

Strategic Sequential Voting in Multi-Issue Domains and Multiple-Election Paradoxes

Lirong Xia
Department of Computer Science
Duke University
Durham, NC 27708, USA
lxia@cs.duke.edu

Vincent Conitzer
Department of Computer Science
Duke University
Durham, NC 27708, USA
conitzer@cs.duke.edu

Jérôme Lang
LAMSADE
Université Paris-Dauphine
75775 Paris Cedex, France
lang@lamsade.dauphine.fr

Abstract

In many settings, a group of agents must come to a joint decision on multiple issues. In practice, this is often done by voting on the issues in sequence. In this paper, we model sequential voting in multi-issue domains as a complete-information extensive-form game, in which the agents are perfectly rational and their preferences are common knowledge. In each step, the voters simultaneously vote on one issue, and the order of the issues is determined before the process. We call this model *strategic sequential voting*.

We focus on domains with binary issues, so that this process leads to a unique outcome under a natural solution concept. We show several *multiple-election paradoxes* in strategic sequential voting: there exists a profile for which the winner under strategic sequential voting is ranked nearly at the bottom in all votes, and the winner is Pareto-dominated by almost every other alternative. We also show that changing the order of the issues cannot completely prevent such paradoxes. We also study paradoxes for strategic sequential voting in which the profiles satisfy domain restrictions such as separability, lexicographicity or \mathcal{O} -legality. Finally, we study other common voting rules (from a non-strategic perspective). For some of them, we show that there exist paradoxes that are similar to the multiple-election paradoxes, and for the others, we show that there are no such paradoxes.

1 Introduction

In a traditional voting system, each voter is asked to report a linear order over the alternatives to represent her preferences. Then, a *voting rule* is applied to the resulting profile of reported preferences, to select a winning alternative.

In practice, the set of alternatives often has a *multi-issue* structure. That is, there are p issues $\mathcal{I} = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$, and each issue can take values in a *local domain*. In other words, the set of alternatives is the Cartesian product of the local domains. For example, in *multiple referenda*, the inhabitants of a local district are asked to vote on multiple inter-related issues [4]. Another example is *voting by committees*, in which the voters select a subset of objects [1], where each object can be seen as a binary issue.

Voting in multi-issue domains has been extensively studied by economists, and more recently has attracted the attention of computer scientists. Previous work has focused on proposing a natural and compact *voting language* for the agents to represent their preferences, as well as designing a sensible voting rule to make decisions based on preferences represented in such languages. A natural approach is to let voters vote on the issues separately, in the following way. For each issue (simultaneously, not sequentially), each voter reports her preferences for that issue, and then, a *local rule* is used to select the winning value that the issue will take. This voting process is called *issue-by-issue* or *seat-by-seat* voting. Computing the winner for issue-by-issue voting rules is easy, and it only requires a modest amount of communication from the agents to the mechanism. Nevertheless, issue-by-issue voting has some drawbacks. First, a voter may feel uncomfortable expressing her preferences over one issue independently of the values that the other issues take [12]. It has been pointed out that issue-by-issue voting avoids this problem if the voters' preferences are *separable* (that is, for any issue i , regardless of the values for the other issues, the voter's preferences over issue i are always the same) [11]. Second, *multiple-election paradoxes* arise in issue-by-issue voting [4, 11, 17, 19]. In models that do not consider strategic (game-theoretic) voting, previous works have shown several types of paradoxes: sometimes the winner is a Condorcet loser; sometimes the winner is Pareto-dominated by another alternative (that is, that alternative is preferred to the winner in all votes); and sometimes the winner is ranked in a very low position by all voters.

A way to (partly) escape these paradoxes consists in organizing the multiple elections *sequentially*: given an order \mathcal{O} over all issues (without loss of generality, we take \mathcal{O} to be $\mathbf{x}_1 > \dots > \mathbf{x}_n$), the voters first vote on issue \mathbf{x}_1 ; then, the value collectively chosen for \mathbf{x}_1 is determined using some voting rule and broadcast to the voters, who then vote on issue \mathbf{x}_2 , and so on. When the issues are all binary, it is natural to choose the majority rule at each stage (plus, in the case of an even number of voters, some tie-breaking mechanism). Such processes are conducted in many real-life situations, such as recruiting committees (suppose there is a full professor position and an assistant professor position to be filled; then, it is realistic to expect that the committee will first decide who gets the full professor position); or, at the executive meeting of the co-owners of a building, important decisions (whether a lift should be installed or not, how much money should be spent to repair the roof) are usually taken before minor decisions. In each of these cases, it is clear that the decision made on one issue influences the votes on later issues, thus the order in which the issues are decided potentially has a strong influence on the final outcome. Now, if voters are assumed to know the preferences of other voters well enough, then we can expect them to vote strategically at each step, forecasting the outcome at later steps conditional on the outcomes at earlier steps.

Our contributions.

In this paper, we analyze the complete-information game-theoretic model of sequential voting that we illustrated by the example above. This model applies to any preferences that the agents may have (not just O -legal ones), though they must be strict orders.

We focus on voting in binary multi-issue domains, that is, for any $i \leq p$, \mathbf{x}_i must take a value in $\{0_i, 1_i\}$. This has the advantage that for each issue, we can use the majority rule as the local rule for that issue. We use a game-theoretic model to analyze outcomes that result from sequential voting. Specifically, we model the sequential voting process as a p -stage complete-information game, as follows. There is an order \mathcal{O} over all issues (w.l.o.g., $\mathcal{O} = \mathbf{x}_1 > \mathbf{x}_2 > \dots > \mathbf{x}_p$), which indicates the order of the issues (the order in which they will be voted on). In stage i , the voters vote on issue \mathbf{x}_i simultaneously, and the majority rule is used to choose the winning value for \mathbf{x}_i . We make the following game-theoretic assumptions: the order \mathcal{O} is common knowledge; all voters' preferences are common knowledge; and, of course, all voters are perfectly rational, which is also common knowledge.

