.² Denote by $W^{-i}(a; r_j) = \times_{k \in \mathcal{P} \setminus \{i\}} W_j^k(a)$ the set of possible winning candidates on all issues besides i , from agent j 's perspective with uncertainty parameter r_j .

Lemma 1. Consider an LDI step $a_j \xrightarrow{j} \hat{a}_j$ over issue i from vote profile a by agent j with uncertainty parameter r_j . Then either (1) $a_j^i \notin H_j^i(a)$; or for every combination of possible winners in $W^{-i}(a; r_j)$, either (2) $a_j^i \prec_j b$ for all $b \in H_j^i(a)$ or (3) $r_j^i = 0$, $\{a_j^i, \hat{a}_j^i\} \subseteq H_0^i(a)$ and $\hat{a}_j^i \succ_j a_j^i$.

The proof of this lemma directly follows that of Lemma 3 in Meir [2015]; see Appendix A for the full proof.

²Without uncertainty, $H_j^i(a)$ (or $H_j^i(a) \cup \{a_j^i\}$ if adding a vote to a_j^i makes it win) is also known as the *chasing set* (excluding $f(a)$) [Rabinovich et al., 2015] or *potential winner set* (including $f(a)$) [Kavner and Xia, 2021] on issue i . $H_j^i(a)$ coincides with Meir et al. [2014] and Meir [2015]'s definition of a possible winner " $W_j(s)$."

Theorem 1. LDI dynamics converge over binary issues when all agents have $\mathcal{O}$ -legal preferences for the common order $\mathcal{O}$.

Proof. Fix an initial vote profile $a(0)$. Suppose for contradiction that there is a cycle among the vote profiles $\mathcal{C} = \{a(t_1), \dots, a(t_T)\}$, where $a(t_T + 1) = a(t_1)$ and $a(t_1)$ is reachable from $a(0)$ via LDI dynamics. Let i be the highest order issue in $\mathcal{O}$ for which any agent changes their vote in $\mathcal{C}$.

Let $t^* \in [t_1, t_T]$ be the first round that some agent j takes an LDI step on issue i , where $a_j \xrightarrow{j} \hat{a}_j$ from vote profile $a(t^*)$; let $t^{**} \in (t^*, t_T]$ be the last round that j switches their vote on i back to a_j^i . It must be the case that $a_j^i \in H_j^i(a(t^*))$, since issues are binary and otherwise, $|H_j^i(a(t^*))| = 1$ and j would not have an improvement step. Hence by Lemma 1, $\hat{a}_j^i \succ_j a_j^i$ for every combination of possible winners in $W^{-i}(a(t^*); r_j)$. Likewise, on round t^{**} , $a_j^i \succ_j \hat{a}_j^i$ for every combination of possible winners in $W^{-i}(a(t^{**}); r_j)$. Thus for some issue k and outcomes $x, y \in \{0, 1\}$, $x \neq y$, we have $W_j^k(a(t^*)) = \{x\}$ and $W_j^k(a(t^{**})) = \{y\}$.

Since j has $\mathcal{O}$ -legal preferences, k must be prior to issue i in the order $\mathcal{O}$. However, no agent changed their vote on issue k between rounds t^* and t^{**} so it must be that $x \in W_j^k(a(t^{**}))$, even if j 's uncertainty parameters changed. This forms a contradiction, so no such cycle can exist. $\square$

The intuition behind Theorem 1 is that as an LDI sequence develops, there is some "foremost" issue i in which no LDI step takes place on any issue prior to i in the order $\mathcal{O}$. Agents' relative preferences for the candidates in i are fixed because their preferences are $\mathcal{O}$ -legal: score vectors for issues prior to i in $\mathcal{O}$ do not change, while scores of issues afterward do not affect agents' preferences for i . Hence, agents' improvement steps over the issue i converge, whereas any cycle must have a sub-sequence of vote profile whose votes for issue i cycles.

Note that $\mathcal{O}$ -legality is not necessary for convergence, as BR dynamics induced from the truthful vote profile in Example 1 converge. Although $\mathcal{O}$ -legality is a strict assumption, loosening this even slightly may lead to cycles. Example 3 demonstrates a cycle in which each agent has an $\mathcal{O}$ -legal ranking but orders differ between agents.

