

Strategy-proof Voting Rules over Multi-issue Domains with Restricted Preferences

Lirong Xia and Vincent Conitzer

Department of Computer Science
Duke University
Durham, NC 27708, USA
{lxia, conitzer@cs.duke.edu}

Abstract. In this paper, we characterize strategy-proof voting rules when the set of alternatives has a multi-issue structure, and the voters' preferences are represented by acyclic CP-nets that follow a common order over issues. Our main result is a simple full characterization of strategy-proof voting rules satisfying non-imposition for a very natural restriction on preferences in multi-issue domains: we show that if the preference domain is lexicographic, then a voting rule satisfying non-imposition is strategy-proof if and only if it can be decomposed into multiple strategy-proof local rules, one for each issue and each setting of the issues preceding it. We also obtain the following variant of Gibbard-Satterthwaite: when there are at least two issues and each of the issues can take at least two values, then there is no non-dictatorial strategy-proof voting rule that satisfies non-imposition, even when the domain of voters' preferences is restricted to linear orders that are consistent with acyclic CP-nets following a common order over issues. This impossibility result follows from either one of two more general new impossibility results we obtained, which are not included in this paper due to the space constraint.

Keywords: Voting, multi-issue domains, strategy-proofness, lexicographic domains

1 Introduction

When agents have conflicting preferences over a set of alternatives, and they want to make a joint decision, a natural way to do so is by *voting*. Each agent (voter) is asked to report his or her preferences. Then, a *voting rule* is applied to the vector of submitted preferences to select a winning alternative. However, in some cases, a voter has an incentive to submit false preferences in order to change the winner to a more preferable alternative (to her). An instance of such misreporting is called a *manipulation*, and the perpetrating voter is called a *manipulator*. If there is no manipulation under a voting rule, then the rule is said to be *strategy-proof*.

Unfortunately, there are some very natural properties that are satisfied by no strategy-proof voting rule, according to the Gibbard-Satterthwaite theorem [16, 27]. The theorem states that when there are three or more alternatives, and any voter can choose *any* linear order over alternatives to represent her preferences, then no non-dictatorial voting rule that satisfies non-imposition is strategy-proof. A voting rule is dictatorial if the same

voter's most-preferred alternative is always chosen; it satisfies non-imposition if for every alternative, there exist *some* reported preferences that make that alternative win.

There are several approaches to circumventing this impossibility result. One that has received significant attention from computer scientists in recent years is to consider whether finding a manipulation is computationally hard under some rules. If so, then even though a manipulation is guaranteed to exist, it will perhaps not occur because the manipulator(s) cannot find it. Indeed, it has been shown that finding a manipulation is computationally hard (more precisely, NP-hard) for various rules, for various definitions of the manipulation problem (*e.g.*, [6, 5, 13, 17, 14, 36]). On the other hand, NP-hardness is a *worst-case* notion of hardness, so that it may very well be the case that *most* manipulations are easy to find. Various recent results suggest that this is indeed the case [25, 12, 24, 15, 37, 31, 30, 28, 34, 29, 18]. This paper does not fall under this line of research.

Instead, this paper falls under another, older, line of research on circumventing the Gibbard-Satterthwaite result. This line, which has been pursued mainly by economists, is to restrict the domain of preferences. That is, we assume that voters' preferences always lie in a restricted class. An example of such a class is that of *single-peaked* preferences [7]. For single-peaked preferences, desirable strategy-proof rules exist, such as the *median* rule. Other strategy-proof rules are also possible in this preference domain: for example, it is possible to add some artificial (*phantom*) votes before running the median rule. In fact, this characterizes all strategy-proof rules for single-peaked preferences [22]. On the other hand, preferences have to be significantly restricted to obtain such positive results: Aswal *et al.* [1] extend the Gibbard-Satterthwaite theorem, showing that if the preference domain is *linked*, then with three or more alternatives the only strategy-proof voting rule that satisfies non-imposition is a dictatorship.

In real life, the set of alternatives often has a multi-issue structure. That is, there are multiple *issues* (or *attributes*), each taking values in its respective domain, and an alternative is completely characterized by the values that the issues take. For example, consider a situation where the inhabitants of a county vote to determine a government plan. The plan is composed of multiple sub-plans for several interrelated issues, such as transportation, environment, and health [10]. Clearly, a voter's preferences for one issue in general depend on the decisions taken on the other issues: if a new highway is constructed through a forest, a voter may prefer a nature reserve to be established; but if the highway is not constructed, the voter may prefer that no nature reserve is established. As another example, in each US presidential election year, the president as well as members of the Senate and the House must be elected. In principle, a voter's preferences for a senator can depend on who is elected as president, for example if the voter prefers a balance of power between the Democratic and Republican parties. A straightforward way to aggregate preferences in multi-issue domains is *issue-by-issue* (a.k.a. *seat-by-seat*) voting, which requires that the voters explicitly express their preferences over each issue separately, after which each issue is decided by applying issue-wise voting rules independently. This makes sense if voters' preferences are *separable*, that is, each voter's preferences over a single issue are independent of her preferences over other issues. However, if preferences are not separable, it is not clear how the voter should vote in such an issue-by-issue election. Indeed, it is known that natural strategies for voting in such a context can lead to very undesirable results [10, 20].