We can solve this game by a type of backward induction: in the last (p th) stage, only two alternatives remain (corresponding to the two possible settings of the last issue), so at this point it is a (weakly) dominant strategy for each voter to vote for her more preferred alternative of the two. Then, in the second-to-last ($(p - 1)$ th) stage, there are two possible local outcomes for the $(p - 1)$ th issue; for each of them, the voters can predict which alternative will finally be chosen, because they can predict what will happen in the p th stage. Thus, the $(p - 1)$ th stage is effectively a majority election between two alternatives, and each voter will vote for her more preferred alternative; etc. We call such a procedure the *strategic sequential voting procedure (SSP)*.

Given the order over issues, this game-theoretic analysis maps every profile of (strict ordinal) preferences to a unique outcome. Since any function from profiles of preferences to alternatives can be interpreted as a voting rule, the voting rule that corresponds to SSP is denoted by $SSP_{\mathcal{O}}$ (this facilitates comparison with results where voters are simply assumed to be voting truthfully).

After the introduction of SSP, we show that, unfortunately, multiple-election paradoxes also arise under SSP. To better present our results, we introduce a parameter called the *minimax satisfaction index (MSI)*. For an election with m alternatives and n voters, it is defined in the following way. For each profile, consider the highest position that the winner obtains across all of the input rankings of the alternatives (corresponding to the most-satisfied voter); this is the *maximum satisfaction index* for this profile. Then, the minimax satisfaction index is obtained by taking the minimum over all profiles of the maximum satisfaction index. A low minimax satisfaction index means that there exists a profile in which the winner is ranked in low positions in all votes, thus there is a multiple-election paradox. Our main theorem is the following.

Theorem 1 *For any $p \in \mathbb{N}$ and any $n \geq 2p^2 + 1$, the minimax satisfaction index of SSP when there are $m = 2^p$ alternatives and n voters is $\lfloor p/2 + 2 \rfloor$. Moreover, in the profile P that we use to prove the upper bound, the winner $SSP_{\mathcal{O}}(P)$ is Pareto-dominated by $2^p - (p + 1)p/2$ alternatives.*

We note that an alternative c Pareto-dominates another alternative c' implies that c beats c' in their pairwise election. Therefore, Theorem 1 implies that the winner for SSP is almost a Condorcet loser. It follows from this theorem that SSP exhibits all three types of multiple-election paradoxes: the winner is ranked almost in the bottom in every vote, the winner is almost a Condorcet loser, and the winner is Pareto-dominated by almost every other alternative. We also show a paradox (Theorem 2) that states that there exists a profile such that for *any* order \mathcal{O} of the issues, in each input ranking, the

SSP winner is ranked almost in the bottom position. We also show that even when the voters' preferences can be represented by CP-nets that are compatible with a common order, multiple-election paradoxes still arise.

To see if there are similar paradoxes for common voting rules (when voters are assumed to vote truthfully), we calculate the minimax satisfaction index for some common voting rules, including dictatorships, positional scoring rules (including k -approval and Borda), plurality with runoff, Copeland, maximin, STV, Bucklin, ranked pairs, and (not necessarily binary) voting trees. We show that for k -approval with large k , voting trees, Copeland₀, and maximin we can find a similar paradox, and for the others there are no such paradoxes.

Related work and discussion.

Our setting is closely related to the *multi-stage sophisticated voting* [13, 14, 10]. They studied the model where the backward induction outcomes correspond to the truthful outcomes of voting trees. Therefore, our SSP is a special case of multi-stage sophisticated voting. However, their work mainly focused on general set of alternatives (while we focus on multi-issue domains), and their results are characterization of the outcomes as the outcomes in the *sophisticated voting* [7]. We, on the other hand, study the multiple-election paradoxes for SSP. Another paper that is closely related to part of this work was written by Dutta and Sen [6]. They show that social choice rules corresponding to binary voting trees can be implemented via backward induction via a sequential voting mechanism. This is closely related to a relationship that we discuss later in this paper, namely an equivalence between the outcome of strategic behavior in sequential voting over multiple binary issues, and a particular type of voting tree. It should be pointed out that the sequential mechanism that Dutta and Sen consider is somewhat different from sequential voting as we consider it—in particular, in the Dutta-Sen mechanism, one agent moves at a time, and a move often consists not of a vote, but rather of choosing the next player to move. Nevertheless, the approaches are related at a high level, though they are motivated quite differently: Dutta and Sen are interested in social choice rules corresponding to voting trees, and are trying to create sequential mechanisms that implement them via backward induction. We, on the other hand again, are primarily interested in the strategic outcome of the natural mechanism for voting sequentially over multiple issues, and use voting trees merely as a useful tool for analyzing the outcome of this process. Also, the relationship between sequential voting and voting trees takes a particularly natural form in the context of domains with multiple binary issues, as we will show. Less closely related, implementation by voting trees has previously been studied at EC: Fischer et al. [8] consider the known result that the Copeland rule (which we define later in this paper) cannot be implemented by a voting tree [16], and set out to *approximate* the the Copeland score using voting trees.

