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

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

Vincent Conitzer
Dept. 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 voters must come to a joint decision on multiple issues. In practice, this is often done by voting on the issues in sequence. We model sequential voting in multi-issue domains as a complete-information extensive-form game, in which the voters 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 given exogenously before the process. We call this model *strategic sequential voting*.

We focus on domains characterized by multiple binary issues, so that strategic sequential voting leads to a unique outcome under a natural solution concept. We show that under some conditions on the preferences, this leads to the same outcome as truthful sequential voting, but in general it can result in very different outcomes. In particular, sometimes the order of the issues has a strong influence on the winner. We also analyze the communication complexity of the corresponding social choice rule.

Most significantly, we illustrate 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 voters' true preferences, and the winner is Pareto-dominated by almost every other alternative. We show that changing the order of the issues cannot completely prevent such paradoxes. We also study the possibility of avoiding the paradoxes for strategic sequential voting by imposing some constraints on the profile, such as separability, lexicographicity or $\mathcal{O}$-legality.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Theory

Keywords

Social choice, strategic voting, multi-issue domains, multiple-election 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 $\mathcal{I} = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$, and each issue can take a value 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 [5]. 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 voters to represent their preferences, as well as designing a sensible voting rule to make decisions based on the reported preferences using such a language. 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 voters 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 [15]. 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) [14]. Second, *multiple-election paradoxes* arise in issue-by-issue voting [5, 14, 21, 23]. 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.

One 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 $\mathbf{x}_1 > \cdots > \mathbf{x}_n$), 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. For instance, 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. Another example is that at the executive meeting of the co-owners of a building, important decisions like 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. Let us consider the following motivating example.

Example 1 Three residents want to vote to decide whether they should build a swimming pool and/or a tennis court. There are two issue $\mathbf{S}$ and $\mathbf{T}$. $\mathbf{S}$ can take the value of s (meaning "to build the swimming pool") or $\bar{s}$ (meaning "not to build the swimming pool"). Similarly, $\mathbf{T}$ takes a value in $\{t, \bar{t}\}$. Suppose the preferences of the three voters are, respectively, $st \succ \bar{s}t \succ s\bar{t} \succ \bar{s}\bar{t}$, $s\bar{t} \succ st \succ \bar{s}t \succ \bar{s}\bar{t}$ and $\bar{s}t \succ \bar{s}\bar{t} \succ s\bar{t} \succ st$. Voter 2 and 3 took the budget constraint into consideration so that they do not rank st as their first choices. Suppose the voters first vote on issue $\mathbf{S}$ then on $\mathbf{T}$. Moreover, since both issues are binary, the local rule used at each step is majority (there will be no ties, because the number of voters is odd). Voter 1 is likely to reason in the following way: *if the outcome of the first step is s, then voters 2 and 3 will vote for $\bar{t}$, since they both prefer $s\bar{t}$ to st, and the final outcome will be $s\bar{t}$; but if the outcome of the first step is $\bar{s}$, then voters 2 and 3 will vote for t, and the final outcome will be $\bar{s}t$; because I prefer $\bar{s}t$ to $s\bar{t}$, I am better off voting for $\bar{s}$, since either it will not make any difference, or it will lead to a final outcome of $\bar{s}t$ instead of $s\bar{t}$.* If voters 2 and 3 reason in the same way, then 2 will vote for s and 3 for $\bar{s}$; hence, the result of the first step is $\bar{s}$, and then, since two voters out of three prefer $\bar{s}t$ to $\bar{s}\bar{t}$, the final outcome will be $\bar{s}t$. Note that the result is fully determined, provided that (1) it is common knowledge that voters behave strategically according to the principle we have stated informally, (2) the order in which the issues are decided, as well as the local voting rules used in all steps, are also common knowledge, and (3) voters' preferences are common knowledge. Therefore, these three assumptions allow the voters and the modeler (provided he knows as much as the voters) to predict the final outcome.

Let us take a closer look at voter 1 in Example 1. Her preferences are *separable*: she prefers s to $\bar{s}$ whatever the value of $\mathbf{T}$ is, and t to $\bar{t}$ whatever the value of $\mathbf{S}$ is. *And yet she strategically votes for* $\bar{s}$, because the outcome for $\mathbf{S}$ affects the outcome for $\mathbf{T}$. Moreover, while voters 2 and 3 have nonseparable preferences, still, all three voters' preferences enjoy the following property: their preferences over the value of $\mathbf{S}$ are independent of the value of $\mathbf{T}$. Such a profile is called a *legal* profile with respect to the order $\mathbf{S} > \mathbf{T}$, meaning that the voters vote on $\mathbf{S}$ first, then on $\mathbf{T}$. Lang and Xia [15] defined a family of sequential voting rules on multi-issue domains, restricted to $\mathcal{O}$-legal profiles for some order $\mathcal{O}$ over the issues, where at each step, each voter is expected to vote for her preferred value for the issue $\mathbf{x}_i$ under consideration given the values of all issues decided so far [1]; then, the value of $\mathbf{x}_i$ is chosen according to a local voting rule, and this local outcome is broadcast to the voters. For example, suppose the local rule used to decide an issue is always majority. For the profile given in Example 1, the outcome of the first step under the sequential voting rule will be s (since two voters out of three prefer s to $\bar{s}$, unconditionally), and the final outcome will be $s\bar{t}$. This outcome is different from the outcome we obtain if voters behave strategically. The reason for this discrepancy is that in [15], voters are not assumed to know the others' preferences and are assumed to vote truthfully.

We have seen that even if the voters' preferences are $\mathcal{O}$-legal, voters may in fact have no incentive to vote truthfully. Consequently, existing results on multiple-election paradoxes are not directly applicable to situations where voters vote strategically.

Our contributions

In this paper, we analyze the complete-information game-theoretic model of sequential voting that we illustrated in Example 1. This model applies to any preferences that the voters may have (not just $\mathcal{O}$-legal ones), though they must be strict orders on the set of all alternatives.

We focus on voting in multi-binary-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 (without loss of generality, let $\mathcal{O} = \mathbf{x}_1 > \mathbf{x}_2 > \cdots > \mathbf{x}_p$), which indicates the order in which these issues will be voted on. For any $1 \leq i \leq p$, 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: it is common knowledge that all voters are perfectly rational; the order $\mathcal{O}$ and the fact that in each step, the majority rule is used to determine the winner are common knowledge; all voters' preferences are common knowledge.

We can solve this game by a type of backward induction already illustrated in Example 1: in the last (pth) 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 pth 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 exogenously the order $\mathcal{O}$ over the 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 $SSP_{\mathcal{O}}$.

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 which we call 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 input rankings of the alternatives (the ranking where this position is obtained corresponds 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,

[1] The $\mathcal{O}$-legality condition ensures that this notion of "preferred value of $\mathbf{x}_i$" is meaningful.

thus indicating a multiple-election paradox. Our main theorem is the following.

Theorem 2 *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 2 implies that the winner for SSP is an almost 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 an almost Condorcet loser, and the winner is Pareto-dominated by almost every other alternative. We further show a paradox (Theorem 3) that states that there exists a profile such that for *any* order $\mathcal{O}$ over the issues, for every voter, the SSP winner w.r.t. $\mathcal{O}$ 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.