The problem of characterizing strategy-proof voting rules in multi-issue domains has already received significant attention. Strategy-proof voting rules for high-dimensional single-peaked preferences (where each dimension can be seen as an issue) have been characterized [8, 2, 3, 23]. Barbera *et al.* [4] characterized strategy-proof voting rules when the voters’ preferences are separable, and each issue is binary (that is, the domain for each issue has two elements). Ju [19] studied multi-issue domains where each issue can take three values: “good”, “bad”, and “null”, and characterized all strategy-proof voting rules that satisfy *null-independence*, that is, if a voter votes “null” on an issue i , then her preferences over other issues do not affect the value of issue i .

The prior research that is closest to ours was performed by Le Breton and Sen [11]. They proved that if the voters’ preferences are separable, and the restricted preference domain of the voters satisfies a *richness* condition, then, a voting rule is strategy-proof if and only if it is an issue-by-issue voting rule, in which each issue-wise voting rule is strategy-proof over its respective domain.

Despite its elegance, the work by Le Breton and Sen is limited by the restrictiveness of separable preferences: as we have argued above, in general, a voter’s preferences on one issue depend on the decision taken on other issues. On the other hand, one would not necessarily expect the preferences for one issue to depend on every other issue. CP-nets [9] were developed in the artificial intelligence community as a natural representation language for capturing limited dependence in preferences over multiple issues. Recent work has started to investigate using CP-nets to represent preferences in voting contexts [26, 21, 35, 32]. If there is an order over issues such that every voter’s preferences for “later” issues depend only on the decisions made on “earlier” issues, then the voters’ CP-nets are acyclic, and a natural approach is to apply issue-wise voting rules *sequentially* [21]. While the assumption that such an order exists is still restrictive, it is much less restrictive than assuming that preferences are separable (for one, the resulting preference domain is exponentially larger [21]). Recent extensions of sequential voting rules include order-independent sequential voting [35], as well as frameworks for voting when preferences are modeled by general (that is, not necessarily acyclic) CP-nets [32, 33]. However, in this paper, we only study acyclic CP-nets that are consistent with a common order over the issues.

Our results. In this paper, we focus on multi-issue domains that are composed of at least two issues with at least two possible values each.¹ We first show that over *lexicographic* preference domains (where earlier issues dominate later issues in terms of importance to the voters), the class of strategy-proof voting rules that satisfy non-imposition is exactly the class of voting rules that can be decomposed into multiple strategy-proof local rules, one for each issue and each setting of the issues preceding it. Technically, it is exactly the class of all *conditional rule nets (CR-nets)*, defined later in this paper but analogous to CP-nets, whose local (issue-wise) entries are strategy-proof voting rules. CR-nets represent how the voting rule’s behavior on one issue depends on the decisions made on all issues preceding it. Conceptually, this is similar to how acyclic

¹ This is the standard assumption for studying voting in multi-issue domains, because otherwise either the domain can be simplified (by removing issues that only take one value), or it has no multi-issue structure (when there is only one issue).

CP-nets represent how a voter’s preferences on one issue depend on the decisions made on all issues preceding it.

Then, we prove an impossibility theorem, which is the following variant of Gibbard-Satterthwaite. When there are at least two issues with at least two values each, the only strategy-proof voting rule that satisfies non-imposition is a dictatorship. This result assumes that each voter is free to choose any linear order that corresponds to an acyclic CP-net that follows a common order over the issues. This impossibility result follows from either one of two more general new impossibility results that we do not include in this paper due to the space constraint.

We are not aware of any previous characterization or impossibility results for strategy-proof voting rules when voters’ preferences display dependencies across issues (that is, when they are modeled by CP-nets).

2 Preliminaries

In a voting setting (not necessarily one with multiple issues), let \mathcal{X} be the set of *alternatives* (or *candidates*). A linear order V on \mathcal{X} is a transitive, antisymmetric, and total relation on \mathcal{X} . The set of all linear orders on \mathcal{X} is denoted by $L(\mathcal{X})$. An n -voter profile P on \mathcal{X} consists of n linear orders on \mathcal{X} . That is, $P = (V_1, \dots, V_n)$, where for every $1 \leq j \leq n$, $V_j \in L(\mathcal{X})$. The set of all profiles on \mathcal{X} is denoted by $P(\mathcal{X})$. In this paper, we let n denote the number of voters. A (*voting*) *rule* r is a mapping from the set of all profiles on \mathcal{X} to \mathcal{X} , that is, $r : P(\mathcal{X}) \rightarrow \mathcal{X}$. For example, the *plurality* rule (also called the *majority* rule, when there are only two alternatives) chooses the alternative that is ranked in the top position in the most votes (with a tie-breaking mechanism, for example, ties are broken in alphabetical order—in this paper, it does not matter which tie-breaking mechanism we use). A voting rule r satisfies

- *unanimity*, if $\text{top}(V) = c$ for all $V \in P$ implies $r(P) = c$.
- *non-imposition*, if for any $c \in \mathcal{X}$, there exists an n -voter profile P such that $r(P) = c$.
- (*strong*) *monotonicity*, if for any pair of profiles $P = (V_1, \dots, V_n)$, $P' = (V'_1, \dots, V'_n)$ such that for any alternative c and any $1 \leq j \leq n$, we have $c \succ_{V'_j} r(P) \Rightarrow c \succ_{V_j} r(P)$, then, $r(P') = r(P)$.
- *strategy-proofness*, if there does not exist a pair (P, V'_j) , where P is a profile, and V'_j is a false vote of voter j , such that $r(P_{-j}, V'_j) \succ_{V'_j} r(P)$. That is, there is no profile where a voter can misrepresent her preferences to make herself better off.

