

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.

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 *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, 34]). 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, 35, 30, 29, 28, 32, 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]. Here, it is assumed that each alternative is associated with a position in some space (for example, the alternative's position on a left-to-right political spectrum), and that voters always prefer alternatives that are closer to their most preferred alternative. That is, if a is voter i 's most-preferred alternative, and we have that a is in the leftmost position, b is in the middle position, and c is in the rightmost position, then voter i must prefer b to c . For single-peaked preferences, desirable strategy-proof rules exist, such as the *median* rule, which, if we assume for simplicity that the number of voters is odd, chooses the median of the voters' peaks (which is also the Condorcet winner). 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 the transportation, environment, and health [10]. Clearly, a voter's preferences for one issue in general depend on the decision 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 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 (because she voted "null" for 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.

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, 33, 31]. 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 [33], as well as frameworks for voting when preferences are modeled by general (that is, not necessarily acyclic) CP-nets [31, ?]. 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 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 included in Appendix C.

We are not aware of any previous characterization or impossibility results of 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 $c \in \mathcal{X}$ and any $n \in \mathbb{N}$, 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, $r(P') = r(P)$.

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

• *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, in each profile, no 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 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 up $CPT(\mathbf{x}_i)$ 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 is 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: **M** is the parent of **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 **M** is b, then, r is the most preferred value for **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) A CP-net \mathcal{N} .(b) The partial order induced by \mathcal{N} .**Fig. 1.** A CP-net \mathcal{N} and its induced partial order.

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 . These notations are 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. $\mathcal{N}|_{\mathbf{w}:b} = r > p > w$. 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 each vote 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, and so on; 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 $Lex(\mathcal{N}) = br > bp > bw > fw > fp > fr > sp > sw > sr$. 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 an 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 ($r \succ p \succ w$). If the caterer uses an unbiased rule, then presumably white wine will be selected. While this is in the interest of slightly more than half the voters and may therefore appear to be a good idea, consider now a setting where not everyone who will enjoy the meal is voting. For example, some people may not have been available at the time of the vote; some people may bring their spouses, who were not present for the voting; perhaps the caterer's (non-voting) crew will be able to eat some of the meal; *etc.* 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.

Of course, if a large majority of the voters prefer white wine given beef, then the caterer should not overrule this. This effectively comes down to a local rule where (say) at least 60% of voters need to prefer white wine for it to be selected given beef (equivalently, the caterer may add some "phantom votes" for red wine given beef, to represent the non-voting diners' likely preferences). Conversely, when fish is chosen, the caterer's rule for deciding the wine based on the votes may be slightly biased towards white wine. Hence, in this situation, it makes good sense for the local rule for wine to depend on the values of its parent (the main course), unlike in a typical sequential voting rule.

There are many other settings where we may wish to bias the rule for one issue conditioned on the decision for an earlier issue. For example, we may consider letting citizens vote for president first, and for vice-president second; but, given the choice of the president, his or her running mate would need to receive less than 30% of the vote to not be elected.

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 \mathbf{x}_i and each setting of the parents of \mathbf{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 $\{\mathbf{x}_1, \dots, \mathbf{x}_p\}$.
2. A set of conditional rule tables (CRTs) in which, for any variable \mathbf{x}_i and any setting \mathbf{d} of $\text{Par}_G(\mathbf{x}_i)$, there is a local conditional voting rule $\mathcal{M}|_{\mathbf{x}_i:\mathbf{d}}$ over D_i .

A CR-net encodes a voting rule over all \mathcal{O} -legal profiles (we recall that we fix $\mathcal{O} = \mathbf{x}_1 > \dots > \mathbf{x}_p$ in this paper). For any $1 \leq i \leq p$, in the i th step, the value d_i is determined by applying $\mathcal{M}|_{\mathbf{x}_i:d_1 \dots d_{i-1}}$ (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 $P|_{\mathbf{x}_i:d_1 \dots d_{i-1}}$ (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: $d_1 = \mathcal{M}|_{\mathbf{x}_1}(P|_{\mathbf{x}_1})$, $d_2 = \mathcal{M}|_{\mathbf{x}_2:d_1}(P|_{\mathbf{x}_2:d_1})$, etc. Finally, $d_p = \mathcal{M}|_{\mathbf{x}_p:d_1 \dots d_{p-1}}(P|_{\mathbf{x}_p:d_1 \dots d_{p-1}})$.