It has been pointed out that typical multiple-election paradoxes partly comes from the incompleteness of information about the preferences of the voters [11]. However, the paradoxes in this paper show that assuming that voters' preferences are common knowledge does not allow to get rid of multiple election paradoxes. Another interpretation of these results is that we may need to move beyond sequential voting to properly address voting in multi-issue domains. However, note that other approaches than sequential voting may be extremely costly in terms of communication and computation, which comes down to saying, one more time, that voting on multiple related issues is an extremely challenging problem for which probably no perfect solution exists.

2 Preliminaries

2.1 Basics of voting

Let \mathcal{X} be the set of *alternatives*, $|\mathcal{X}| = m$. A vote is a linear order over \mathcal{X} . The set of all linear orders over \mathcal{X} is denoted by $L(\mathcal{X})$. For any $c \in \mathcal{X}$ and $V \in L(\mathcal{X})$, we let $\text{rank}_V(c)$ denote the position of c in V from the top. An n -*profile* P is a collection of n votes for some $n \in \mathbb{N}$, that is, $P \in L(\mathcal{X})^n$. For any $c, d \in \mathcal{X}$ and any profile P , we say c *Pareto-dominates* d , if for any $V \in P$, c is ranked higher than d in V , that is, $c \succ_V d$. A *voting rule* r is a mapping that assigns to each profile a unique winning alternative. That is, $r : L(\mathcal{X}) \cup L(\mathcal{X})^2 \cup \dots \rightarrow \mathcal{X}$. Some common voting rules are listed below.

- *Dictatorship*. for every $n \in \mathbb{N}$ there exists a voter $j \leq n$ such that the winner is always the alternative that is ranked in the top position in V_i .

- *(Positional) scoring rules*: Given a *scoring vector*

$\vec{v} = (v(1), \dots, v(m))$, for any vote $V \in L(\mathcal{X})$ and any $c \in \mathcal{X}$, let $s(V, c) = v(j)$, where j is the rank of c in V . For any profile $P = (V_1, \dots, V_n)$, let $s(P, c) = \sum_{i=1}^n s(V_i, c)$. The rule will select $c \in \mathcal{X}$ so that $s(P, c)$ is maximized. Some examples of

positional scoring rules are *Borda*, for which the scoring vector is $(m-1, m-2, \dots, 0)$, k -*approval* (App_k , with $k \leq m$), for which the scoring vector is $(\underbrace{1, \dots, 1}_k, 0, \dots, 0)$,

plurality, for which the scoring vector is $(1, 0, \dots, 0)$, and *veto*, for which the scoring vector is $(1, \dots, 1, 0)$.

- *Copeland $_\alpha$* ($0 \leq \alpha \leq 1$): For any two alternatives c_i and c_j , we can simulate a *pairwise election* between them, by seeing how many votes prefer c_i to c_j , and how many prefer c_j to c_i ; the winner of the pairwise election is the one preferred more often. Then, an alternative receives one point for each win in a pairwise election, α points for each draw, and 0 point for each loss. The winner is the alternative that has the highest score.

- *Plurality with runoff* (Plu_o): The election has two rounds. In the first round, the alternatives are ranked from high to low according to the number of times they are ranked in the top position in the votes of the profile (that is, according to their plurality scores). Only the top two alternatives enter the second (runoff) round. In the runoff, we simulate a pairwise election between these two alternatives, and the alternative that wins the pairwise election is the winner.

- *Maximin*: Let $N(c_i, c_j)$ denote the number of votes that rank c_i ahead of c_j . The winner is the alternative c that maximizes $\min\{N(c, c') : c' \in \mathcal{X}, c' \neq c\}$.

- *STV*: The election has $m - 1$ rounds. In each round, we count for each remaining alternative how many votes rank it highest among the remaining alternatives; then, the alternative with the lowest count drops out. The last remaining alternative is the winner.

- *Bucklin*: An alternative c 's Bucklin score is the smallest number k such that more than half of the votes rank c among the top k alternatives. The winner is the alternative that has the lowest Bucklin score. If multiple alternatives have the lowest score k , then ties are broken by the number of votes that rank an alternative among the top k .

- *Ranked pairs*: This rule first creates an entire ranking of all the alternatives. $N(c_i, c_j)$ is defined as for the maximin rule. In each step, we will consider a pair of alternatives c_i, c_j that we have not previously considered; specifically, we choose the remaining pair with the highest $N(c_i, c_j)$. We then fix the ordering $c_i \succ c_j$, unless this contradicts orderings that we fixed previously (that is, it violates transitivity). We continue until we have considered all pairs of alternatives (hence we end up with a full ranking). The alternative at the top of the ranking wins.

• *Voting trees*: A voting tree is a binary tree with m leaves, where each leaf is associated with an alternative. In each round, there is a pairwise election between an alternative c_i and its sibling c_j : if a majority of voters prefers c_i to c_j , then c_j is eliminated, and c_i is associated with the parent of these two nodes; similarly, if a majority of voters prefers c_j to c_i , then c_i is eliminated, and c_j is associated with the parent of these two nodes. The alternative that is associated with the root of the tree (wins all its rounds) is the winner.

2.2 Multi-issue domains

In this paper (except Section 7), the set of all alternatives \mathcal{X} is a *binary multi-issue domain*. That is, let $\mathcal{I} = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ ($p \geq 2$) be a set of *issues*, where each issue \mathbf{x}_i takes values in a binary *local domain* $D_i = \{0_i, 1_i\}$. The set of alternatives is $\mathcal{X} = D_1 \times \dots \times D_p$, that is, an alternative is uniquely identified by its values on all issues. If $Y \subseteq \mathcal{I}$ then $D_Y = \prod_{y \in Y} D_y$.