Separately, the theorem describes that LDI steps over the issue i eventually terminate, thus enabling each subsequent issue in $\mathcal{O}$ to converge. This seems to suggest that IV under $\mathcal{O}$ -legal preferences is the same as *truthful sequential voting*, where agents vote for their preferred alternative on each issue o_i given the known previous outcomes of $\{o_1, \dots, o_{i-1}\}$ [Lang and Xia, 2009]. Although the procedures' outcomes could be the same, there are two notable differences. First, the initial vote profile could have an issue whose outcome differs from the truthful sequential outcome and no agent has an improvement step on that issue. Second, depending on the scheduler, agents may not have further improvement steps over an issue intermediately before IV reaches the same outcome as in truthful sequential voting.

This convergence result does not extend to the multi-candidate case, as declared in the following proposition and proved with the subsequent example.

Proposition 4. *LDI dynamics may not converge for multiple issues, even if agents have the same constant uncertainty parameters and $\mathcal{O}$ -legal preferences for the common order $\mathcal{O}$.*

Example 4. Consider $p = 2$ issues and $n = 15$ agents who each use the ℓ_∞ uncertainty metric with common fixed uncertainty parameters $(r_j^1, r_j^2) = (2, 1) \forall j \leq n$. Label the candidates $\{0, 1\}$ and $\{a, b, c, d\}$ respectively. Agent j has preferences: if $f^1 = 0$ then $b \succ_j c \succ_j a \succ_j d$ on the second issue; otherwise $c \succ_j b \succ_j a \succ_j d$. Agent k always prefers $a \succ_k d \succ_k b \succ_k c$ on the second issue. These preferences are $\mathcal{O}$ -legal for $\mathcal{O} = \{1, 2\}$.

Define $a(0)$ so $s(a(0)) = \{(7, 8), (3, 5, 5, 2)\}$ and $a_j(0) = a_k(0) = (0, a)$. There are four LDI steps involved in this cycle: (i) $(0, a) \xrightarrow{j} (0, d)$, (ii) $(0, a) \xrightarrow{k} (0, d)$, (iii) $(0, d) \xrightarrow{j} (0, a)$, and (iv) $(0, d) \xrightarrow{k} (0, a)$. We prove these steps are valid in Appendix B.

Note that $H_j^2(a(0)) = \{a, b, c\}$ and $H_j^2(a(2)) = \{b, c, d\}$. In contrast to the single-issue setting (see Lemma 4 of Meir [2015]), agent j takes LDI steps to candidates not in the potential winning set. This results from j 's uncertainty over whether b or c is most-preferred, even as both are preferable to a and d . Hence, we get the following corollary:

Corollary 1. *LDI dynamics may not converge for plurality over a single issue for agents with partial order preferences.*

4.3 Alternating Uncertainty

In Proposition 3 we found that for binary issues, agents may have fewer LDI steps over an issue i if that issue has less uncertainty and other issues have more. This suggests that LDI steps occur from a relative lack of information about the current issue's score vector than for other issues. If agents can gather more information about the current issue before changing their vote, thereby decreasing its uncertainty relative to other issues, then they may not have an LDI step.

We therefore consider a specific form of dynamics over agents' uncertainty parameters where agents can gather this information and consider themselves pivotal only with respect to the lowered uncertainty. Agents are assumed to subsequently forget this relative information since it may be outdated by the time they change their vote again. We show in the following theorem that this eliminates cycles.

Definition 6. (Alternating Uncertainty.) Fix two parameters r_j^c, r_j^o for each agent j such that $r_j^c < r_j^o$. Define each agent j 's uncertainty parameters such that whenever they are scheduled to change their vote on issue i , j 's uncertainty for i is r_j^c and for each other issue $k \neq i$ the uncertainty is r_j^o .

Theorem 2. *Given binary issues, LDI dynamics converges for agents with alternating uncertainty.*

Proof. Fix an initial vote profile $a(0)$ and uncertainty parameters r_j^c, r_j^o for each agent $j \leq n$. Suppose for contradiction that there is a cycle among the vote profiles $\mathcal{C} = \{a(t_1), \dots, a(t_T)\}$, where $a(t_T + 1) = a(t_1)$ and $a(t_1)$ is reachable from $a(0)$ via LDI dynamics. Without loss of generality, suppose all issues and agents are involved in the cycle.

Consider the agent j with the largest $r_j^o = \max_{u \leq n} r_u^o$. Let $t^* \in [t_1, t_T]$ be the first round that j takes an LDI step

on issue i , where $a_j \xrightarrow{j} \hat{a}_j$ from vote profile $a(t^*)$; let $t^{**} \in (t^*, t_T]$ be the last round that j switches their vote on i back to a_j^i . It must be the case that $a_j^i \in H_j^i(a(t^*))$, since issues are binary and otherwise, $|H_j^i(a(t^*))| = 1$ and j would not have an improvement step. Hence by Lemma 1, $\hat{a}_j^i \succ_j a_j^i$ for every combination of possible winners in $\mathcal{W}^{-i}(a(t^*); r_j)$. Likewise, on round t^{**} , $a_j^i \succ_j \hat{a}_j^i$ for every combination of possible winners in $\mathcal{W}^{-i}(a(t^{**}); r_j)$. Thus for some issue k and outcomes $x, y \in \{0, 1\}$, $x \neq y$, we have $W_j^k(a(t^*)) = \{x\}$ and $W_j^k(a(t^{**})) = \{y\}$.