In this paper, the set of all alternatives \mathcal{X} is a *multi-issue domain*. That is, let $\mathcal{I} = \{\mathbf{x}_1, \dots, \mathbf{x}_p\}$ be a set of *issues*, where each issue \mathbf{x}_i takes values in a *local domain*, denoted by D_i . An alternative is uniquely identified by its values on all issues, that is, $\mathcal{X} = D_1 \times \dots \times D_p$.

Example 1 *A group of people must make a joint decision on the menu for dinner (the caterer can only serve a single menu to everyone). The menu is composed of two issues: the main course (**M**) and the wine (**W**). There are three choices for the main course: beef (b), fish (f), or salad (s). The wine can be either red wine (r), white wine (w), or pink wine (p). The set of alternatives is a multi-issue domain: $\mathcal{X} = \{b, f, s\} \times \{r, w, p\}$.*

CP-nets [9] are a compact representation that captures dependencies across issues. In this paper, we use them not for their representational compactness, but rather as useful

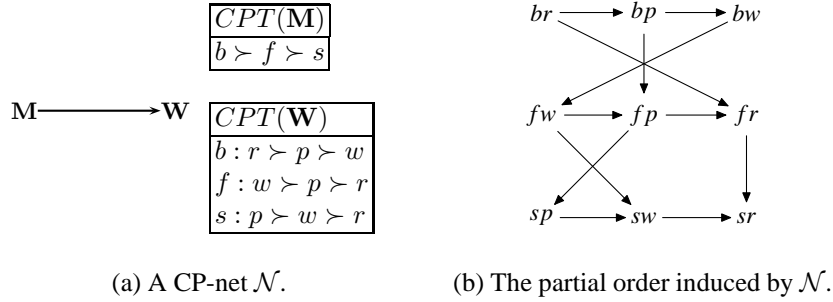


Fig. 1. A CP-net \mathcal{N} and its induced partial order.

mathematical notation for describing preferences in multi-issue domains, where preferences over one issue can depend on the values of earlier issues.

A CP-net \mathcal{N} over \mathcal{X} consists of two parts: (a) a directed graph $G = (\mathcal{I}, E)$ and (b) a set of conditional linear preferences $\succeq_{\mathbf{d}}^i$ over D_i , for each setting \mathbf{d} of the parents of \mathbf{x}_i in G . Let $CPT(\mathbf{x}_i)$ be the set of the conditional preferences of a voter on D_i ; this is called a *conditional preference table (CPT)*.

A CP-net \mathcal{N} captures dependencies across issues in the following sense. \mathcal{N} induces a partial preorder $\succeq_{\mathcal{N}}$ over the alternatives \mathcal{X} as follows: for any $a_i, b_i \in D_i$, any setting \mathbf{d} of the set of parents of \mathbf{x}_i (denoted by $Par_G(\mathbf{x}_i)$), and any setting \mathbf{z} of $\mathcal{I} \setminus (Par_G(\mathbf{x}_i) \cup \{\mathbf{x}_i\})$, $(a_i, \mathbf{d}, \mathbf{z}) \succeq_{\mathcal{N}} (b_i, \mathbf{d}, \mathbf{z})$ if and only if $a_i \succeq_{\mathbf{d}}^i b_i$. In words, the preferences over issue \mathbf{x}_i only depend on the setting of the parents of \mathbf{x}_i (but not on any other issues). For any $1 \leq i \leq p$, $CPT(\mathbf{x}_i)$ specifies conditional preferences over \mathbf{x}_i . Now, if we obtain an alternative \mathbf{d}' from \mathbf{d} by only changing the value of the i th issue of \mathbf{d} , we can look at $CPT(\mathbf{x}_i)$ to conclude whether the voter prefers \mathbf{d}' to \mathbf{d} , or vice versa. In general, however, from the CP-net, we will not always be able to conclude which of two alternatives a voter prefers, if the alternatives differ on two or more issues. This is why \mathcal{N} usually induces a partial preorder rather than a linear order.

We note that when the graph of \mathcal{N} is acyclic, $\succeq_{\mathcal{N}}$ is transitive and asymmetric, that is, a strict partial order. Let $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$. We say that a CP-net \mathcal{N} is *compatible* with (or, *follows*) \mathcal{O} , if \mathbf{x}_i being a parent of \mathbf{x}_j in the graph implies that $i < j$. That is, preferences over issues only depend on the values of earlier issues in \mathcal{O} . A CP-net is *separable* if there are no edges in its graph, which means that there are no preferential dependencies among issues.