Related work and discussion

Our setting is closely related to the *multi-stage sophisticated voting*, studied by McKelvey and Niemi [17], Moulin [18], and Gretlein [13]. They investigated 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 focused on the characterization of the outcomes as the outcomes in the *sophisticated voting* [10], and therefore did not shed much light on the quality of the equilibrium outcome. We, on the other hand, are primarily interested in the strategic outcome of the natural procedure of voting sequentially over multiple issues. 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. More importantly, we illustrate several multiple-election paradoxes for SSP, indicating that the equilibrium outcome could be extremely undesirable.

Another paper that is closely related to part of this work was written by Dutta and Sen [9]. They showed 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 the relationship revealed for multi-stage sophisticated voting and will also be mentioned later in this paper, that is, 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 voter moves at a time, and a move consists not of a vote, but rather of choosing the next player to move (or in some states, choosing the winner).

Nevertheless, the approach by Dutta and Sen and our approach 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.

Less closely related, implementation by voting trees has previously been studied at EC: Fischer et al. [11] consider the known result that the Copeland rule cannot be implemented by a voting tree [20], and set out to *approximate* the the Copeland score using voting trees.

It has been pointed out that typical multiple-election paradoxes partly come from the incompleteness of information about the preferences of the voters [14]. 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.

Lastly, in a recent paper [25], Xia and Conitzer studied a voting game with a different type of sequential nature: in it, the voters cast their votes one after another (strategically), and after all the voters have cast their votes, a common voting rule (not necessarily the plurality rule) is used to select the winner. This type of voting games has been studied in the literature [24, 7, 2, 8]. In [25], a strong general paradox was shown for these voting games, implying that for most common voting rules, there exists a profile such that the unique winner in subgame-perfect equilibrium is ranked within the bottom two positions in almost all the voters' true preferences. Desmedt and Elkind [8] showed similar paradoxes for such voting games with the plurality rule, where random tie-breaking is used and the voters seek to maximize their expected utility. We note that the voting games studied in [25] are quite different from the voting games studied in this paper: there, the voters move in sequence, the set of alternatives does not need to have a combinatorial structure, and a voter casts her complete vote all at once; here in this paper, the voters move simultaneously, the set of alternatives has a combinatorial structure, and the voters vote on one issue at a time.

2. PRELIMINARIES

2.1 Basics of voting

Let $\mathcal{X}$ be the set of *alternatives*, $|\mathcal{X}| = m$. A vote is a linear order (that is, a transitive, antisymmetric, and total relation) 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 $rank_V(c)$ denote the position of c in V from the top. For any $n \in \mathbb{N}$, an *n-profile* P is a collection of n votes, 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 every $V \in P$, 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, $r : L(\mathcal{X}) \cup L(\mathcal{X})^2 \cup \ldots \to \mathcal{X}$. For example, when there are two alternatives, the *majority* rule selects the alternative that is preferred by the majority of voters.

2.2 Multi-issue domains

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

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 $\mathcal{O} = \mathbf{x}_1 > \cdots > \mathbf{x}_p$. A preference relation $\succ$ is *$\mathcal{O}$-legal* if for any $i \leq p$, $\mathbf{x}_i$ is preferentially independent of $\{\mathbf{x}_{i+1}, \ldots, \mathbf{x}_p\}$ given $\{\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}\}$. In words, 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 $\mathcal{X} \setminus \{\mathbf{x}_i\}$. That is, we do not need to know the value of any other issue to find out whether a particular change in the value of an issue will make the voter better or worse off. It follows directly that 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 $\vec{u} \in D_1 \times \cdots \times D_{i-1}$, any $a_i, b_i \in D_i$, and any $\vec{d}_1, \vec{d}_2, \vec{e}_1, \vec{e}_2 \in D_{i+1} \times \cdots \times D_p$, $(\vec{u}, a_i, \vec{d}_1) \succ (\vec{u}, b_i, \vec{e}_1)$ if and only if she prefers $(\vec{u}, a_i, \vec{d}_2) \succ (\vec{u}, b_i, \vec{e}_2)$. In words, 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 $\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}$ will be the same, they still matter in that they affect the voter's preferences on $\mathbf{x}_i$.) We note that $\mathcal{O}$-lexicographicity and separability are incomparable notions. For example, $0_1 0_2 \succ 1_1 0_2 \succ 0_1 1_2 \succ 1_1 1_2$ is separable (flipping 1_1 or 1_2 always makes the alternative rank higher) but not ($\mathbf{x}_1 > \mathbf{x}_2$)-lexicographic ($0_1 0_2 \succ 1_1 1_2$ but $1_1 0_2 \succ 0_1 1_2$). On the other hand, $0_1 0_2 \succ 0_1 1_2 \succ 1_1 1_2 \succ 1_1 0_2$ is ($\mathbf{x}_1 > \mathbf{x}_2$)-lexicographic but not separable. 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* $(r_1, \ldots, r_p)$ (where for any $i \leq p$, r_i is a voting rule for preferences over D_i), the *sequential composition* of $r_1, \ldots, r_p$ with respect to $\mathcal{O}$, denoted by $Seq_{\mathcal{O}}(r_1, \ldots, r_p)$, is defined for all $\mathcal{O}$-legal profiles as follows: $Seq_{\mathcal{O}}(r_1, \ldots, r_p)(P) = (d_1, \ldots, d_p) \in \mathcal{X}$, where for any $i \leq p$, $d_i = r_i(P|_{\mathbf{x}_i:d_1 \cdots d_{i-1}})$, where $P|_{\mathbf{x}_i:d_1 \cdots d_{i-1}}$ is composed of the voters' local preferences over $\mathbf{x}_i$, given that the issues preceding it take values $d_1, \ldots, d_{i-1}$. 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 $d_1, \ldots, d_{i-1}$ that have already been determined for the issues that precede $\mathbf{x}_i$. In this paper, we focus on the case where every r_i is the majority rule, because it is the most natural voting rule for two alternatives.

3. STRATEGIC SEQUENTIAL VOTING

3.1 Formal definition

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 (in an optimal manner that will be explained later), 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 rules), 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), in that sense, *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. Of course, for incomplete information setting in general, we need a more elaborate model to reason about voters' strategic behavior.

Given these assumptions, the voting process can be modeled as a 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 $\mathbf{x}_1 > \cdots > \mathbf{x}_p$. Let P be 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.

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_i, and what it will be if it is 1_i. More concretely, suppose that steps 1 to $i-1$ resulted in issues $\mathbf{x}_1, \ldots, \mathbf{x}_{i-1}$ taking the values $d_1, \ldots, d_{i-1}$, and let $\vec{d} = (d_1, \ldots, d_{i-1})$. Suppose also that if $\mathbf{x}_i$ takes the value 0_i (respectively, 1_i), then, recursively, the remaining issues will take the tuple of values $\vec{a}$ (respectively, $\vec{b}$). 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 the majority of voters prefer $(\vec{d}, 0_i, \vec{a})$ 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 (SSP)* procedure, and for any profile P, the winner with respect to 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. The following definition and two propositions merely serve to make the game-theoretic solution concept that we use precise; a reader who is not interested in this may safely skip them.