Given a preference relation \succ in $L(\mathcal{X})$, an issue \mathbf{x}_i , and a subset of issues $W \subseteq \mathcal{I}$, let $U = \mathcal{I} \setminus (W \cup \{\mathbf{x}_i\})$; then, \mathbf{x}_i is *preferentially independent of W given U* (with respect to \succ) if for any $\vec{u} \in D_U$, any $a_i, b_i \in D_i$, and any $\vec{w}, \vec{w}' \in D_W$, $(\vec{u}, a_i, \vec{w}) \succ (\vec{u}, b_i, \vec{w})$ if and only if $(\vec{u}, a_i, \vec{w}') \succ (\vec{u}, b_i, \vec{w}')$. Informally, if we wish to find out whether changing the value of \mathbf{x}_i from a_i to b_i (while keeping everything else fixed) will make the voter better or worse off, we only need to know the values of the issues in U .

Let $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$. A preference relation \succ is *\mathcal{O} -legal* if for any $i \leq p$, \mathbf{x}_i is preferentially independent of $\{\mathbf{x}_{i+1}, \dots, \mathbf{x}_p\}$ given $\{\mathbf{x}_1, \dots, \mathbf{x}_{i-1}\}$. Informally, to find out whether a particular change in the value of an issue will make the voter better or worse off, we only need to know the values of earlier issues. A preference relation \succ is *separable* if for any $i \leq p$, \mathbf{x}_i is preferentially independent of $\mathcal{X} \setminus \{\mathbf{x}_i\}$. Informally, to find out whether a particular change in the value of an issue will make the voter better or worse off, we do not need to know the value of any other issue. A separable preference relation is *\mathcal{O} -legal* for any \mathcal{O} .

A preference relation \succ is *\mathcal{O} -lexicographic* if for any $i \leq p$, any $\vec{u} \in D_1 \times \dots \times D_{i-1}$, any $a_i, b_i \in D_i$, and any $\vec{d}_1, \vec{d}_2, \vec{e}_1, \vec{e}_2 \in D_{i+1} \times \dots \times D_p$, $(\vec{u}, a_i, \vec{d}_1) \succ (\vec{u}, b_i, \vec{e}_1)$ if and only if she prefers $(\vec{u}, a_i, \vec{d}_2) \succ (\vec{u}, b_i, \vec{e}_2)$. Informally, if a profile is *\mathcal{O} -lexicographic*, then it is *\mathcal{O} -legal*, and moreover, earlier issues are more important—that is, to compare two alternatives, it suffices to know the values of the issues up to and including the first issue \mathbf{x}_i on which they differ. (While the values of $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$ will be the same, they still matter in that they affect the preference on \mathbf{x}_i .) *\mathcal{O} -lexicographicity* and *separability* are incomparable notions.

A profile is *separable/ \mathcal{O} -lexicographic/ \mathcal{O} -legal* if it is composed of preference relations that are all *separable/ \mathcal{O} -lexicographic/ \mathcal{O} -legal*.

We can now define sequential composition of local voting rules. Given a vector of *local rules* (r_1, \dots, r_p) (where for any $i \leq p$, r_i is a voting rule on D_i), the *sequential composition* of r_1, \dots, r_p w.r.t. \mathcal{O} , denoted by $Seq_{\mathcal{O}}(r_1, \dots, r_p)$, is defined for all *\mathcal{O} -legal profiles* as follows: $Seq_{\mathcal{O}}(r_1, \dots, r_p)(P) = (d_1, \dots, d_p) \in \mathcal{X}$, where for any $i \leq p$, $d_i = r_i(P|_{\mathbf{x}_i: d_1 \dots d_{i-1}})$, where $P|_{\mathbf{x}_i: d_1 \dots d_{i-1}}$ is composed of the voters' local preferences over \mathbf{x}_i , given that the issues preceding it take values d_1, \dots, d_{i-1} . Thus, the winner is selected in p steps, one for each issue, in the following way: in step i , d_i is selected by applying the local rule r_i to the preferences of voters over D_i , conditioned on the values d_1, \dots, d_{i-1} that have already been determined for the issues that precede

\mathbf{x}_i . In this paper, we only consider the case where every r_i is the majority rule over two alternatives.

3 Strategic sequential voting

Sequential voting on multi-issue domains can be seen as a game where in each step, the voters decide whether to vote for or against the issue under consideration after reasoning about what will happen next. We make the following assumptions.

(1) All voters act strategically, and this is common knowledge.

(2) The order in which the issues will be voted upon, as well as the local voting rules used at the different steps (namely, majority), are common knowledge.

(3) All voters' preferences on the set of alternatives are common knowledge.

Assumption 1 is standard in game theory. Assumption 2 merely means that the rule has been announced. Assumption 3 (complete information) is the most significant assumption. It may be interesting to consider more general settings with incomplete information, resulting in a Bayesian game. Nevertheless, because the complete-information setting is a special case of the incomplete-information setting (where the prior distribution is degenerate), *all negative results obtained for the complete-information setting also apply to the incomplete-information setting*. That is, the restriction to complete information only strengthens negative results.

Given these assumptions, the voting process can be modeled as an game that is composed of p stages where in each stage, the voters vote simultaneously on one issue. Let \mathcal{O} be the order over the set of issues, which without loss of generality we assume to be $\mathbf{x}_1 > \dots > \mathbf{x}_p$, and P the profile of preferences over \mathcal{X} . The game is defined as follows: for each $i \leq p$, in stage i the voters vote simultaneously on issue i ; then, the value of \mathbf{x}_i is determined by the majority rule (plus, in the case of an even number of voters, some tie-breaking mechanism), and this local outcome is broadcast to all voters. (Equivalently, all voters could broadcast their votes at each stage.)