Example 2 Let \mathcal{X} be the multi-issue domain defined in Example 1. We define a CP-net \mathcal{N} as follows: \mathbf{M} is the parent of \mathbf{W} , and the CPTs consist of the following conditional preferences: $CPT(\mathbf{M}) = \{b \succ f \succ s\}$, $CPT(\mathbf{W}) = \{b : r \succ p \succ w, f : w \succ p \succ r, s : p \succ w \succ r\}$, where $b : r \succ p \succ w$ is interpreted as follows: “when \mathbf{M} is b , then, r is the most preferred value for \mathbf{W} , p is the second most preferred value, and w is the least preferred value.” \mathcal{N} and its induced partial order $\succeq_{\mathcal{N}}$ are illustrated in Figure 1. \mathcal{N} is compatible with $\mathbf{M} > \mathbf{W}$. \mathcal{N} is not separable.

A linear order V over \mathcal{X} *extends* a CP-net \mathcal{N} , denoted by $V \sim \mathcal{N}$, if it extends the partial order that \mathcal{N} induces. (This is merely saying that V is consistent with the preferences implied by the CP-net \mathcal{N} .) V is *separable* if it extends a separable CP-net. The set of all linear orders that extend CP-nets that are compatible with \mathcal{O} is denoted by

$Legal(\mathcal{O})$. Throughout the paper, we make the following assumption about multi-issue domains and the voters' preferences.

Assumption 1 *In this paper, each multi-issue domain is composed of at least two issues ($p \geq 2$), and each issue can take at least two values. Moreover, all CP-nets are compatible with $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$, and the voters' preferences are always in $Legal(\mathcal{O})$ (that is, a voter's preferences over an issue do not depend on the values of later issues).*

To present our results, we will frequently use notations that represent the projection of a vote/CP-net/profile to an issue \mathbf{x}_i (that is, the voter's local preferences over \mathbf{x}_i) given the setting of all issues preceding \mathbf{x}_i , defined as follows. For any issue \mathbf{x}_i , any setting \mathbf{d} of $Par_G(\mathbf{x}_i)$, and any linear order V that extends \mathcal{N} , we let $V|_{\mathbf{x}_i:\mathbf{d}}$ and $\mathcal{N}|_{\mathbf{x}_i:\mathbf{d}}$ denote the the projection of V (or, equivalently \mathcal{N}) to \mathbf{x}_i , given \mathbf{d} . That is, each of these notations evaluates to the linear order $\succeq_{\mathbf{d}}^i$ in the CPT associated with \mathbf{x}_i . For example, let \mathcal{N} be the CP-net defined in Example 2. $\mathcal{N}|_{\mathbf{W}:b} = r \succ p \succ w$. For any \mathcal{O} -legal profile P , $P|_{\mathbf{x}_i:\mathbf{d}}$ is the profile over D_i that is composed of the projections of all votes in P on \mathbf{x}_i , given \mathbf{d} . That is, $P|_{\mathbf{x}_i:\mathbf{d}} = (V_1|_{\mathbf{x}_i:\mathbf{d}}, \dots, V_n|_{\mathbf{x}_i:\mathbf{d}}) = (\mathcal{N}_1|_{\mathbf{x}_i:\mathbf{d}}, \dots, \mathcal{N}_n|_{\mathbf{x}_i:\mathbf{d}})$, where $P = (V_1, \dots, V_n)$, and for any $1 \leq i \leq p$, V_i extends \mathcal{N}_i .

The *lexicographic extension* of a CP-net \mathcal{N} , denoted by $Lex(\mathcal{N})$, is a linear order V over \mathcal{X} such that for any $1 \leq i \leq p$, any $\mathbf{d}_i \in D_1 \times \dots \times D_{i-1}$, any $a_i, b_i \in D_i$, and any $\mathbf{y}, \mathbf{z} \in D_{i+1} \times \dots \times D_p$, if $a_i \succ_{\mathcal{N}|_{\mathbf{x}_i:\mathbf{d}_i}} b_i$, then $(\mathbf{d}_i, a_i, \mathbf{y}) \succ_V (\mathbf{d}_i, b_i, \mathbf{z})$. Intuitively, in the lexicographic extension of \mathcal{N} , \mathbf{x}_1 is the most important issue, \mathbf{x}_2 is the next important issue, etc; a desirable change to an earlier issue always outweighs any changes to later issues. We note that the lexicographic extension of any CP-net is unique w.r.t. the order \mathcal{O} . We say that $V \in L(\mathcal{X})$ is *lexicographic* if it is the lexicographic extension of a CP-net \mathcal{N} . For example, let \mathcal{N} be the CP-net defined in Example 2. We have $Lex(\mathcal{N}) = br \succ bp \succ bw \succ fw \succ fp \succ fr \succ sp \succ sw \succ sr$. A profile P is \mathcal{O} -legal/separable/lexicographic, if each of its votes is in $Legal(\mathcal{O})$ / is separable/ is lexicographic.

Given a vector of *local rules* (r_1, \dots, r_p) (that is, for any $1 \leq i \leq p$, 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 $Seq(r_1, \dots, r_p)$, is defined for all \mathcal{O} -legal profiles as follows: $Seq(r_1, \dots, r_p)(P) = (d_1, \dots, d_p) \in \mathcal{X}$, so that for any $1 \leq i \leq p$, $d_i = r_i(P|_{\mathbf{x}_i:d_1 \dots d_{i-1}})$. That is, 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 issues that precede \mathbf{x}_i . When the input profile is separable, $Seq(r_1, \dots, r_p)$ becomes an *issue-by-issue* voting rule.