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 value 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 \mathbf{d}_i of the issues before issue \mathbf{x}_i , there is a set of “reasonable” (or: possible, admissible) preferences over \mathbf{x}_i , which we call $\mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$. Formally, *admissible conditional preference sets*, which encode all possible conditional preferences of voters, are defined as follows.

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 $\mathbf{d}_i \in D_1 \times \dots \times D_{i-1}$, $\mathcal{S}|_{\mathbf{x}_i:\mathbf{d}_i}$ is a set of (not necessarily all) linear orders over D_i .

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

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 \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 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 the voter's restricted preferences. 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 preferred local alternative in D_i at will. We note that $|\mathcal{S}_{|\mathbf{x}_i:\mathbf{d}_i}|$ 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).

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.

We now revisit our example and restrict the voters' preferences in a reasonable manner. We let the voters' preferences over any issue be single-peaked.

Example 3 Let the multi-issue domain \mathcal{X} be defined as in Example 1. Let \mathcal{S} be the admissible conditional preference set whose local conditional preference sets are single-peaked, as illustrated in Figure 2. That is, $\mathcal{S}|_{\mathbf{M}} = \{(b \succ s \succ f), (s \succ b \succ f), (s \succ f \succ b), (f \succ s \succ b)\}$ is the single-peaked preference domain in which the positions of b , s , and f are listed from left to right in the order on a straight line, and the preferences of a voter is specified by the distance from each alternative to the “peak” (the voter’s most preferred point) along the line; $\mathcal{S}|_{\mathbf{w}:b} = \mathcal{S}|_{\mathbf{w}:f} = \mathcal{S}|_{\mathbf{w}:s}$ are the single-peaked preference domains in which the positions of r , p , and w are listed from the left to the right in the order on a straight line (we note that in this example, these three local conditional preference sets are the same, but they can be different in general). \mathcal{S} is rich, because in single-peaked domains, each alternative is ranked in the top position in at least one linear order. The CP-net \mathcal{N} defined in Example 2 is not in $\text{CPnets}(\mathcal{S})$, because $(b \succ f \succ s) \notin \mathcal{S}|_{\mathbf{M}}$. Let \mathcal{N}' be a CP-net in which $\mathcal{N}'|_{\mathbf{M}} = b \succ s \succ f$, and all other conditional preferences are the same as in \mathcal{N} . Then, $\mathcal{N}' \in \text{CPnets}(\mathcal{S})$, and $\text{Lex}(\mathcal{N}') \in \text{Pref}(\mathcal{S})$.



Fig. 2. An admissible conditional preference set \mathcal{S} in which all local domains are single-peaked. Positions of the alternatives are shown in the figure.

Throughout the paper, we focus on the following preference domains: for each voter j (with $1 \leq j \leq n$), there is an admissible conditional preference set \mathcal{S}_j , and voter j 's preferences are restricted to a set of linear orders L_j that extends \mathcal{S}_j . We say all votes in L_j are *admissible*. Let L_Π be the set of all profiles, in each of which voter j 's preferences are chosen from L_j for any $1 \leq j \leq n$, that is, $L_\Pi = \prod_{j=1}^n L_j$. A CR-net \mathcal{M} is *locally strategy-proof* if all its local conditional rules are strategy-proof over respective local domains (we remember that the voters' local preferences must be in the corresponding local conditional preference set). That is, for any $1 \leq i \leq p$, $\mathbf{d}_i \in D_1 \times \cdots \times D_{i-1}$, $\mathcal{M}|_{\mathbf{x}_i:\mathbf{d}_i}$ is strategy-proof over $\prod_{j=1}^n \mathcal{S}_j|_{\mathbf{x}_i:\mathbf{d}_i}$.

We now propose a locally strategy-proof rule for our example that captures the idea of the caterer biasing the choice of wine.

Example 4 Let the multi-issue domain \mathcal{X} be defined as in Example 1, and let \mathcal{S} be defined as in Example 3. For any $1 \leq j \leq n$, let $\mathcal{S}_j = \mathcal{S}$. For any $0 \leq t \leq 1$, let r_t be the voting rule over a single-peaked preference domain that selects the alternative that is closest to the $(\lfloor t(n-1) \rfloor + 1)$ th leftmost value within the set of all voters' favorite values (peaks). For example, $r_{0.5}$ selects the alternative that is closest to the median value. Let \mathcal{M} be a CR-net defined as follows: $\mathcal{M}|_{\mathbf{M}} = r_{0.5}$, $\mathcal{M}|_{\mathbf{w}:b} = r_{0.1}$, $\mathcal{M}|_{\mathbf{w}:f} = r_{0.9}$.