We now show how to solve the game. Because of assumptions 1 to 3, at step i the voters vote strategically, by recursively figuring out what the final outcome will be if the local outcome for \mathbf{x}_i is 0, and what it will be if it is 1. More concretely, suppose that steps 1 to $i - 1$ resulted in issues $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$ taking the values d_1, \dots, d_{i-1} , and let $\vec{d} = (d_1, \dots, d_{i-1})$. Suppose also that if \mathbf{x}_i takes the value 0_i (resp., 1_i), then, recursively, the remaining issues will take the tuple of values \vec{a} (resp., \vec{b}). Then, \mathbf{x}_i is determined by a pairwise comparison between $(\vec{d}, 0_i, \vec{a})$ and $(\vec{d}, 1_i, \vec{b})$ in the following way: if a majority of voters prefers $(\vec{d}, 0_i, \vec{a})$ over $(\vec{d}, 1_i, \vec{b})$, then \mathbf{x}_i takes the value 0_i ; in the opposite case, \mathbf{x}_i takes the value 1_i . This process, which corresponds to the strategic behavior in the sequential election, is what we call the *strategic sequential voting* procedure, and for any profile P , the winner w.r.t. the order \mathcal{O} is denoted by $SSP_{\mathcal{O}}(P)$.

As we shall see later, SSP can not only be thought of as the strategic outcome of sequential voting, but also as a voting rule in its own right. In fact, SSP corresponds to a particular balanced voting tree, as illustrated in Figure 1 for the case $p = 3$. In this voting tree, in the first round, each alternative is paired up against the alternative that differs only on the p th issue; each alternative that wins the first round is then paired up with the unique other remaining alternative that differs only on the $(p - 1)$ th and possibly the p th issue; etc.

Of course, there are many voting trees that do *not* correspond to an SSP election; this is easily seen by observing that there are only $p!$ different SSP elections (corre-

sponding to the different orders of the issues), but many more voting trees. The voting tree corresponding to the order $\mathcal{O} = x_1 > \dots > x_p$ is defined by the property that for any node v whose depth is i (where the root has depth 1), the alternative associated with any leaf in the left (resp., right) subtree of v gives the value 0_i (resp., 1_i) to x_i .

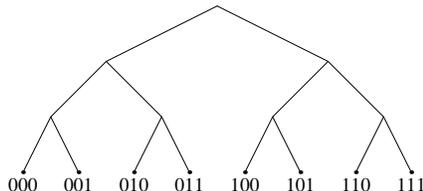


Figure 1: A voting tree that is equivalent to the strategic sequential voting procedure ($p = 3$). 000 is the abbreviation for $0_1 0_2 0_3$, etc.

4 Minimax Satisfaction Index

In the rest of this paper, we will show that strategic sequential voting on multi-issue domains is prone to paradoxes that are almost as severe as previously studied multiple-election paradoxes under models that are not game-theoretic [4, 11]. To facilitate the presentation of these results, we define an index that is intended to measure one aspect of the quality of a voting rule, called *minimax satisfaction index*.

In words, the minimum satisfaction index can be defined as follows. For each profile, consider the highest position that the winner obtains across all of the input rankings of the alternatives (corresponding to the most-satisfied voter); this is the *maximum satisfaction index* for this profile. Then, the minimax satisfaction index is obtained by taking the minimum over all profiles of the maximum satisfaction index.

Definition 1 For any voting rule r , the minimax satisfaction index (MSI) of r is defined by $MSI_r(m, n) = \min_{P \in L(\mathcal{X})^n} \max_{i \leq n} (m + 1 - \text{rank}_{V_i}(r(P)))$ where m is the number of alternatives and n is the number of voters.

The MSI of a voting rule is not the final word on it. For example, the MSI for dictatorships is m , the maximum possible value, which is not to say that dictatorships are desirable. However, if the MSI of a voting rule is low, then this implies the existence of a paradox for it, namely, a profile that results in a winner that makes all voters unhappy.

Many of the multiple-election paradoxes known so far implicitly refer to such an index. For example, Lacy and Niou [11] and Benoit and Kornhauser [2] showed that for multiple referenda, if voters vote on issues separately (under some assumptions on how voters vote), then there exists a profile such that in each vote, the winner is ranked near the bottom—therefore the rule has a very low MSI.

5 Multiple-Election Paradoxes for Strategic Sequential Voting

In this section, we show that over multi-issue domains, for any n that is sufficiently large (we will specify the number in our theorems), there exists an n -profile P such that $SSP_{\mathcal{O}}(P)$ is ranked almost in the bottom position in each vote in P . That is, the minimax satisfaction index is extremely low for the strategic sequential voting procedure SSP.

We first calculate the MSI for $SSP_{\mathcal{O}}$ when the winner does not depend on the tie-breaking mechanism. That is, either n is odd, or n is even and there is never a tie in the process of running the election. (This is our main multiple-election paradox result.) **Because of the space constraint, we omit all proofs. The full version of this paper can be found online.**

Theorem 1 *For any $p \in \mathbb{N}$ ($p \geq 2$) and any $n \geq 2p^2 + 1$, $MSI_{SSP_{\mathcal{O}}}(m, n) = \lfloor p/2 + 2 \rfloor$.¹ Moreover, in the profile P that we use to prove the upper bound, the winner $SSP_{\mathcal{O}}(P)$ is Pareto-dominated by $2^p - (p + 1)p/2$ alternatives.*

We note that the number of alternatives is $m = 2^p$. Therefore, $\lfloor p/2 + 2 \rfloor$ is exponentially smaller than the number of alternatives, which means that there exists a profile for which every voter ranks the winner very close to the bottom. Moreover, $(p + 1)p/2$ is still exponentially smaller than 2^p , which means that the winner is Pareto-dominated by almost every other alternative.