Definition 1 *Consider a finite extensive-form game which transitions among* states. *In each nonterminal state s, all players simultaneously take an action; this joint local action profile $(a_1^s, \ldots, a_n^s)$ determines the next state s'.*[2] *Terminal states t are associated with payoffs for the players (alternatively, players have ordinal preferences over the terminal states). The current state is always common knowledge among the players.*[3]

Suppose that in every final nonterminal state s (that is, every state that has only terminal states as successors), every player i has a (weakly) dominant action a_i^s. At each final nonterminal state, its local profile of dominant actions $(a_1^s, \ldots, a_n^s)$ results in a terminal state $t(s)$ and associated payoffs. We then replace each final nonterminal state s with the terminal state $t(s)$ that its dominant-strategy profile leads to. Furthermore suppose that in the resulting smaller tree, again, in every final nonterminal state, every player

[2]In the extensive-form representation of the game, each state is associated with multiple nodes, because in the extensive form only one player can move at a node.

[3]Hence, the only imperfect information in the extensive form of the game is due to simultaneous moves within states.

has a (weakly) dominant strategy. Then, we can repeat this procedure, etc. *If we can repeat this all the way to the root of the tree, then we say that the game is solvable by* within-state dominant-strategy backward induction (WSDSBI).

We note that the backward induction in perfect-information extensive-form games is just the special case of WSDSBI where in each state only one player acts.

Proposition 1 *If a game is solvable by WSDSBI, then the solution is unique.*

Proposition 2 *The complete-information sequential voting game with binary issues (with majority as the local rule everywhere) is solvable by WSDSBI when voters have strict preferences over the alternatives.*

Proposition 1 is obviously true. Due to the space constraints, most proofs are omitted.

We note that 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 pth 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 pth issue; etc. This bottom-up procedure corresponds exactly to the backward induction (WSDSBI) process.

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 (corresponding to the different orders of the issues), but many more voting trees. The voting tree corresponding to the order $\mathcal{O} = \mathbf{x}_1 > \cdots > \mathbf{x}_p$ 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 (respectively, right) subtree of v gives the value 0_i (respectively, 1_i) to $\mathbf{x}_i$.

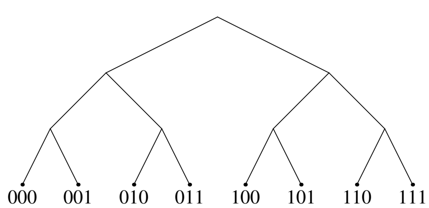

Figure 1: A voting tree that is equivalent to the strategic sequential voting procedure ($p = 3$). 000 is the abbreviation for $0_1 0_2 0_3$, etc.

3.2 Strategic sequential voting vs. truthful sequential voting

We have seen on Example 1 that even when the profile P is $\mathcal{O}$-legal, $SSP_{\mathcal{O}}(P)$ can be different from $Seq_{\mathcal{O}}(maj, \ldots, maj)(P)$. This means that even if the profile is $\mathcal{O}$-legal, voters may be better off voting strategically than truthfully. However, $SSP_{\mathcal{O}}(P)$ and $Seq_{\mathcal{O}}(maj, \ldots, maj)(P)$ are guaranteed to coincide under the further restriction that P is $\mathcal{O}$-lexicographic.

Proposition 3 *For any $\mathcal{O}$-lexicographic profile P, $SSP_{\mathcal{O}}(P) = Seq_{\mathcal{O}}(maj, \ldots, maj)(P)$.*

The intuition for Proposition 3 is as follows: if P is $\mathcal{O}$-lexicographic, then, as is shown in the proof of the proposition, when voters vote strategically under sequential voting (the Seq process), they are best off voting according to their true preferences in each round (their preferences in each round are well-defined because voters have $\mathcal{O}$-legal preferences in this case). When voters with $\mathcal{O}$-legal preferences vote truthfully in each round under sequential voting, the outcome is $Seq_{\mathcal{O}}(maj, \ldots, maj)(P)$; when they vote strategically, the outcome is $SSP_{\mathcal{O}}(P)$; and so, these must be the same when preferences are $\mathcal{O}$-lexicographic.

Now, there is another interesting domain restriction under which $SSP_{\mathcal{O}}(P)$ and $Seq(maj, \ldots, maj)(P)$ coincide, namely when P is $inv(\mathcal{O})$- legal, where $inv(\mathcal{O}) = (\mathbf{x}_p > \ldots > \mathbf{x}_1)$.

Proposition 4 *Let $inv(\mathcal{O}) = \mathbf{x}_p > \ldots > \mathbf{x}_1$. For any $inv(\mathcal{O})$-legal profile P, $SSP_{\mathcal{O}}(P) = Seq_{inv(\mathcal{O})}(maj, \ldots, maj)(P)$.*

As a consequence, when P is separable, it is *a fortiori* $inv(\mathcal{O})$-legal, and therefore, $SSP_{\mathcal{O}}(P) = Seq_{inv(\mathcal{O})}(maj, \ldots, maj)(P)$, which in turn is equal to $Seq_{\mathcal{O}}(maj, \ldots, maj)(P)$ and coincides with seat-by-seat voting [3].

Corollary 1
If P is separable, then $SSP_{\mathcal{O}}(P) = Seq_{\mathcal{O}}(maj, \ldots, maj)(P)$.

3.3 A second interpretation of SSP

The first interpretation of SSP (that we follow in this paper) is the one we have discussed so far, namely, SSP consists in modeling sequential voting as a complete-information game, which allows us to analyze sequential voting on multi-issue domains from a game-theoretic point of view. For this, assumptions 1, 2, and 3 above are crucial. Under this interpretation, $SSP_{\mathcal{O}}(P)$ is a (specific kind of) equilibrium for sequential voting.

However, there is a second interpretation of SSP. It consists in seeing $SSP_{\mathcal{O}}$ as a new voting rule on multi-issue domains (which is implementable in complete-information contexts by using sequential voting).[4] This defines a family of voting rules (one for each order over issues), which can be applied to any profile. The family of voting rules thus defined is a distinguished subset of the family of voting trees. This interpretation does not say anything about how preferences are to be elicited; unlike in the game-theoretic interpretation, the p-step protocol does not apply here. The communication complexity of finding the outcome of $SSP_{\mathcal{O}}$ (without any complete-information assumption, of course)[5] is given as follows.

Proposition 5 *When the voters' preferences over alternatives are unrestricted, the communication complexity of $SSP_{\mathcal{O}}$ is $\Theta(2^p \cdot n)$.*

Proof of Proposition 5: This now follows immediately from a result in [6], where it is established that the communication complexity for balanced voting trees is $\Theta(m \cdot n)$ for m alternatives and n voters. Since we do not place any restrictions on the preferences in the multi-issue domain in the statement of the proposition, the communication complexity is identical, and $m = 2^p$. □

The upper bound in this proposition is obtained simply by eliciting the voters preferences for every pair of alternatives that face each other in the voting tree.

Now, Propositions 3 and 4 immediately give us conditions under which this communication complexity can be reduced. Indeed, these Propositions say that when P is $\mathcal{O}$-lexicographic or $inv(\mathcal{O})$-legal, then the SSP winner coincides with the sequential election winner in the sense of [15]. Now, the sequential election winner in the sense of [15]. can be found with $O(pn)$ communication, simply by having each agent vote for a value for the issue at each round. This leads immediately to the following two corollaries (to Propositions 3 and 4, respectively).