Let $t' \in (t^*, t^{**})$ be the first round since t^* that some agent h changes their vote on issue k . Then $H_h^k(a(t')) = \{0, 1\}$. Since $W_j^k(a(t')) = W_j^k(a(t^*)) = \{x\} \subsetneq \{0, 1\}$ and distance functions are candidate-wise, $r_h^c \geq r_j^o$. This entails $r_h^o > r_j^o$ by definition of alternating uncertainty, which contradicts the assertion that j is the agent u with the largest r_u^o . $\square$

This convergence result does not extend to the multi-candidate case, as Example 4 also covers this setting.

5 Experiments

Our computational experiments investigate the effects of uncertainty and numbers of binary issues and agents on LDI dynamics. Specifically, we ask how often truthful vote profiles are themselves in equilibrium, how often LDI dynamics do not converge, and the path length to equilibrium given that LDI dynamics do converge. Our inquiry focuses on whether cycles are commonplace in practice even though convergence is not guaranteed.

We answer these questions for a broad cross-section of inputs, with $n \in \{7, 11, 15, 19\}$ agents, $p \in \{2, 3, 4, 5\}$ binary issues, and $r \in \{0, 1, 2, 3\}$ uncertainty that is constant for all agents, issues, and rounds. We generate 10,000 preference profiles for each combination by sampling agents' preferences uniformly and independently at random. We simulate LDI dynamics from the truthful vote profile using a scheduler that selects profiles uniformly at random from the set of valid LDI steps among all agents and issues. If there are no such steps, we say the sequence has converged. Otherwise, we take 50,000 rounds as a sufficiently large stopping condition to declare the sequence has cycled.

Our results are presented in Figures 1 – 3 with respect to n . As uncertainty is introduced and r increases, given $p = 5$, the availability of LDI steps diminishes significantly from the initial vote profile (Figure 1) and throughout the dynamics to eliminate (almost) all cycles and shorten the path length to convergence (Figure 3). Figure 2 presents the number of initial vote profiles whose LDI sequence cycles for $r = 0$, given that they are not themselves in equilibrium; only five of the sampled $r \geq 1$ profiles' sequences cycle. Therefore, cycles with uncertainty are the exception rather than the norm.

These findings corroborate our theoretical analysis. As uncertainty increases, more issues are perceived by agents to have more than one possible winner. Since issues are interdependent for many preference rankings, fewer agents have LDI steps. On the other hand, as n increases, more agents




Figure 1: Percentage of truthful vote profiles not in equilibrium as n increases.




Figure 2: Number of truthful vote profiles whose LDI sequences cycle as n increases.

have rankings without these interdependencies, thus increasing the availability of LDI steps.

As an additional inquiry, we studied how IV affects the quality of outcomes by comparing the social welfare of equilibrium to truthful vote profiles.³ We find in Figure 4 that IV improves average welfare, but at a rate decreasing in r . This finding agrees with experiments by Bowman *et al.* [2014] and Grandi *et al.* [2022], suggesting that IV may reduce multiple-election paradoxes by helping agents choose better outcomes. However, further work will be needed to generalize this conclusion, as it contrasts experiments of single-issue IV by Meir *et al.* [2020] and Koolyk *et al.* [2017].

6 Discussion and Open Questions

We have introduced a novel model of strategic behavior in iterative voting (IV) for multiple issues under uncertainty. We find that for binary issues, the existence of cycles hinges on the interdependence of issues in agents' preference rankings. Specifically, once an agent j takes an LDI step on an issue i , they only subsequently revert their vote if their preference for i changes. This occurs if the possible winning candidates among other issues that affect j 's preference for i change. Without this interdependence, agents' preference over indi-

³Measured by the percent change in Borda welfare. The *Borda utility* of outcome a for ranking R is 2^p minus the index of a 's position in R ; the *Borda welfare* is the sum of utilities across agents.




Figure 3: Average number steps for LDI sequences to converge as n increases; log scale; 95% CI (too small to show).