Naturally, we wish to avoid such paradoxes. One may wonder if the paradox occurs only if the ordering of the issues is particularly unfortunate with respect to the preferences of the voters. If not, then, for example, perhaps a good approach is to randomly choose the order of the issues.² Unfortunately, our next result shows that we can construct a single profile that results in a paradox for *all* orderings. While it works for all orders, the result is otherwise somewhat weaker than Theorem 1: it does not show a Pareto-dominance result, it requires a number of voters that is at least twice the number of alternatives, the upper bound shown on the MSI is very slightly higher than in Theorem 1, and unlike Theorem 1 no matching lower bound is shown.

Theorem 2 *For any $p, n \in \mathbb{N}$ (with $p \geq 2$ and $n \geq 2^{p+1}$), there exists an n -profile P such that for any order \mathcal{O} over $\{\mathbf{x}_1, \dots, \mathbf{x}_p\}$, $SSP_{\mathcal{O}}(P) = 1_1 \dots 1_p$, and any $V \in P$ ranks $1_1 \dots 1_p$ somewhere in the bottom $p + 2$ positions.*

6 Multiple-election paradoxes for SSP with restrictions on preferences

The paradoxes exhibited so far placed no restriction on the voters' preferences. While SSP is perfectly well defined for any preferences that the voters may have over the alternatives, we may yet wonder what happens if the voters' preferences over alternatives are restricted in a way that is natural with respect to the multi-issue structure of the setting. In particular, are paradoxes avoided by such restrictions? It is well known that natural restrictions on preferences sometimes lead to much more positive results in social choice and mechanism design—for example, single-peaked preferences allow for good strategy-proof mechanisms [3, 15].

In this section, we study the MSI for $SSP_{\mathcal{O}}$ for the following three cases: (1) voters' preferences are separable; (2) voters' preferences are \mathcal{O} -lexicographic; and (3) voters' preferences are \mathcal{O} -legal. For case (1), we show a mild paradox (and that this is effectively the strongest paradox that can be obtained); for case (2), we show a positive result; for case (3), we show a paradox that is very nearly as bad as the unrestricted case.

Theorem 3 *For any $n \geq 2p$, when the profile is separable, the MSI for $SSP_{\mathcal{O}}$ is between $2^{\lceil p/2 \rceil}$ and $2^{\lfloor p/2 \rfloor + 1}$.*

¹If n is even, then to prove $MSI_{SSP_{\mathcal{O}}}(m, n) \geq \lfloor p/2 + 2 \rfloor$, we restrict attention to profiles without ties.

²Of course, for any ordering of the issues, there exists a profile that results in the paradoxes in Theorem 1; but this does not directly imply that there exists a single profile that works for all orderings.

That is, the MSI of $SSP_{\mathcal{O}}$ when votes are separable is $\Theta(\sqrt{m})$. We still have that $\lim_{m \rightarrow \infty} \Theta(\sqrt{m})/m = 0$, so in that sense this is still a paradox. However, the convergence to 0 is much slower than for $\Theta(\log m)/m$, which corresponds to the convergence rate for the earlier paradoxes.

Theorem 4 *For any $p \in \mathbb{N}$ ($p \geq 2$) and any $n \geq 5$, when the profile is \mathcal{O} -lexicographic, $MSI(SSP_{\mathcal{O}}) = 3 \cdot 2^{p-2} + 1$. Moreover, $SSP_{\mathcal{O}}(P)$ is ranked somewhere in the top 2^{p-1} positions in at least $n/2$ votes.*

Naturally $\lim_{m \rightarrow \infty} (3m/4 + 1)/m = 3/4$, so in that sense there is no paradox when votes are \mathcal{O} -lexicographic.

Finally, we study the MSI for $SSP_{\mathcal{O}}$ when the profile is \mathcal{O} -legal. Theorem 5 shows that it is very nearly as bad as the unrestricted case (Theorem 1).

Theorem 5 *For any $p, n \in \mathbb{N}$ with $n \geq 2p^2 + 2p + 1$, there exists an \mathcal{O} -legal profile such that in each vote, no more than $\lceil p/2 \rceil + 4$ alternatives are ranked lower than $SSP_{\mathcal{O}}(P)$. Moreover, $SSP_{\mathcal{O}}(P)$ is Pareto-dominated by at least $2^p - 4p^2$ alternatives.*

Of course, the lower bound on the MSI from Theorem 1 still applies when the profile is \mathcal{O} -legal, so together with Theorem 5 this proves that the MSI for $SSP_{\mathcal{O}}$ when the profile is \mathcal{O} -legal is $\Theta(\log m)$, just as in the unrestricted case.

7 Minimax satisfaction index of other common voting rules

So far, we have focused strictly on strategic sequential voting (SSP) in multi-issue domains (and voting trees, but only in the sense of their equivalence to strategic sequential voting). Hence, at this point, it may not be clear whether the paradoxes (or, in some cases, lack of paradoxes) that we have shown are due to the sequential, multi-issue nature of the process, or whether they are due to the strategic behavior, or whether such paradoxes are prevalent throughout voting settings.

First, let us address the question of to what extent they are due to strategic behavior. To answer this, it is most natural to compare to $Seq_{\mathcal{O}}(maj, \dots, maj)$ (“truthful” sequential voting), which is only well defined when the profile is \mathcal{O} -legal. We answer this question by the following Proposition, which shows a much milder paradox.