[4]Of course, by Gibbard-Satterthwaite [12, 22], SSP is not strategy-proof.

[5]The communication complexity of a voting rule is the smallest number of bits that must be transmitted to compute the winner of that rule (i.e., taking the minimum across all correct protocols). See [6].

Corollary 2 *When the voters' preferences over alternatives are $\mathcal{O}$-lexicographic, the communication complexity of $SSP_{\mathcal{O}}$ is $O(pn)$.*

Corollary 3 *When the voters' preferences over alternatives are $inv(\mathcal{O})$-legal, the communication complexity of $SSP_{\mathcal{O}}$ is $O(pn)$.*

3.4 The winner is sensitive to the order over the issues

In the definition of SSP, we simply fixed the order $\mathcal{O}$ to be $\mathbf{x}_1 > \mathbf{x}_2 > \cdots > \mathbf{x}_p$. A question worth addressing is, to what extent is the outcome of SSP sensitive to the variation of the order $\mathcal{O}$? More precisely, given a profile P, let $\text{PW}(P) = |\{\vec{d} \in \mathcal{X} \mid \vec{d} = SSP_{\mathcal{O}'}(P)$ for some order $\mathcal{O}'\}|$. PW(P) is the number of different alternatives that can be made SSP winners by choosing a particular order $\mathcal{O}'$. Then, for a given number of binary issues p, we look for the maximal value of PW(P), for all profiles P on $\mathcal{X} = D_1 \times \ldots \times D_p$; we denote this number by MW(p).

A first observation is that there are $p!$ different choices for $\mathcal{O}'$. Therefore, a trivial upper bound on MW(p) is $p!$. Since there are 2^p alternatives, the $p!$ upper bound is only interesting when $p! < 2^p$, that is, $p \leq 3$. Example 2 shows that when $p = 2$ or $p = 3$, this trivial upper bound is actually tight, i.e. MW$(2) = 2!$ and MW$(3) = 3!$: there exists a profile such that by changing the order over the issues, all $p!$ different alternatives can be made winners. Due to McGarvey's theorem [16], any complete and asymmetric directed graph G over the alternatives corresponds to the majority graph of some profile (we recall that the majority graph of a profile P is the directed graph whose vertices are the alternatives and containing an edge from c to c' if and only if a majority of voters in P prefer c to c'). Therefore, in the example, we only show the majority graph instead of explicitly constructing the whole profile.

Example 2 *The majority graphs for $p = 2$ and $p = 3$ are shown in Figure 2. Let P (respectively, P') denote an arbitrary profile whose majority graph is the same as Figure 2(a) (respectively, Figure 2(b)). It is not hard to verify that $SSP_{\mathbf{x}_1>\mathbf{x}_2}(P) = 00$ and $SSP_{\mathbf{x}_2>\mathbf{x}_1}(P) = 01$. For P', the value of $SSP_{\mathcal{O}'}(P')$ for the six possible orders is shown on Table 1. Note that $2! = 2$ and $3! = 6$. It follows that when $p = 2$ or $p = 3$, there exists a profile for which the SSP winners w.r.t. different orders over the issues are all different from each other.*

The order	$\mathbf{x}_1 > \mathbf{x}_2 > \mathbf{x}_3$	$\mathbf{x}_1 > \mathbf{x}_3 > \mathbf{x}_2$	$\mathbf{x}_2 > \mathbf{x}_1 > \mathbf{x}_3$
SSP winner	010	011	001
The order	$\mathbf{x}_2 > \mathbf{x}_3 > \mathbf{x}_1$	$\mathbf{x}_3 > \mathbf{x}_1 > \mathbf{x}_2$	$\mathbf{x}_3 > \mathbf{x}_2 > \mathbf{x}_1$
SSP winner	100	110	101

Table 1: The SSP winners for P' w.r.t. different orders over the issues.

When $p \geq 4$, $p! > 2^p$. However, it is not immediately clear whether MW$(p) = 2^p$ or not, *i.e*, whether each of the 2^p alternatives can be made a winner by changing the order over the issues. The next theorem shows that this can actually be done, that is, MW$(p) = 2^p$.

Theorem 1 *For any $p \geq 4$ and any $n \geq 142 + 4p$, there exists an n-profile P such that for every alternative $\vec{d}$, there exists an order $\mathcal{O}'$ over $\mathcal{I}$ such that $SSP_{\mathcal{O}'}(P) = \vec{d}$.*

The proof of the theorem is by induction on p. Surprisingly, the hardest part in the inductive proof is the base case: when $p = 4$, we need to construct a profile P such that each of the 16 alternatives can be made a $SSP_{\mathcal{O}'}$ winner for at least one $\mathcal{O}'$ out of the 24 orders over four issues. We prove directly that it satisfies the desired property, and we also wrote a computer program to verify its correctness. The code (implemented in Java) can be found via the following link.
`http://www.cs.duke.edu/~lxia/Files/SSP.zip`

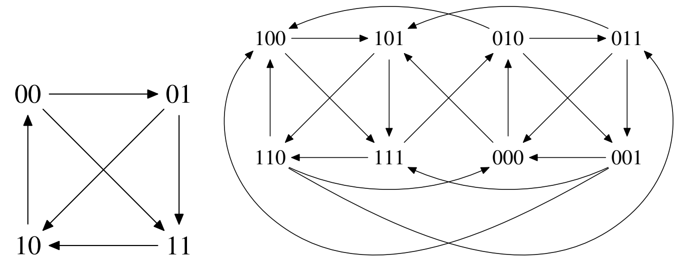

Figure 2: (a) The majority graph for $p = 2$. (b) The majority graph for $p = 3$, where four edges are not shown in the graph: $100 \to 000$, $101 \to 001$, $110 \to 010$, and $111 \to 011$. The directions of the other edges are defined arbitrarily. 000 is the abbreviation for $0_1 0_2 0_3$, etc.

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 [5, 14]. 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*.

Definition 2 *For any voting rule r, the* minimax satisfaction index (MSI) *of r is defined by*

$$MSI_r(m,n) = \min_{P \in L(\mathcal{X})^n} \max_{V \in P} \left(m + 1 - rank_V(r(P))\right)$$

where m is the number of alternatives, n is the number of voters, and $rank_V(r(P))$ is the position of $r(P)$ in vote V.

We note that in this paper $m = 2^p$, where p is the number of issues. 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 [14] and Benoit and Kornhauser [3] 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-binary-issue domains, for any natural number 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.

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 any stage of running the election sequentially. This is our main multiple-election paradox result.

Theorem 2 *For any $p \in \mathbb{N}$ ($p \geq 2$) and any $n \geq 2p^2+1$, $MSI_{SSP_{\mathcal{O}}}(m,n) = \lfloor p/2+2 \rfloor$.*[6] *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.*

Proof of Theorem 2: The upper bound on $MSI_{SSP_{\mathcal{O}}}(m,n)$ is constructive, that is, we explicitly construct a paradox.