3 Conditional rule nets (CR-nets)

We now move on to the contributions of this paper. In a sequential voting rule, the local voting rule that is used for a given issue is always the same, that is, the local voting *rule* does not depend on the decisions made on earlier issues (though, of course, the voters' *preferences* for this issue do depend on those decisions).

However, in many cases, it makes sense to let the local voting rules depend on the values of preceding issues. For example, let us consider again the setting in Example 1, and let us suppose that the caterer is collecting the votes and making the decision based on some rule. Suppose the order of voting is $\mathbf{M} > \mathbf{W}$. Suppose the main course is

determined to be beef. One would expect that, conditioning on beef being selected, most voters prefer red wine (e.g., $r \succ p \succ w$). Still, it can happen that even conditioned on beef being selected, surprisingly, slightly more than half the voters vote for white wine ($w \succ p \succ r$), and slightly less than half vote for red ($r \succ p \succ w$). In this case, the caterer, who knows that in the general population most people prefer red to white given a meal of beef, may “overrule” the preference for white wine among the slight majority of the voters, and select red wine anyway. While this may appear somewhat snobbish on the part of the caterer, in fact she may be acting in the best interest of social welfare if we take the non-voting agents (who are likely to prefer red given beef) into account.

In this section, we introduce *conditional rule nets* (CR-nets) to model voting rules where the local rules depend on the values chosen for earlier issues. A CR-net is defined similarly to a CP-net—the difference is that CPTs are replaced by conditional rule tables (CRTs), which specify a local voting rule over D_i for each issue x_i and each setting of the parents of x_i .²

Definition 1 An (acyclic) conditional rule net (CR-net) \mathcal{M} over \mathcal{X} is composed of the following two parts.

1. A directed acyclic graph G over $\{x_1, \dots, x_p\}$.

2. A set of conditional rule tables (CRTs) in which, for any variable x_i and any setting d of $\text{Par}_G(x_i)$, there is a local conditional voting rule $\mathcal{M}|_{x:d}$ over D_i .

A CR-net encodes a voting rule over all \mathcal{O} -legal profiles (we recall that we fix $\mathcal{O} = x_1 > \dots > x_p$ in this paper). For any $1 \leq i \leq p$, in the i th step, the value d_i is determined by applying $\mathcal{M}|_{x_i:d_1 \dots d_{i-1}}$ (the local rule specified by the CR-net for the i th issue given that the earlier issues take the values $d_1 \dots d_{i-1}$) to $P|_{x_i:d_1 \dots d_{i-1}}$ (the profile of preferences over the i th issue, given that the earlier issues take the values $d_1 \dots d_{i-1}$). Formally, for any \mathcal{O} -legal profile P , $\mathcal{M}(P) = (d_1, \dots, d_p)$ is defined as follows: $d_1 = \mathcal{M}|_{x_1}(P|_{x_1})$, $d_2 = \mathcal{M}|_{x_2:d_1}(P|_{x_2:d_1})$, etc. Finally, $d_p = \mathcal{M}|_{x_p:d_1 \dots d_{p-1}}(P|_{x_p:d_1 \dots d_{p-1}})$.

A CR-net \mathcal{M} is *separable* if there are no edges in the graph of \mathcal{M} . That is, the local voting rule for any issue is independent of the values of all other issues (which corresponds to a sequential voting rule).

4 Restricting voters’ preferences

We now consider restrictions on preferences. A restriction on preferences (for a single voter) rules out some of the possible preferences in $L(\mathcal{X})$. Following the convention of [11], a *preference domain* is a set of all admissible profiles, which represents the restricted preferences of the voters. Usually a preference domain is the Cartesian product of the sets of restricted preferences for individual voters. A natural way to restrict preferences in a multi-issue domain is to restrict the preferences on individual issues. For example, we may decide that $r \succ w \succ p$ is not a reasonable preference for wine (regardless of the choice of main course), and therefore rule it out (assume it away). More generally, which preferences are considered reasonable for one issue may depend on the decisions for the other issues. Hence, in general, for each i , for each setting d_i of the issues before issue x_i , there is a set of “reasonable” (or: possible, admissible) preferences over x_i , which we call $\mathcal{S}|_{x_i:d_i}$. Formally, *admissible conditional preference sets*, which encode all possible conditional preferences of voters, are defined as follows.

² It is not clear how a cyclic CR-net could be useful, so we only define acyclic CR-nets.

Definition 2 An admissible conditional preference set \mathcal{S} over \mathcal{X} is composed of multiple local conditional preference sets, denoted by $\mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$, such that for any $1 \leq i \leq p$ and any $\mathbf{d}_i \in D_1 \times \cdots \times D_{i-1}$, $\mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$ is a set of (not necessarily all) linear orders over D_i .

That is, for any $1 \leq i \leq p$ and any $\mathbf{d}_i \in D_1 \times \cdots \times D_{i-1}$, $\mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$ encodes the voter's local language over issue i , given the preceding issues taking values \mathbf{d}_i . In other words, if \mathcal{S} is the admissible conditional preference set for a voter, then we require the voter's preferences over \mathbf{x}_i given \mathbf{d}_i to be in $\mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$.

An admissible conditional preference set restricts the possible CP-nets, preferences, and lexicographic preferences. We note that Le Breton and Sen [11] defined a similar structure, which works specifically for separable votes.