Figure 4: Average percent change in Borda welfare as n increases; 95% CI (too small to show).

vidual issues change only finite times, so LDI dynamics converge (Theorem 1). We also find that as uncertainty increases over issues other than the one agents are changing, fewer preference rankings admit LDI steps, eliminating cycles (Theorem 2). Convergence does not extend to multi-candidate issues since LDI dynamics may cycle if agents only have partial order preference information (Corollary 1). Our experiments confirm that convergence is practically guaranteed with uncertainty, despite its possibility, and suggests IV improves agents' social welfare over truthful outcomes.

There are several open directions for future work. First, our empirical study was limited by sampling agents' preferences from the impartial culture preference distribution and measuring additive social welfare (see e.g., [Tsetlin *et al.*, 2003; Sen, 1999]). Proving IV's welfare properties for more realistic preference distributions and welfare functions may follow the research in dynamic price of anarchy of Brânzei *et al.* [2013] and Kavner and Xia [2021] and in smoothed analysis of Xia [2020]. Second, IV is useful for protecting agents' privacy, in part, as it does not explicitly reveal agents' truthful preferences. However, agents implicitly reveal partial information through their improvement steps. Studying IV when agents account for others' preferences based on current information is an interesting open direction. A third direction is detailing the axiomatic properties of multi-issue IV rules that may be inherited by the decision rule used locally on each issue, as studied for sequential voting by Xia *et al.* [2011].

Acknowledgments

Lirong Xia acknowledges NSF #1453542, #2007994, and #2106983, and a Google Research Award for support. Reshef Meir is supported by the Israel Science Foundation (ISF; Grant No. 2539/20).

References

- [Airiau and Endriss, 2009] Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *International Conference on Algorithmic Decision Theory*, pages 38–49. Springer, 2009.
- [Apt and Simon, 2015] Krzysztof R Apt and Sunil Simon. A classification of weakly acyclic games. *Theory and Decision*, 78(4):501–524, 2015.
- [Bowling, 2005] Michael Bowling. Convergence and no-regret in multiagent learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, pages 209–216, Vancouver, Canada, 2005.
- [Bowman et al., 2014] Clark Bowman, Jonathan K Hodge, and Ada Yu. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and Decision*, 77(1):111–124, 2014.
- [Brânzei et al., 2013] Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, pages 138–144, 2013.
- [Chopra et al., 2004] Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In *European Workshop on Logics in Artificial Intelligence*, pages 18–30. Springer, 2004.
- [Conitzer et al., 2009] Vincent Conitzer, Jérôme Lang, and Lirong Xia. How hard is it to control sequential elections via the agenda? In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 103–108, Pasadena, CA, USA, 2009.
- [Conitzer et al., 2011] Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 638–643, San Francisco, CA, USA, 2011.
- [Desmedt and Elkind, 2010] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic Commerce*, pages 347–356, 2010.
- [Elkind et al., 2020] Edith Elkind, Umberto Grandi, Francesca Rossi, and Arkadii Slinko. Cognitive hierarchy and voting manipulation in k-approval voting. *Mathematical Social Sciences*, 108:193–205, 2020.
- [Endriss et al., 2016] Ulle Endriss, Svetlana Obraztsova, Maria Polukarov, and Jeffrey S. Rosenschein. Strategic voting with incomplete information. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, page 236–242. AAAI Press, 2016.
- [Fudenberg and Levine, 2009] Drew Fudenberg and David K Levine. Learning and equilibrium. *Annu. Rev. Econ.*, 1(1):385–420, 2009.
- [Grandi et al., 2013] Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *International Conference on Algorithmic Decision Theory*, pages 181–192. Springer, 2013.
- [Grandi et al., 2019] Umberto Grandi, Daniel Hughes, Francesca Rossi, and Arkadii Slinko. Gibbard–Satterthwaite games for k-approval voting rules. *Mathematical Social Sciences*, 99:24–35, 2019.
- [Grandi et al., 2022] Umberto Grandi, Jérôme Lang, Ali I Ozkes, and Stéphane Airiau. Voting behavior in one-shot and iterative multiple referenda. *Social Choice and Welfare*, pages 1–35, 2022.
- [Hodge, 2002] Jonathan K Hodge. *Separable preference orders*. PhD thesis, Western Michigan University, 2002.
- [Kavner and Xia, 2021] Joshua Kavner and Lirong Xia. Strategic behavior is bliss: iterative voting improves social welfare. *Advances in Neural Information Processing Systems*, 34:19021–19032, 2021.
- [Koolyk et al., 2017] Aaron Koolyk, Tyrone Strangway, Omer Lev, and Jeffrey S. Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 273–279, 2017.
- [Kukushkin, 2011] Nikolai S Kukushkin. Acyclicity of improvements in finite game forms. *International Journal of Game Theory*, 40(1):147–177, 2011.
- [Lacy and Niou, 2000] D. Lacy and E. Niou. A problem with referenda. *Journal of Theoretical Politics*, 12(1):5–31, 2000.
- [Lang and Xia, 2009] Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.
- [Lang and Xia, 2016] Jérôme Lang and Lirong Xia. Voting in Combinatorial Domains. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel Procaccia, editors, *Handbook of Computational Social Choice*, chapter 9. Cambridge University Press, 2016.
- [Lang, 2007] Jérôme Lang. Vote and aggregation in combinatorial domains with structured preferences. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1366–1371, Hyderabad, India, 2007.
- [Lev and Rosenschein, 2012] Omer Lev and Jeffrey S Rosenschein. Convergence of iterative voting. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 611–618, 2012.
- [Marden et al., 2007] Jason R. Marden, Gürdal Arslan, and Jeff S. Shamma. Regret based dynamics: Convergence in