For any n-profile $P=(V_1,\ldots,V_n)$, we define the mapping $f_P:\mathcal{X}\to\mathbb{N}^n$ as follows: for any $c\in\mathcal{X}$, $f_P(c)=(h_1,\ldots,h_n)$ such that for any $i\leq n$, h_i is the number of alternatives that are ranked below c in V_i. For any $l\leq p$, we denote $\mathcal{X}_l = D_l\times\cdots\times D_p$ and $\mathcal{O}_l = \mathbf{x}_l > \mathbf{x}_{l+1} > \cdots > \mathbf{x}_p$. For any vector $\vec{h}=(h_1,\ldots,h_n)$ and any $l\leq p$, we say that $\vec{h}$ is *realizable* over $\mathcal{X}_l$ (through a balanced binary tree) if there exists a profile $P_l=(V_1,\ldots,V_n)$ over $\mathcal{X}_l$ such that $f_{P_l}(SSP_{\mathcal{O}_l}(P_l))=\vec{h}$. We first prove the following lemma.

Lemma 1 *For any l such that $1\leq l<p$,*

$$\vec{h}_* = (\underbrace{0,\ldots,0}_{\lfloor n/2\rfloor-p+l},\underbrace{1,\ldots,1}_{p-l+1},\underbrace{2^{p-l+1}-1,\ldots,2^{p-l+1}-1}_{\lceil n/2\rceil-1})$$

is realizable over $\mathcal{X}_l$.

Proof of Lemma 1: We prove that there exists an n-profile P_l over $\mathcal{X}_l$ such that $SSP_{\mathcal{O}_l}(P_l)=1_l\cdots 1_p$ and $\vec{h}_*$ is realized by P_l. For any $1\leq i\leq p-l+1$, we let $\vec{b}_i = 1_l\cdots 1_{p-i}0_{p+1-i}1_{p+2-i}\cdots 1_p$. That is, $\vec{b}_i$ is obtained from $1_l\cdots 1_p$ by flipping the value of $\mathbf{x}_{p+1-i}$. We obtain $P_l=(V_1,\ldots,V_n)$ in the following steps.

1. Let $W_1,\ldots,W_n$ be null partial orders over $\mathcal{X}_l$. That is, for any $i\leq n$, the preference relation W_i is empty.
2. For any $j\leq\lfloor n/2\rfloor-p+l$, we put $1_l\cdots 1_p$ in the bottom position in W_j; we put $\{\vec{b}_1,\ldots,\vec{b}_{p-l+1}\}$ in the top positions in W_j.
3. For any j with $\lfloor n/2\rfloor+2\leq j\leq n$, we put $1_l\cdots 1_p$ in the top position of W_j, and we put $\{\vec{b}_1,\ldots,\vec{b}_{p-l+1}\}$ in the positions directly below the top.
4. For j with $\lfloor n/2\rfloor-p+l+1\leq j\leq\lfloor n/2\rfloor+1$, we define preferences as follows. For any $i\leq p-l+1$, in $W_{\lfloor n/2\rfloor-p+l+i}$, we put $\vec{b}_i$ in the bottom position, $1_l\cdots 1_p$ in the second position from the bottom, and all the remaining b_j (with $j\neq i$) at the very top.
5. Finally, we complete the profile arbitrarily: for any $j\leq n$, we let V_j be an arbitrary extension of W_j.

Let $P_l=(V_1,\ldots,V_n)$. We note that for any $i\leq p-l+1$, $\vec{b}_i$ beats any alternative in $\mathcal{X}_l\setminus\{1_l\cdots 1_p,\vec{b}_1,\ldots,\vec{b}_{p-l+1}\}$ in pairwise elections. Therefore, for any $i\leq p-l+1$, the ith alternative that meets $1_l\cdots 1_p$ is $\vec{b}_i$, which loses to $1_l\cdots 1_p$ (just barely). It follows that $1_l\cdots 1_p$ is the winner, and it is easy to check that $f_{P_l}(1_l\cdots 1_p)=\vec{h}_*$. This completes the proof of the lemma. □

Because the majority rule is anonymous, for any permutation π over $1,\ldots,n$ and any $l<p$, if $(h_1,\ldots,h_n)$ is realizable over $\mathcal{X}_l$, then $(h_{\pi(1)},\ldots,h_{\pi(n)})$ is also realizable over $\mathcal{X}_l$. For any $k\in\mathbb{N}$, we define $H_k=\{\vec{h}\in\{0,1\}^n:\sum_{j\leq n}h_j\geq k\}$. That is, H_k is composed of all n-dimensional binary vectors in each of which at least k components are 1. We next show a lemma to derive a realizable vector over $\mathcal{X}_{l-1}$ from two realizable vectors over $\mathcal{X}_l$.

Lemma 2 *Let $l<p$, and let $\vec{h}_1$, $\vec{h}_2$ be vectors that are realizable over $\mathcal{X}_l$. For any $\vec{h}\in H_{\lfloor n/2\rfloor+1}$, $\vec{h}_1+(\vec{h}_2+\vec{1})\cdot\vec{h}$ is realizable over $\mathcal{X}_{l-1}$, where $\vec{1}=(1,\ldots,1)$, and for any $\vec{a}=(a_1,\ldots,a_n)$ and any $\vec{b}=(b_1,\ldots,b_n)$, we have $\vec{a}\cdot\vec{b}=(a_1b_1,\ldots,a_nb_n)$.*

[6] If 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.

Proof of Lemma 2: Without loss of generality, we prove the lemma for $\vec{h}=(\underbrace{0,\ldots,0}_{\lceil n/2\rceil-1},\underbrace{1,\ldots,1}_{\lfloor n/2\rfloor+1})$. Let P_1,P_2 be two profiles over $\mathcal{X}_l$, each of which is composed of n votes, such that $f(P_1)=\vec{h}_1$ and $f(P_2)=\vec{h}_2$. Let $P_1=(V_1^1,\ldots,V_n^1)$, $P_2=(V_1^2,\ldots,V_n^2)$, $\vec{a}=SSP_{\mathcal{O}_l}(P_1)$, $\vec{b}=SSP_{\mathcal{O}_l}(P_2)$. We define a profile $P=(V_1,\ldots,V_n)$ over $\mathcal{X}_{l-1}$ as follows.

1. Let $W_1,\ldots,W_n$ be n null partial orders over $\mathcal{X}_{l-1}$.
2. For any $j\leq n$ and any $\vec{e}_1,\vec{e}_2\in\mathcal{X}_l$, we let $(1_{l-1},\vec{e}_1)\succ_{W_j}(1_{l-1},\vec{e}_2)$ if $\vec{e}_1\succ_{V_j^1}\vec{e}_2$; and we let $(0_{l-1},\vec{e}_1)\succ_{W_j}(0_{l-1},\vec{e}_2)$ if $\vec{e}_1\succ_{V_j^2}\vec{e}_2$.
3. For any $\lceil n/2\rceil\leq j\leq n$, we let $(1_{l-1},\vec{a})\succ_{W_j}(0_{l-1},\vec{b})$.
4. Finally, we complete the profile arbitrarily: for any $j\leq n$, we let V_j be an (arbitrary) extension of W_j such that $(1_{l-1},\vec{a})$ is ranked as low as possible.