Now we are ready to define the restricted preferences of a voter over \mathcal{X} . Let \mathcal{S} be the admissible conditional preference set for the voter. A voter's admissible vote can be generated in the following two steps: first, a CP-net \mathcal{N} is constructed such that for any $1 \leq i \leq p$ and any $\mathbf{d}_i \in D_1 \times \cdots \times D_{i-1}$, the restriction of \mathcal{N} on \mathbf{x}_i given \mathbf{d}_i is chosen from $\mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$; second, an extension of \mathcal{N} is chosen as the voter's vote. By restricting the freedom in either of the two steps (or both), we obtain a set of restricted preferences for the voter. Hence, we have the following definitions.

Definition 3 Let \mathcal{S} be an admissible conditional preference set over \mathcal{X} .

- $CPnets(\mathcal{S}) = \{\mathcal{N} : \mathcal{N} \text{ is a CP-net over } \mathcal{X}, \text{ and } \forall i \forall \mathbf{d}_i \in D_1 \times \cdots \times D_{i-1}, \mathcal{N}|_{\mathbf{x}_i:\mathbf{d}_i} \in \mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}\}$.
- $Pref(\mathcal{S}) = \{V : V \sim \mathcal{N}, \mathcal{N} \in CPnets(\mathcal{S})\}$.
- $LD(\mathcal{S}) = \{Lex(\mathcal{N}) : \mathcal{N} \in CPnets(\mathcal{S})\}$.

That is, $CPnets(\mathcal{S})$ is the set of all CP-nets over \mathcal{X} corresponding to preferences that are consistent with the admissible conditional preference set \mathcal{S} . $Pref(\mathcal{S})$ is the set of all linear orders that are consistent with the admissible conditional preference set \mathcal{S} . $LD(\mathcal{S})$, which we call the *lexicographic preference domain*, is the subset of linear orders in $Pref(\mathcal{S})$ that are lexicographic. For any $L \subseteq Pref(\mathcal{S})$, we say that L extends \mathcal{S} if for any CP-net in $CPnets(\mathcal{S})$, there exists at least one linear order in L consistent with that CP-net. It follows that $LD(\mathcal{S})$ extends \mathcal{S} ; in this case, for any CP-net \mathcal{N} in $CPnets(\mathcal{S})$, there exists exactly one linear order in $LD(\mathcal{S})$ that extends \mathcal{N} . Lexicographic preference domains are natural extensions of admissible conditional preference sets, but they are also quite restrictive, since any CP-net only has one lexicographic extension.

We now define a notion of richness for admissible conditional preference sets. This notion says that for any issue, given any setting of the earlier issues, each value of the current issue can be the most-preferred one.³

Definition 4 An admissible conditional preference set \mathcal{S} is rich if for each $1 \leq i \leq p$, each valuation \mathbf{d}_i of the preceding issues, and each $a_i \in D_i$, there exists $V^i \in \mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$ such that a_i is ranked in the top position of V^i .

We remark that richness is a natural requirement, and it is also a very weak restriction in the following sense. It only requires that when a voter is asked about her (local) preferences over \mathbf{x}_i given \mathbf{d}_i , she should have the freedom to at least specify her most

³ This is *not* the same richness notion as the one proposed by Le Breton and Sen, which applies to preferences over all alternatives rather than to admissible conditional preference sets.

preferred local alternative in D_i at will. We note that $|\mathcal{S}|_{\mathbf{x}_i:d_i}$ can be as small as $|D_i|$ (by letting each alternative in D_i be ranked in the top position exactly once), which is in sharp contrast to $|L(D_i)| = |D_i|!$ (when all local orders are allowed).

A CR-net \mathcal{M} is *locally strategy-proof* if all its local conditional rules are strategy-proof over their respective local domains (we recall that the voters' local preferences must be in the corresponding local conditional preference set). That is, for any $1 \leq i \leq p$, $\mathbf{d}_i \in D_1 \times \cdots \times D_{i-1}$, $\mathcal{M}|_{\mathbf{x}_i:d_i}$ is strategy-proof over $\prod_{j=1}^n \mathcal{S}_j|_{\mathbf{x}_i:d_i}$.

5 Strategy-proof voting rules in lexicographic preference domains

In this section, we present our main theorem, which characterizes strategy-proof voting rules that satisfy non-imposition, when the voters' preferences are restricted to lexicographic preference domains. Our main theorem states the following: if each voter's preferences are restricted to the lexicographic preference domain for a rich admissible conditional preference set, then a voting rule that satisfies non-imposition is strategy-proof if and only if it is a locally strategy-proof CR-net. We recall that in this paper, there are at least two issues with at least two possible values each, and the lexicographic preference domain for a rich admissible conditional preference set \mathcal{S} is composed of all lexicographic extensions of the CP-nets that are constructed from \mathcal{S} .