$\mathcal{M}|_{\mathbf{w}:s} = r_{0.5}$.⁴ \mathcal{M} is locally strategy-proof given this restriction of preferences, because the local rules are strategy-proof for single-peaked preferences [22].

5 Strategy-proof voting rules in lexicographic preference domains

In this section, we present our main theorem. We characterize strategy-proof voting rules that satisfy non-imposition, when the voters' preferences are restricted to lexicographic preference domains.

We are now ready to present our main result, which 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 *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.*

The proofs of all theorems are relegated to the appendices.

Theorem 1 has some interesting corollaries. First, we note that a CR-net is computationally easy to apply, as long as each local rule is easy to apply. This suggests that strategy-proof voting rules that satisfy non-imposition over a lexicographic preference domain tend to be easy to apply.

Second, 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. This displays an interesting contrast to 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.

Of course, Theorem 1 allows for other strategy-proof voting rules besides sequences of majority rules, when preferences are lexicographic. For example, with binary issues, we can set different thresholds (instead of the 50% threshold of majority), and the threshold for an issue can depend on the decisions on the previous issues. With non-binary issues, if the preferences on each local domain are restricted to be single-peaked, then a sequence of median-voter rules is also strategy-proof; we can also add phantom

⁴ This rule is strongly biased towards red wine if beef is chosen, and towards white wine if fish is chosen, corresponding to a very snobby caterer.

voters [22], and again, which phantom voters we add for an issue can depend on the decisions on the previous issues.

Moreover, from Theorem 1 we immediately obtain the following impossibility theorem.

Corollary 1 *Suppose that each local domain has at least three elements and any voter is free to choose any lexicographic linear order. Any strategy-proof voting rule that satisfies non-imposition must select the winner in a sequence of p steps, as follows: in step i where $1 \leq i \leq p$, the value for \mathbf{x}_i is determined by applying a dictatorship on the voters' local preferences over \mathbf{x}_i . (The voters' preferences as well as which dictatorship is used can depend on the values of preceding issues.)*

Conversely, any local-dictatorship CR-net of the form described in Corollary 1 in fact is strategy-proof and satisfies non-imposition.

Of course, 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 proposition. 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 *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 appendices: one is for extensions of lexicographical domains (Theorem 6 in Appendix B), and the other is for extensions of the “rich” domains defined by Le Breton and Sen [11] (Appendix C). Due to the space constraint and heavy technicality and notation of the two impossibility theorems, they are relegated to the appendix.

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 $Legal(\mathcal{O})$ 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 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 *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 627–634, Boston, MA, USA, 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. Early versions in AAAI-02 and TARK-03.
14. Piotr Faliszewski, Edith Hemaspaandra, and Henning Schnoor. Copeland voting: Ties matter. In *Proceedings of AAMAS '08*, pages 983–990, 2008.
15. Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *FOCS-08*, 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-10*, 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. H. 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 *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Honolulu, HI, USA, 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 *Proceedings of AAI-04*, 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-09*, pages 324–329, 2009.
29. Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 109–118, Chicago, IL, USA, 2008.
30. Lirong Xia and Vincent Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 99–108, Chicago, IL, USA, 2008.
31. Lirong Xia, Vincent Conitzer, and Jérôme Lang. Voting on multiattribute domains with cyclic preferential dependencies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 202–207, Chicago, IL, USA, 2008.
32. Lirong Xia, Vincent Conitzer, and Ariel D. Procaccia. A scheduling approach to coalitional manipulation. In *Proc. of EC-10*, pages 275–284, 2010.
33. Lirong Xia, Jérôme Lang, and Mingsheng Ying. Strongly decomposable voting rules on multiattribute domains. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, BC, Canada, 2007.
34. Lirong Xia, Michael Zuckerman, Ariel D. Procaccia, Vincent Conitzer, and Jeffrey Rosenschein. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*, Pasadena, CA, USA, 2009.
35. Michael Zuckerman, Ariel D. Procaccia, and Jeffrey S. Rosenschein. Algorithms for the coalitional manipulation problem. *Artif. Intell.*, 173(2):392–412, 2009.