We note that $(1_{l-1},\vec{a})$ is the winner of the subtree in which $\mathbf{x}_{l-1}=1_{l-1}$, $(0_{l-1},\vec{b})$ is the winner of the subtree in which $\mathbf{x}_{l-1}=0_{l_1}$, and $(1_{l-1},\vec{a})$ beats $(0_{l-1},\vec{b})$ in their pairwise election (because the votes from $\lceil n/2\rceil$ to n rank $(1_{l-1},\vec{a})$ above $(0_{l-1},\vec{b})$). Therefore, $SSP_{\mathcal{O}_{l-1}}(P)=(1_{l-1},\vec{a})$.

Finally, we have that $f_P((1_{l-1},\vec{a}))=\vec{h}_1+(\vec{h}_2+\vec{1})\cdot\vec{h}$. This is because $(1_{l-1},\vec{a})$ is ranked just as low as in the profile P_1 for voters 1 through $\lceil n/2\rceil-1$; for any voter j with $\lceil n/2\rceil\leq j\leq n$, additionally, $(0_{l-1},\vec{b})$ needs to be placed below $(1_{l-1},\vec{a})$, which implies that also, all the alternatives $(0_{l-1},\vec{b}')$ for which j ranked $\vec{b}'$ below $\vec{b}$ in P_2 must be below $(1_{l-1},\vec{a})$ in j's new vote in P. This completes the proof of the lemma. □

Now we are ready to prove the main part of the theorem. It suffices to prove that for any $n\geq 2p^2+1$, there exists a vector $\vec{h}_p\in\mathbb{N}^n$ such that each component of $\vec{h}_p$ is no more than $\lfloor p/2+1\rfloor$, and $\vec{h}_p$ is realizable over $\mathcal{X}$. We first prove the theorem for the case in which n is odd. We show the construction by induction in the proof of the following lemma.

Lemma 3 *Let n be odd. For any $l'<p$ (such that l' is odd),*

$$\vec{h}_{l'} = (\underbrace{\lfloor l'/2\rfloor,\ldots,\lfloor l'/2\rfloor}_{\lceil n/2\rceil-(l'^2+1)/2},\underbrace{\lceil l'/2\rceil,\ldots,\lceil l'/2\rceil}_{\lfloor n/2\rfloor+(l'^2+1)/2})$$

is realizable over $\mathcal{X}_{p-l'+1}$, and if $l'<p$, then

$$\vec{h}_{l'+1} = (\underbrace{\lfloor l'/2\rfloor,\ldots,\lfloor l'/2\rfloor}_{l'+1},\underbrace{\lceil l'/2\rceil,\ldots,\lceil l'/2\rceil}_{n-(l'+5)(l'+1)/2},\underbrace{\lceil l'/2\rceil+1,\ldots,\lceil l'/2\rceil+1}_{(l'+3)(l'+1)/2})$$

is realizable over $\mathcal{X}_{p-l'}$.

Proof of Lemma 3: The base case in which $l'=1$ corresponds to a single-issue majority election over two alternatives, where $\lceil n/2\rceil-1$ voters vote for one alternative, and $\lfloor n/2\rfloor+1$ vote for the other, so that only the latter get their preferred alternative.

Now, suppose the claim holds for some $l'\leq p-2$; we next show that the claim also holds for $l'+2$. To this end, we apply Lemma 2 twice. Let $l=p-l'+1$.

First, let

$$\vec{h}_* = (\underbrace{1,\ldots,1}_{l'},\underbrace{2^{l'}-1,\ldots,2^{l'}-1}_{l'+1},\underbrace{0,\ldots,0}_{\lfloor n/2\rfloor-l'+1},\underbrace{2^{l'}-1,\ldots,2^{l'}-1}_{\lceil n/2\rceil-l'-2})$$

By Lemma 1, $\vec{h}_*$ is realizable over $\mathcal{X}_l$ (via a permutation of the voters). Let $\vec{h} = (\underbrace{1,\ldots,1}_{l'},\underbrace{0,\ldots,0}_{l'+1},\underbrace{1,\ldots,1}_{\lfloor n/2\rfloor - l'+1},\underbrace{0,\ldots,0}_{\lceil n/2\rceil - l'-2})$.

Then, by Lemma 2, $\vec{h}_{l'} + (\vec{h}_* + \vec{1}) \cdot \vec{h}$ is realizable over $\mathcal{X}_{l-1}$. We have the following calculation.

$$
\begin{aligned}
&\vec{h}_{l'} + (\vec{h}_* + \vec{1}) \cdot \vec{h}\\
=&(\underbrace{\lceil l'/2\rceil+1,\ldots,\lceil l'/2\rceil+1}_{l'},\\
&\underbrace{\lfloor l'/2\rfloor,\ldots,\lfloor l'/2\rfloor}_{l'+1},\underbrace{\lceil l'/2\rceil,\ldots,\lceil l'/2\rceil}_{\lceil n/2\rceil-(l'+3)(l'+1)/2},\\
&\underbrace{\lceil l'/2\rceil+1,\ldots,\lceil l'/2\rceil+1}_{(l'+1)^2/2+1},\underbrace{\lceil l'/2\rceil,\ldots,\lceil l'/2\rceil}_{\lfloor n/2\rfloor - l'-1})
\end{aligned}
$$

The partition of the set of voters into these five groups uses the fact that $n \geq 2p^2+1$ implies $\lceil n/2\rceil - (l'+3)(l'+1)/2 \geq 0$. After permuting the voters in this vector, we obtain the following vector which is realizable over $\mathcal{X}_{l-1}$:

$$
\begin{aligned}
\vec{h}_{l'+1} =&(\underbrace{\lfloor l'/2\rfloor,\ldots,\lfloor l'/2\rfloor}_{l'+1},\underbrace{\lceil l'/2\rceil,\ldots,\lceil l'/2\rceil}_{n-(l'+5)(l'+1)/2},\\
&\underbrace{\lceil l'/2\rceil+1,\ldots,\lceil l'/2\rceil+1}_{(l'+3)(l'+1)/2})
\end{aligned}
$$

We next let $\vec{h}' = (\underbrace{1,\ldots,1}_{\lfloor n/2\rfloor+1},\underbrace{0,\ldots,0}_{\lceil n/2\rceil-1})$ and

$$
\vec{h}'_* = (\underbrace{1,\ldots,1}_{l'+1},\underbrace{0,\ldots,0}_{\lfloor n/2\rfloor - l'},\underbrace{2^{l'+1}-1,\ldots,2^{l'+1}-1}_{\lceil n/2\rceil-1})
$$

By Lemma 1, the latter is realizable over $\mathcal{X}_{l-1}$. Thus, by Lemma 2, $\vec{h}_{l'+1} + (\vec{h}'_* + \vec{1}) \cdot \vec{h}'$ is realizable over $\mathcal{X}_{l-2}$. Through a permutation over the voters, we obtain the desired vector:

$$
\vec{h}_{l'+2} = (\underbrace{\lfloor l'/2\rfloor+1,\ldots,\lfloor l'/2\rfloor+1}_{\lceil n/2\rceil-(l'+2)(l'+1)/2-1},\underbrace{\lceil l'/2\rceil+1,\ldots,\lceil l'/2\rceil+1}_{\lfloor n/2\rfloor+(l'+2)(l'+1)/2+1})
$$

which is realizable over $\mathcal{X}_{l-2}$. Therefore, the claim holds for $l'+2$. This completes the proof of the lemma. □

If p is odd, from Lemma 3 we know that the theorem is true, by setting $l' = p$. If p is even, then we first set $l' = p-1$; then, the maximum component of $\vec{h}_{l'+1}$ is $\lceil l'/2\rceil + 1 = \lceil (p-1)/2\rceil + 1 = p/2+1$. Thus we have proved the upper bound in the theorem when n is odd.