- weakly acyclic games. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '07, New York, NY, USA, 2007. Association for Computing Machinery.
- [Meir *et al.*, 2010] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Convergence to Equilibria of Plurality Voting. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 823–828, Atlanta, GA, USA, 2010.
- [Meir *et al.*, 2014] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A Local-Dominance Theory of Voting Equilibria. In *Proceedings of the 15th ACM Conference on Electronic Commerce*, pages 313–330, Palo Alto, CA, USA, 2014.
- [Meir *et al.*, 2017] Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nicholas R. Jennings. Iterative voting and acyclic games. *Artificial Intelligence*, 252:100–122, 2017.
- [Meir *et al.*, 2020] Reshef Meir, Kobi Gal, and Maor Tal. Strategic voting in the lab: compromise and leader bias behavior. *Autonomous Agents and Multi-Agent Systems*, 34(1):1–37, 2020.
- [Meir, 2015] Reshef Meir. Plurality Voting under Uncertainty. In *Proceedings of the 29th AAAI on Artificial Intelligence*, 2015.
- [Meir, 2018] Reshef Meir. *Strategic Voting*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool, 2018.
- [Milchtaich, 1996] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
- [Monderer and Shapley, 1996] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [Myerson and Weber, 1993] Roger B. Myerson and Robert J. Weber. A theory of voting equilibria. *American Political Science Review*, 87(1):102–114, 1993.
- [Obraztsova *et al.*, 2015] Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R. Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, page 993–999, 2015.
- [Rabinovich *et al.*, 2015] Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey Rosenschein. Analysis of equilibria in iterative voting schemes. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*. AAAI Press, 2015.
- [Reijngoud and Endriss, 2012] Annemieke Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, volume 1, pages 635–644, 06 2012.
- [Reyhani and Wilson, 2012] Reyhaneh Reyhani and Mark C. Wilson. Best reply dynamics for scoring rules. In *Proceedings of the 20th European Conference on Artificial Intelligence*, ECAI'12, page 672–677, NLD, 2012. IOS Press.
- [Sen, 1999] Amartya Sen. The Possibility of Social Choice. *American Economic Review*, 89(3):349–378, 1999.
- [Tsetlin *et al.*, 2003] Ilia Tsetlin, Michel Regenwetter, and Bernard Grofman. The impartial culture maximizes the probability of majority cycles. *Social Choice and Welfare*, 21(3):387–398, 2003.
- [Van Ditmarsch *et al.*, 2013] Hans Van Ditmarsch, Jérôme Lang, and Abdallah Saffidine. Strategic voting and the logic of knowledge. *arXiv preprint arXiv:1310.6436*, 2013.
- [Xia and Conitzer, 2010] Lirong Xia and Vincent Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 921–926, Atlanta, GA, USA, 2010.
- [Xia *et al.*, 2011] Lirong Xia, Vincent Conitzer, and Jérôme Lang. Strategic sequential voting in multi-issue domains and multiple-election paradoxes. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 179–188, San Jose, CA, USA, 2011.
- [Xia, 2020] Lirong Xia. The smoothed possibility of social choice. *Advances in Neural Information Processing Systems*, 33:11044–11055, 2020.
- [Young, 1993] H. Peyton Young. The Evolution of Conventions. *Econometrica*, 61(1):57–84, January 1993.
- [Zou *et al.*, 2015] James Zou, Reshef Meir, and David Parkes. Strategic voting behavior in doodle polls. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW '15*, page 464–472, New York, NY, USA, 2015.