Theorem 1 *Under Assumption 1, for any $1 \leq j \leq n$, suppose \mathcal{S}_j is a rich admissible conditional preference set, and voter j 's preferences are restricted to the lexicographic preference domain of \mathcal{S}_j . Then, a voting rule r that satisfies non-imposition is strategy-proof if and only if r is a locally strategy-proof CR-net.*

Sketch of Proof: The “if” part is easy. The “only if” part is proved by induction on p (the number of issues). More precisely, suppose the theorem holds for p issues. For $p+1$ issues, let r be a strategy-proof voting rule that satisfies non-imposition. We first prove that r can be decomposed in the following way: there exists a local rule r_1 over D_1 and a voting rule $r_{\mathbf{x}_{-1}:a_1}$ over $D_2 \times \cdots \times D_{p+1}$ for each $a_1 \in D_1$, such that for any profile P , the first component of $r(P)$ is determined by applying r_1 to the projection of P on \mathbf{x}_1 , and the remaining components are determined by applying $r_{\mathbf{x}_{-1}:a_1}$ to the restriction of P on the remaining issues given $\mathbf{x}_1 = a_1$, where a_1 is the first component of $r(P)$ (just determined by r_1). Moreover, we prove that r_1 and $r_{\mathbf{x}_{-1}:a_1}$ (for all $a_1 \in D_1$) satisfy non-imposition and strategy-proofness. Therefore, by the induction hypothesis, for each $a_1 \in D_1$, $r_{\mathbf{x}_{-1}:a_1}$ is a locally strategy-proof CR-net over $D_2 \times \cdots \times D_{p+1}$. It follows that r is a locally strategy-proof CR-net over $D_1 \times \cdots \times D_{p+1}$, in which the (unconditional) rule for \mathbf{x}_1 is r_1 , and given any $a_1 \in D_1$, the sub-CR-net conditioned on $\mathbf{x}_1 = a_1$ is $r_{\mathbf{x}_{-1}:a_1}$. \square

The proofs of all theorems are omitted due to the space constraint. All proofs can be found in the long version of this paper on the first author's website.

It follows from Theorem 1 that any sequential voting rule that is composed of locally strategy-proof voting rules is strategy-proof over lexicographic preference domains, because a sequential voting rule is a separable CR-net. Specifically, when the multi-issue domain is binary (that is, for any $1 \leq i \leq p$, $|D_i| = 2$), the sequential composition of majority rules is strategy-proof when the profiles are lexicographic. It is interesting to view this in the context of previous works on the strategy-proofness of sequential

composition of majority rules: Lacy and Niou [20] and Le Breton and Sen [11] showed that the sequential composition of majority rules is strategy-proof when the profile is restricted to the set of all separable profiles; on the other hand, Lang and Xia [21] showed that this rule is not strategy-proof when the profile is restricted to the set of all \mathcal{O} -legal profiles.

The restriction to lexicographic preferences is still limiting. Next, we investigate whether there is any other preference domain for the voters on which the set of strategy-proof voting rules that satisfy non-imposition is equivalent to the set of all locally strategy-proof CR-nets. The answer to this question is “No,” as shown in the next result. More precisely, over any preference domain that extends an admissible conditional preference set, the set of strategy-proof voting rules satisfying non-imposition and the set of locally strategy-proof CR-nets satisfying non-imposition are identical *if and only if* the preference domain is lexicographic.

Theorem 2 *Under Assumption 1, for any $1 \leq j \leq n$, suppose \mathcal{S}_j is a rich admissible conditional preference set, $L_j \subseteq \text{Pref}(\mathcal{S}_j)$, and L_j extends \mathcal{S}_j . If there exists $1 \leq j \leq n$ such that L_j is not the lexicographic preference domain of \mathcal{S}_j , then there exists a locally strategy-proof CR-net \mathcal{M} that satisfies non-imposition and is not strategy-proof over $\prod_{j=1}^n L_j$.*

6 An impossibility theorem

In this section, we present an impossibility theorem for strategy-proof voting rules when voters’ preferences are restricted to be \mathcal{O} -legal.

Theorem 3 *When the set of alternatives is a multi-issue domain, if each voter can choose any linear order in $\text{Legal}(\mathcal{O})$ to represent her preferences, then there is no strategy-proof voting rule that satisfies non-imposition, except a dictatorship.*

This impossibility theorem is a variant of the Gibbard-Satterthwaite theorem. We emphasize that there are at least two issues with at least two possible values each, and $\text{Legal}(\mathcal{O})$ is much smaller than the set of all linear orders over \mathcal{X} . Therefore, the theorem does *not* follow directly from Gibbard-Satterthwaite. It follows directly from either of the two stronger impossibility theorems proved in the full version of the paper: one is for extensions of lexicographic domains, and the other is for extensions of the “rich” domains defined by Le Breton and Sen [11]. Due to the space constraint and the heavy technicality and notation of the two impossibility theorems, we omit them.

We recall that Lang and Xia [21] showed that a specific sequential voting rule (the sequential composition of majority rules) is not strategy-proof when each voter can choose any linear order in $\text{Legal}(\mathcal{O})$ to represent her preferences. Theorem 3 is much stronger, in that it states that over such a preference domain, not only does the sequential composition of majority rules fail to be strategy-proof, but in fact all non-dictatorial voting rules that satisfy non-imposition fail to be strategy-proof; moreover, this holds for non-binary multi-issue domains as well.