When n is even, we have the following lemma (the proof is similar to the proof of Lemma 3, so we omitted its proof).

Lemma 4 *Let n be even. For any $l' < p$ (such that l' is odd),*

$$
\vec{h}_{l'} = (\underbrace{\lfloor l'/2\rfloor,\ldots,\lfloor l'/2\rfloor}_{n/2-(l'^2-l'+1)/2},\underbrace{\lceil l'/2\rceil,\ldots,\lceil l'/2\rceil}_{n/2+(l'^2-l'+1)/2})
$$

is realizable over $\mathcal{X}_{p-l'+1}$, and if $l'+1 \leq p$, then

$$
\begin{aligned}
\vec{h}_{l'+1} =&(\underbrace{\lfloor l'/2\rfloor,\ldots,\lfloor l'/2\rfloor}_{l'+1},\underbrace{\lceil l'/2\rceil,\ldots,\lceil l'/2\rceil}_{n-1-(l'+4)(l'+1)/2},\\
&\underbrace{\lceil l'/2\rceil+1,\ldots,\lceil l'/2\rceil+1}_{(l'+2)(l'+1)/2+1})
\end{aligned}
$$

is realizable over $\mathcal{X}_{p-l'}$.

The upper bound in the theorem when n is even follows from Lemma 4. Moreover, we note that in the step from l' to $l'+1$ (respectively, from $l'+1$ to $l'+2$), no more than l' new alternatives are ranked lower than the winner in the profile that realizes $\vec{h}_{l'+1}$ (respectively, $\vec{h}_{l'+2}$). It follows that in the profile that realizes $\vec{h}_{l'+1}$ (respectively, $\vec{h}_{l'+2}$) in Lemma 3 or Lemma 4, the number of alternatives that are ranked lower than the winner by at least one voter is no more than $(l'+1)l'/2 + l' + 1 = (l'+1)(l'+2)/2$ (respectively, $(l'+2)(l'+3)/2$), which equals $(p+1)p/2$ if $l'+1 = p$ (respectively, $(p+1)p/2$ if $l'+2 = p$). Therefore, in the profile that we use to obtain the upper bound, the winner under $SSP_{\mathcal{O}}$ is Pareto-dominated by $2^p - (p+1)p/2$ alternatives.

Finally, we show that $\lfloor p/2+2\rfloor$ is a lower bound on $MSI_{SSP_{\mathcal{O}}}(m,n)$. Let P be an n-profile; let $SSP_{\mathcal{O}}(P) = \vec{a}$, and let $\vec{b}_1,\ldots,\vec{b}_p$ be the alternatives that $\vec{a}$ defeats in pairwise elections in rounds $1,\ldots,p$. It follows that in round j, more than half of the voters prefer $\vec{a}$ to $\vec{b}_j$, because we assume that there are no ties in the election. Therefore, summing over all votes, there are at least $p \times (\lfloor n/2\rfloor + 1)$ occasions where $\vec{a}$ is preferred to one of $\vec{b}_1,\ldots,\vec{b}_p$. It follows that there exists some $V \in P$ in which $\vec{a}$ is ranked higher than at least $\lceil p \times (\lfloor n/2\rfloor + 1)/n\rceil \geq \lfloor p/2+1\rfloor$ of the alternatives $\vec{b}_1,\ldots,\vec{b}_p$. Thus $MSI_{SSP_{\mathcal{O}}}(m,n) \geq \lfloor p/2+2\rfloor$.

(End of proof for Theorem 2.) □

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 whether 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 ordering of the issues.[7] Unfortunately, our next result shows that we can construct a single profile that results in a paradox for *all* orderings of the issues. While it works for all orders, the result is otherwise somewhat weaker than Theorem 2: 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 slightly higher than in Theorem 2, and unlike Theorem 2, no matching lower bound is shown.

Theorem 3 *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 $\{\mathbf{x}_1,\ldots,\mathbf{x}_p\}$, $SSP_{\mathcal{O}}(P) = 1_1 \cdots 1_p$, and any $V \in P$ ranks $1_1 \cdots 1_p$ somewhere in the bottom $p+2$ positions.*

Proof of Theorem 3: We first prove a lemma.

Lemma 5 *For any $c \in \mathcal{X}$, $\mathcal{C} \subset \mathcal{X}$ such that $c \notin \mathcal{C}$, and any $n \in \mathbb{N}$ ($n \geq 2m = 2^{p+1}$), there exists an n-profile that satisfies the following conditions. Let $F = \mathcal{X} \setminus (\mathcal{C} \cup \{c\})$.*

- *For any $c' \in \mathcal{C}$, c defeats c' in their pairwise election.*
- *For any $c' \in \mathcal{C}$ and $d \in F$, c' defeats d in their pairwise election.*
- *For any $V \in P$, c is ranked somewhere in the bottom $|\mathcal{C}|+2$ positions.*

Proof of Lemma 5: We let $P = (V_1,\ldots,V_n)$ be the profile defined as follows. Let $F_1,\ldots,F_{\lfloor n/2\rfloor+1}$ be a partition of F such

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

that for any $j \leq \lfloor n/2 \rfloor + 1$, $|F_j| \leq \lceil 2m/n \rceil = 1$. For any $j \leq \lfloor n/2 \rfloor + 1$, we let $V_j = [(F \setminus F_j) \succ c \succ \mathcal{C} \succ F_j]$. For any $\lfloor n/2 \rfloor + 2 \leq j \leq n$, we let $V_j = [\mathcal{C} \succ F \succ c]$. It is easy to check that P satisfies all conditions in the lemma. □

Now, let $c = 1_1 \cdots 1_p$ and $\mathcal{C} = \{0_1 1_2 \cdots 1_p, 1_1 0_2 1_3 \cdots 1_p, \ldots, 1_1 \cdots 1_{p-1} 0_p\}$. By Lemma 5, there exists a profile P such that c beats any alternative in $\mathcal{C}$ in pairwise elections, any alternative in $\mathcal{C}$ beats any alternative in $\mathcal{X} \setminus (\mathcal{C} \cup \{c\})$ in pairwise elections, and c is ranked somewhere in the bottom $p+2$ positions. This is the profile that we will use to prove the paradox.

Without loss of generality , we assume that $\mathcal{O} = \mathbf{x}_1 > \mathbf{x}_2 \cdots > \mathbf{x}_p$. (This is without loss of generality because all issues have been treated symmetrically so far.) c beats $1_1 \cdots 1_{p-1} 0_p$ in the first round; c will meet $1_1 \cdots 1_{p-2} 0_{p-1} 1_p$ in the next pairwise election, because $1_1 \cdots 1_{p-2} 0_{p-1} 1_p$ beats every other alternative in that branch (they are all in F), and c will win; and so on. It follows that $c = SSP_{\mathcal{O}}(P)$. Moreover, all voters rank c in the bottom $p+2$ positions.

(End of proof for Theorem 3.) □

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, we may wonder if paradoxes are 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 [4, 19].

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 nearly as bad as the unrestricted case.