A Proofs

The next two easy lemmas are part of the folklore of strategy-proof voting and will be frequently used in the proofs of the main theorems (we include the proofs in the appendix for the convenience of the reader). Lemma 1 states that any strategy-proof rule r satisfies monotonicity, that is, for any profile P , if each voter changes her vote by ranking $r(P)$ higher, then the winner is still $r(P)$.

Lemma 1 (Known) *Any strategy-proof voting rule satisfies monotonicity.*

Proof of Lemma 1: Suppose for the sake of contradiction r is strategy-proof but does not satisfy monotonicity. It follows that there exists a profile P , i , and V'_i such that V'_i is obtained from V_i by raising $r(P)$, and $r(P_{-i}, V'_i) \neq r(P)$. If $r(P_{-i}, V'_i) \succ_{V'_i} r(P)$, then we must have that $r(P_{-i}, V'_i) \succ_{V_i} r(P)$, which means that voter i has incentive to falsely report that her true preferences are V'_i ; if $r(P) \succ_{V'_i} r(P_{-i}, V'_i)$, then when voter i 's true preferences are V'_i and the other voters' profile is P_{-i} , she has incentive to falsely report that her preferences are V_i . In either case there is a manipulation, which contradicts the assumption that r is strategy-proof. \square

Lemma 2 states that any strategy-proof rule r satisfying non-imposition always satisfies unanimity, that is, if all votes rank the same alternative first, that alternative wins.

Lemma 2 (Known) *Any strategy-proof voting rule that satisfies non-imposition also satisfies unanimity.*

Proof of Lemma 2: Suppose for the sake of contradiction r is strategy-proof and satisfies non-imposition, but r does not satisfy unanimity. There exist an alternative c and a profile $P = (V_1, \dots, V_n)$ such that c is ranked in the top position in each of V_j , but $r(P) \neq c$. Now, because r satisfies non-imposition, there exists a profile $Q = (W_1, \dots, W_n)$ such that $r(Q) = c$. For any $0 \leq j \leq n$, we let $P_j = (W_1, \dots, W_j, V_{j+1}, \dots, V_n)$. We note that $P_0 = P$ and $P_n = Q$. Therefore, there exists $j^* \leq n$ such that $r(P_{j^*-1}) \neq c$ and $r(P_{j^*}) = c$. It follows that when the true preferences of voter j^* is V_{j^*} , and the preferences of the other voters are as in P_{j^*} , voter j^* has incentive to falsely report that her true preferences is W_{j^*} , which can improve the outcome from $r(P_{j^*-1}) \neq c$ to c . This contradicts the assumption that r is strategy-proof. \square

Proof of Theorem 1: In the proofs of this paper, for any $1 \leq i \leq p$, we let \mathbf{x}_{-i} denote $\mathcal{I} \setminus \{\mathbf{x}_i\}$, and we let D_{-i} denote $D_1 \times \dots \times D_{i-1} \times D_{i+1} \times \dots \times D_p$. For any $1 \leq j \leq n$, any profile P of n votes, we let P_{-j} denote the profile that consists of all votes in P except the vote by voter j .

First, we prove the ‘‘only if’’ part by induction on p . When $p = 1$, the theorem is immediate. Now, suppose that the theorem holds when $p = k$. When $p = k + 1$, for any strategy-proof rule r that satisfies non-imposition, over $\mathcal{X}_{k+1} = D_1 \times \dots \times D_{k+1}$, we prove that this rule can be decomposed into two parts: first, it applies a local voting rule r_1 for \mathbf{x}_1 , and subsequently, it applies a rule $r|_{\mathbf{x}_{-1}:a_1}$ for \mathbf{x}_{-1} , which depends on the outcome of r_1 . Thus, we have the property that for any $P \in L_\Pi$, we have $r(P) = (r_1(P|_{\mathbf{x}_1}), r|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})}(P|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})}))$. Then, we will show that the induction assumption can be applied to the second part.