Proposition 1 *For any $n \geq 2p$, when the profile is \mathcal{O} -legal, the MSI for $Seq_{\mathcal{O}}(maj, \dots, maj)$ is between $2^{\lceil p/2 \rceil}$ and $2^{\lfloor p/2 \rfloor + 1}$.*

Having settled the effect of strategic behavior, we next investigate the effect of the multi-issue nature of the setting. We do this by studying the MSI of common voting rules in non-combinatorial settings, where there is a single issue (but one that can take more than two values). In this context, studying strategic behavior seems intractable. By the Gibbard-Satterthwaite theorem [9, 18], without restrictions on preferences, no strategy-proof rules exist other than dictatorships and rules that exclude certain alternatives *ex ante*. Moreover, even with complete information, common voting rules have many different equilibria. Hence, we focus on studying the extent to which paradoxes occur when voters vote truthfully.

Specifically, we investigate the minimax satisfaction indices of positional scoring rules (including k -approval and Borda), plurality with runoff (*Pluo*), Copeland $_{\alpha}$, maximin, ranked pairs, Bucklin, STV, and (not necessarily balanced) voting trees. Of

course, these rules can be applied to multi-issue domains as well as to any other domain, but they do not make use of multi-issue structure; in general, we just have a set of alternatives $C = \{c_1, \dots, c_m\}$. Throughout the remainder of this section, we assume that $m \geq 3$, and that ties are broken in the order $c_1 \succ c_2 \succ \dots \succ c_m$.

Proposition 2 *Let $m, n \in \mathbb{N}$.*

- $MSI_{Dict}(m, n) = m$;
- for any $k \leq m$, $MSI_{Appk}(m, n) = m + 1 - k$;
- $MSI_{Pluo}(m, n) = m$;
- $MSI_{STV}(m, n) = m$;
- $MSI_{Bucklin}(m, n) \geq m/2$.

Proposition 3 (Borda) *Let $m \in \mathbb{N}$. For any $n \in \mathbb{N}$ such that n is even, $MSI_{Borda}(m, n) = \lfloor m/2 + 1 \rfloor$; for any $n \in \mathbb{N}$ such that $n \geq m$, and n is odd, $MSI_{Borda}(m, n) = \lceil m/2 + 1 \rceil$.*

Proposition 4 (Copeland) *Let $m, n \in \mathbb{N}$. If either $0 < \alpha \leq 1$, or n is odd and $\alpha = 0$, then $MSI_{Copeland_\alpha}(m, n) \geq \alpha m/4$. For any $n \geq 2m$ such that n is even, $MSI_{Copeland_0}(m, n) = 2$.*

Proposition 5 (Maximin) *Let $m, n \in \mathbb{N}$ with $n \geq m - 1$. $MSI_{maximin}(m, n) \leq 3$.*

Proposition 6 (Ranked pairs) *Let $m, n \in \mathbb{N}$ with $n \geq \sqrt{m}$. $MSI_{rp}(m, n) \geq \sqrt{m}$.*

Proposition 7 (Voting trees) *Let T be a voting tree; let c be the alternative whose corresponding leaf is closest to the root among all leaves in T , and let its distance to the root be denoted l . If $l = 1$, then for any $n \geq 2m$, $MSI_{r_T}(m, n) = 3$; if $l \geq 2$, then for any $n \geq 2m$, $MSI_{r_T}(m, n) = \lfloor l/2 + 2 \rfloor$.*

Proposition 7 implies that among all voting trees for m alternatives, balanced voting trees have the highest MSI.

8 Conclusion and future work

In this paper, we considered a complete-information game-theoretic analysis of sequential voting on binary issues, which we called strategic sequential voting. Specifically, given that agents have complete information about each other's preferences and their preferences are strict, the game can be solved by a natural backward induction process, which leads to a unique solution.

We showed that strategic sequential voting is prone to multiple-election paradoxes; to do so, we introduced a minimax satisfaction index, which measures the degree to which at least one voter is made happy by the outcome of the election. We showed that the minimax satisfaction index for strategic sequential voting is exponentially small, which means that there exists a profile for which the winner is ranked almost in the bottom position in all votes; even worse, the winner is Pareto-dominated by almost every other alternative. We showed that changing the order of the issues in sequential voting cannot completely avoid the paradoxes. These negative results indicate that the solution of the sequential game can be very bad. We also showed that multiple-election paradoxes can be avoided to some extent by restricting voters' preferences to be separable or lexicographic, but the paradoxes still exist when the voters' preferences are \mathcal{O} -legal.

Voting rule	MSI
Dictatorships	m (Proposition 2)
Plu w/ runoff	m (Proposition 2)
STV	m (Proposition 2)
Copeland $_{\alpha}$ ($0 < \alpha \leq 1$)	$\Theta(m)$ (Proposition 4)
Borda ($n \geq m$)	$\Theta(m)$ (Proposition 3)
Bucklin	$\Theta(m)$ (Proposition 2)
$Seq_{\mathcal{O}}(maj, \dots, maj)$ (\mathcal{O} -lexico profiles)	$3m/4 + 1$ (Theorem 4)
$SSP_{\mathcal{O}}$ (\mathcal{O} -lexico profiles)	$3m/4 + 1$ (Theorem 4)
k -Approval (incl. Plurality and Veto)	$m + 1 - k$ (Proposition 2)
Ranked pairs ($n \geq \sqrt{m}$)	$\Omega(\sqrt{m})$ (Proposition 6)
$Seq_{\mathcal{O}}(maj, \dots, maj)$ (separable profiles)	$\Theta(\sqrt{m})$ (Theorem 3)
$SSP_{\mathcal{O}}$ (separable profiles)	$\Theta(\sqrt{m})$ (Theorem 3)
$Seq_{\mathcal{O}}(maj, \dots, maj)$ (\mathcal{O} -legal profiles)	$\Theta(\sqrt{m})$ (Proposition 1)
$SSP_{\mathcal{O}}$ (\mathcal{O} -legal profiles)	between $\lfloor \log m/2 + 2 \rfloor$ and $\lfloor \log m/2 + 5 \rfloor$ (Theorem 5)
$SSP_{\mathcal{O}}$ ³	$\lfloor \log m/2 + 2 \rfloor$ (Theorem 1)
Voting tree ($n \geq 2m$)	$\lfloor l/2 + 2 \rfloor$ ⁴ (Proposition 7)
Maximin ($n \geq m - 1$)	≤ 3 (Proposition 5)
Copeland $_0$ (n is even)	2 (Proposition 4)