7 Conclusion

In settings where a group of agents needs to make a joint decision, the set of alternatives often has a multi-issue structure. In this paper, we characterized strategy-proof

voting rules when the voters' preferences are represented by acyclic CP-nets that follow a common order over issues. We showed that if each voter's preferences are restricted to a lexicographic preference domain, then a voting rule satisfying non-imposition is strategy-proof if and only if it is a locally strategy-proof CR-net. We then proved that if the profile is allowed to be any \mathcal{O} -legal profile, then the only strategy-proof voting rules satisfying non-imposition are dictatorships.

Our result for lexicographic preferences is quite positive; however, beyond that, our results do not inspire much hope for desirable strategy-proof voting rules in multi-issue domains. Of course, it is well known that it is difficult to obtain strategy-proofness in voting settings in general, and this does not mean that we should abandon voting as a general method. Similarly, difficulties in obtaining desirable strategy-proof voting rules in multi-issue domains should not prevent us from studying voting rules for multi-issue domains altogether. From a mechanism design perspective, strategy-proofness is a very strong criterion, which corresponds to implementation in dominant strategies. It may well be the case that rules that are not strategy-proof still result in good outcomes in practice—or, more formally, in (say) Bayes-Nash equilibrium.

Acknowledgements

We thank Jérôme Lang and the anonymous reviewers for EC and WINE for helpful discussions and comments. Lirong Xia is supported by a James B. Duke Fellowship and Vincent Conitzer is supported by an Alfred P. Sloan Research Fellowship. This work is also supported by NSF under award numbers IIS-0812113 and CAREER 0953756.

References

1. Navin Aswal, Shurojit Chatterji, and Arunava Sen. Dictatorial domains. *Economic Theory*, 22(1):45–62, 2003.
2. Salvador Barberà, Faruk Gul, and Ennio Stacchetti. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61(2):262–289, 1993.
3. Salvador Barberà, Jordi Masso, and Alejandro Neme. Voting under constraints. *Journal of Economic Theory*, 76(2):298–321, 1997.
4. Salvador Barberà, Hugo Sonnenschein, and Lin Zhou. Voting by committees. *Econometrica*, 59(3):595–609, 1991.
5. John Bartholdi, III and James Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.
6. John Bartholdi, III, Craig Tovey, and Michael Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
7. Duncan Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.
8. Kim C. Border and J. S. Jordan. Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1):153–170, January 1983.
9. Craig Boutilier, Ronen Brafman, Carmel Domshlak, Holger Hoos, and David Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.
10. S. Brams, D. Kilgour, and W. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.
11. Michel Le Breton and Arunava Sen. Separable preferences, strategyproofness, and decomposability. *Econometrica*, 67(3):605–628, 1999.
12. Vincent Conitzer and Tuomas Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. of AAI*, pages 627–634, 2006.

13. Vincent Conitzer, Tuomas Sandholm, and Jérôme Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 1–33, 2007.
14. Piotr Faliszewski, Edith Hemaspaandra, and Henning Schnoor. Copeland voting: Ties matter. In *Proc. of AAMAS*, pages 983–990, 2008.
15. Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *Proc. of FOCS*, pages 243–249, 2008.
16. Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41:587–602, 1973.
17. Edith Hemaspaandra and Lane A. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.
18. Marcus Isaksson, Guy Kindler, and Elchanan Mossel. The geometry of manipulation – a quantitative proof of the Gibbard Satterthwaite theorem. In *Proc. of FOCS*, 2010.
19. Biung-Ghi Ju. A characterization of strategy-proof voting rules for separable weak orderings. *Social Choice and Welfare*, 21(3):469–499, 2003.
20. Dean Lacy and Emerson M.S. Niou. A problem with referendums. *Journal of Theoretical Politics*, 12(1):5–31, 2000.
21. Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.
22. Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
23. Klaus Nehring and Clemens Puppe. Efficient and strategy-proof voting rules: A characterization. *Games and Economic Behavior*, 59(1):132–153, 2007.
24. Ariel D. Procaccia and Jeffrey S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proc. of AAMAS*, 2007.
25. Ariel D. Procaccia and Jeffrey S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, February 2007.
26. F. Rossi, K. Venable, and T. Walsh. mCP nets: Representing and reasoning with preferences of multiple agents. In *Proc. of AAAI*, pages 729–734, 2004.
27. 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.
28. Toby Walsh. Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In *Proc. of IJCAI*, pages 324–329, 2009.
29. Toby Walsh. An empirical study of the manipulability of single transferable voting. In *Proc. of ECAI*, 2010.
30. Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proc. of EC*, pages 109–118, 2008.
31. Lirong Xia and Vincent Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proc. of EC*, pages 99–108, 2008.
32. Lirong Xia, Vincent Conitzer, and Jérôme Lang. Voting on multiattribute domains with cyclic preferential dependencies. In *Proc. of AAAI*, pages 202–207, 2008.
33. Lirong Xia, Vincent Conitzer, and Jérôme Lang. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proc. of AAMAS*, pages 399–408, 2010.
34. Lirong Xia, Vincent Conitzer, and Ariel D. Procaccia. A scheduling approach to coalitional manipulation. In *Proc. of EC*, pages 275–284, 2010.
35. Lirong Xia, Jérôme Lang, and Mingsheng Ying. Strongly decomposable voting rules on multiattribute domains. In *Proc. of AAAI*, 2007.
36. Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proc. of IJCAI*, 2009.
37. Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artif. Intell.*, 173(2):392–412, 2009.