Theorem 4 *For any $n \geq 2p$, when the profile is separable, the MSI for $SSP_{\mathcal{O}}$ is between $2^{\lceil p/2 \rceil}$ and $2^{\lfloor p/2 \rfloor + 1}$.*

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

Theorem 5 *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 $\lim_{m \to \infty} (3m/4 + 1)/m = 3/4$, so in that sense there is no paradox when votes are $\mathcal{O}$-lexicographic.

Under the previous two restrictions (separability and $\mathcal{O}$-lexicographicity), $SSP_{\mathcal{O}}$ coincides with $Seq(maj, \ldots, maj)$ (by Corollary 1 and Proposition 3, respectively). Therefore, Theorems 4 and 5 also apply to sequential voting rules as defined in [15]; furthermore, Theorem 4 also applies to seat-by-seat voting [3].

Finally, we study the MSI for $SSP_{\mathcal{O}}$ when the profile is $\mathcal{O}$-legal. Theorem 6 shows that it is nearly as bad as the unrestricted case (Theorem 2). The proof of Theorem 6 is the most involved proof in the paper and is omitted due to the space constraint. The idea of the proof is similar to that of the proof for Theorem 2, but now we cannot apply Lemma 2, because $\mathcal{O}$-legality must be preserved. We start with a simpler result that shows the idea of the construction.

Claim 1 *There exists a way to break ties in $SSP_{\mathcal{O}}$ such that the following is true. Let $SSP'_{\mathcal{O}}$ be the rule corresponding to $SSP_{\mathcal{O}}$ plus the tiebreaking mechanism. For any $p \in \mathbb{N}$, there exists an $\mathcal{O}$-legal profile that consists of two votes, such that in one of the two votes, no more than $\lceil p/2 \rceil$ alternatives are ranked lower than the winner $SSP'_{\mathcal{O}}(P)$; and in the other vote, no more than $\lfloor p/2 \rfloor$ alternatives are ranked lower than $SSP'_{\mathcal{O}}(P)$.*

We emphasize that, unlike any of our other results, Claim 1 is based on a specific tie-breaking mechanism. The next theorem studies the more general and complicated case in which n can be either odd or even, and the winner does not depend on the tie-breaking mechanism. That is, there are no ties in the election. The situation is almost the same as in Theorem 2.

Theorem 6 *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 2 still applies when the profile is $\mathcal{O}$-legal, so together with Theorem 6 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. CONCLUSION AND FUTURE WORK

Combinatorial voting settings, in which the space of all alternatives is exponential in size, constitute an important area in which techniques from computer science can be fruitfully applied. Perhaps the simplest and most natural combinatorial voting setting is that of multi-issue domains, where the space of alternatives is the Cartesian product of the local domains. In practice, common decisions on multiple issues are often reached by voting on the issues sequentially. 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 voters have complete information about each other's preferences and their preferences are strict, the game can be solved by a natural backward induction process (WSDSBI), which leads to a unique solution. We showed that under some conditions on the preferences, this process leads to the same outcome as the truthful sequential voting, but in general it can result in very different outcomes. We analyzed the effect of changing the order over the issues that voters vote on and showed that, in some elections, every alternative can be made a winner by voting according to an appropriate order over the issues.

Most significantly, we showed that strategic sequential voting is prone to multiple-election paradoxes; to do so, we introduced a concept called 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 positions 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 extremely undesirable for every voter. 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.

There are many topics for future research. For example, is there any criterion on the selection of the order over the issues? Perhaps more importantly, how can we get around the multiple-election paradoxes in sequential voting games? For example, Theorem 5 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 4 and Theorem 6). 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?[8] How can we extend these results to incomplete-information settings?[9] Also, beyond 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?[10]

Acknowledgements

Lirong Xia is supported by a James B. Duke Fellowship. Lirong Xia and Vincent Conitzer acknowledge NSF CAREER 0953756 and IIS-0812113, and an Alfred P. Sloan fellowship for support. Jérôme Lang thanks the ANR project ComSoc (ANR-09-BLAN-0305). We thank Michel Le Breton, all EC-10, LOFT-10, and EC-11 reviewers, as well as the participants of the Dagstuhl Seminar 10101 "Computational Foundations of Social Choice", the 2010 Bertinoro Workshop on Frontiers in Mechanism Design, LOFT-10, and University of Amsterdam ILLC Computational Social Choice Seminar, for their helpful comments.

8. REFERENCES

[1] Salvador Barbera, Hugo Sonnenschein, and Lin Zhou. Voting by committees. *Econometrica*, 59(3):595–609, 1991.

[2] Marco Battaglini. Sequential voting with abstention. *Games and Economic Behavior*, 51:445–463, 2005.

[3] Jean-Pierre Benoit and Lewis A. Kornhauser. Only a dictatorship is efficient or neutral. Technical report, 2006.

[4] Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

[5] Steven J. Brams, D. Marc Kilgour, and William S. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

[6] Vincent Conitzer and Tuomas Sandholm. Communication complexity of common voting rules. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 78–87, Vancouver, BC, Canada, 2005.

[7] Eddie Dekel and Michele Piccione. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, 108:34–55, 2000.

[8] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 347–356, Cambridge, MA, USA, 2010.

[9] Bhaskar Dutta and Arunava Sen. Implementing generalized Condorcet social choice functions via backward induction. *Social Choice and Welfare*, 10(2):149–160, 1993.

[10] Robin Farquharson. *Theory of Voting*. Yale University Press, 1969.

[11] Felix Fischer, Ariel D. Procaccia, and Alex Samorodnitsky. A new perspective on implementation by voting trees. In *Random Structures and Algorithms*, forthcoming, 2011. Early version appeared in *EC-10*.

[12] Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–601, 1973.

[13] Rodney J. Gretlei. Dominance elimination procedures on finite alternative games. *International Journal of Game Theory*, 12(2):107–113, 1983.

[14] Dean Lacy and Emerson M.S. Niou. A problem with referendums. *Journal of Theoretical Politics*, 12(1):5–31, 2000.

[15] Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

[16] David C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.

[17] 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.

[18] Hervé Moulin. Dominance solvable voting schemes. *Econometrica*, 47(6):1137–51, 1979.

[19] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

[20] Hervé Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3:271–291, 1986.

[21] 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.

[22] 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.

[23] Marco Scarsini. A strong paradox of multiple elections. *Social Choice and Welfare*, 15(2):237–238, 1998.

[24] Birgitte Sloth. The theory of voting and equilibria in noncooperative games. *Games and Economic Behavior*, 5:152–169, 1993.

[25] 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.

[26] Lirong Xia and Vincent Conitzer. Strategy-proof voting rules over multi-issue domains with restricted preferences. In *Proceedings of the Sixth Workshop on Internet and Network Economics (WINE)*, pages 402–414, Stanford, CA, USA, 2010.

[8]In [26], Xia and Conitzer studied strategy-proof voting rules over multi-issue domains composed of not-necessarily-binary issues. Among other contributions, Xia and Conitzer characterized all strategy-proof voting rules when the votersąŕ preferences are lexicographic and their possible local preferences over each issue are restricted.

[9]Because the complete-information setting is a special case of incomplete-information-settings, our results, which are worst-case results, still hold in the latter.

[10]This may require quite restrictive conditions or a different notion of a paradox.