Table 1: The minimax satisfaction index for strategic sequential voting (SSP), truthful sequential voting (Seq), and common voting rules, ranked roughly from high to low.

For the sake of benchmarking our results, we also study the minimax satisfaction index for some common voting rules (under truthful voting). The results are summarized in Table 1. For a voting rule with a low (high) MSI, we can (cannot) find a paradox that is similar to the multiple-election paradoxes—that is, a profile for which the winner is ranked in extremely low positions in all votes.

From this table, it may be concluded that: (1) in sequential voting, the paradoxes are stronger when voting is strategic than when it is truthful, though of course this is no longer true if we are in a restricted setting where truthful and strategic voting lead to identical results (that is, when the profile is separable or lexicographic); (2) the strength of the paradoxes for sequential voting ranks somewhere in the middle, though perhaps somewhat more on the strong side, among standard social choice rules (when voters are assumed to vote truthfully).

There are many topics for future research. For example, given a profile, can we characterize the set of alternatives that win for some order over the issues?⁵ Perhaps more importantly, how can we get around the multiple-election paradoxes in sequential games? For example, Theorem 4 shows that if the voters’ preferences are lexicographic, then we can avoid the paradoxes. It is not clear if there are other ways to avoid the paradoxes (paradoxes occur even if we restrict voters’ preferences to be separable or \mathcal{O} -legal, as shown in Theorem 3 and Theorem 5). Another approach is to consider other, non-sequential voting procedures for multi-issue domains. What are good examples of such procedures? Will these avoid paradoxes? What is the effect of strategic behavior for such procedures? How should we even define “strategic behavior” for such procedures, or for sequential voting with non-binary issues, or for voting rules in general? How can we extend these results to incomplete-information settings?⁶ Also, beyond

³Additionally, there exists a profile P such that for any order \mathcal{O} over issues, the maximum satisfaction index of $SSP_{\mathcal{O}}$ for P is no more than $\log m + 2$ (Theorem 2).

⁴ l is the minimum distance from the root to a leaf, $l \leq \log m$. If $l = 1$, then $MSI_{r_T}(m, n) = 3$.

⁵This results in a *social choice set* or *correspondence*; social choice sets have recently attracted attention from computer scientists [5].

⁶Of course, because the complete-information setting is a special case of incomplete-information settings,

proving paradoxes for individual rules, is it possible to show a general impossibility result that shows that under certain minimal conditions, paradoxes cannot be avoided?⁷

References

- [1] Salvador Barberà, Hugo Sonnenschein, and Lin Zhou. Voting by committees. *Econometrica*, 59(3):595–609, 1991.
- [2] Jean-Pierre Benoit and Lewis Kornhauser. Only a dictatorship is efficient or neutral. Technical report, 2006.
- [3] Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
- [4] Steven Brams, David Kilgour, and William Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.
- [5] Felix Brandt, Felix Fischer, and Paul Harrenstein. The computational complexity of choice sets. *Mathematical Logic Quarterly*, 55:444–459, 2009.
- [6] Bhaskar Dutta and Arunava Sen. Implementing generalized Condorcet social choice functions via backward induction. *Social Choice and Welfare*, 10(2):149–160, 1993.
- [7] R. Farquharson. *Theory of Voting*. Yale University Press, 1969.
- [8] Felix Fischer, Ariel D. Procaccia, and Alex Samorodnitsky. A new perspective on implementation by voting trees. In *Proc. of EC-09*, pages 31–40, 2009.
- [9] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–602, 1973.
- [10] R. J. Gretlein. Dominance elimination procedures on finite alternative games. *International Journal of Game Theory*, 12(2):107–113, 1983.
- [11] Dean Lacy and Emerson M.S. Niou. A problem with referendums. *J. of Theoretical Politics*, 12(1):5–31, 2000.
- [12] Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, pages 304–324, 2009.
- [13] Richard D. McKelvey and Richard G. Niemi. A multistage game representation of sophisticated voting for binary procedures. *Journal of Economic Theory*, 18(1):1 – 22, 1978.
- [14] Herve Moulin. Dominance solvable voting schemes. *Econometrica*, 47(6):1137–51, 1979.
- [15] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [16] Hervé Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.
- [17] Donald G. Saari and Katri K. Sieberg. The sum of the parts can violate the whole. *The American Political Science Review*, 95(2):415–433, 2001.
- [18] Mark Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [19] Marco Scarsini. A strong paradox of multiple elections. *Social Choice and Welfare*, 15(2):237–238, 1998.

things can only get worse in the latter.

⁷This may require quite restrictive conditions or a different notion of a paradox—for example, we have already shown that several natural voting rules have MSI m , albeit under truthful voting.