First, we claim that for any strategy-proof voting rule r satisfying non-imposition, and any $P \in L_{\Pi}$, the value of issue \mathbf{x}_1 for the winning alternative only depends on the restriction of the profile to \mathbf{x}_1 . That is, we show that for any pair of profiles $P, Q \in L_{\Pi}$, $P = (V_1, \dots, V_n)$, $Q = (W_1, \dots, W_n)$ and $P|_{\mathbf{x}_1} = Q|_{\mathbf{x}_1}$, we must have $r(P)|_{\mathbf{x}_1} = r(Q)|_{\mathbf{x}_1}$. Suppose on the contrary that $r(P)|_{\mathbf{x}_1} \neq r(Q)|_{\mathbf{x}_1}$. For any $0 \leq j \leq n$, we define $P_j = (W_1, \dots, W_j, V_{j+1}, \dots, V_n)$. It follows that $P_0 = P$ and $P_n = Q$. We claim that for any $0 \leq j \leq n-1$, $r(P_j)|_{\mathbf{x}_1} = r(P_{j+1})|_{\mathbf{x}_1}$. For the sake of contradiction, suppose $r(P_j)|_{\mathbf{x}_1} \neq r(P_{j+1})|_{\mathbf{x}_1}$ for some $j \leq n-1$. Let $a_1 = r(P_j)|_{\mathbf{x}_1}$ and $b_1 = r(P_{j+1})|_{\mathbf{x}_1}$. If $a_1 \succ_{V_{j+1}|_{\mathbf{x}_1}} b_1$, then, because $V_{j+1}|_{\mathbf{x}_1} = W_{j+1}|_{\mathbf{x}_1}$, (P_{j+1}, V_{j+1}) is a successful manipulation; on the other hand, if $b_1 \succ_{V_{j+1}|_{\mathbf{x}_1}} a_1$, then, (P_j, W_{j+1}) is a successful manipulation. This contradicts the strategy-proofness of r . Thus, we have shown that the value of issue \mathbf{x}_1 for the winning alternative only depends on the restriction of the profile to \mathbf{x}_1 .

Therefore, we can define a voting rule r_1 over D_1 as follows. For any $P^1 \in \prod_{j=1}^n \mathcal{S}_j|_{\mathbf{x}_1}$, $r_1(P^1) = r(P)|_{\mathbf{x}_1}$, where $P \in L_{\Pi}$ and $P|_{\mathbf{x}_1} = P^1$. Such a P exists because $LD(\mathcal{S}_j)$ extends \mathcal{S}_j for all j , and this is well-defined by the observation from the previous paragraph. r_1 satisfies non-imposition because r satisfies non-imposition.

Next, we prove that r_1 is strategy-proof. If we assume for the sake of contradiction that r_1 is not strategy-proof, then there exists a successful manipulation (P^1, \hat{V}_l^1) over D_1 , where voter l is the manipulator, and $P^1 = (V_1^1, \dots, V_n^1)$. Let $\mathcal{N}_1, \dots, \mathcal{N}_n, \hat{\mathcal{N}}_l$ be $n+1$ CP-nets satisfying the following conditions.

- For any $1 \leq j \leq n$, $\mathcal{N}_j|_{\mathbf{x}_1} = V_j^1$; $\hat{\mathcal{N}}_l|_{\mathbf{x}_1} = \hat{V}_l^1$.
- For any $1 \leq j \leq n$, $\mathcal{N}_j \in \text{CPnets}(\mathcal{S}_j)$, $\hat{\mathcal{N}}_l \in \text{CPnets}(\mathcal{S}_l)$.

For $1 \leq j \leq n$, let V_j be the lexicographic extension of \mathcal{N}_j . Let \hat{V}_l be the lexicographic extension of $\hat{\mathcal{N}}_l$. Let $P = (V_1, \dots, V_n)$. We note that the \mathbf{x}_1 component of $r(P_{-l}, \hat{V}_l)$ is $r_1(P_{-l}^1, \hat{V}_l^1) \succ_{V_l^1} r_1(P^1)$, which is the \mathbf{x}_1 component of $r(P)$. Because V_l is the lexicographic extension of \mathcal{N}_l , and $\mathcal{N}_l|_{\mathbf{x}_1} = V_l^1$, we have that $r(P_{-l}, \hat{V}_l) \succ_{V_l} r(P)$, which means that (P, \hat{V}_l) is a successful manipulation. This contradicts the strategy-proofness of r . So, we have shown that r_1 is strategy-proof.

We next show that the second part of r can be written as $r|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})}(P|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})})$ —that is, the rule for the remaining issues \mathbf{x}_{-1} only depends on the outcome for \mathbf{x}_1 . For any $V \in \text{Legal}(\mathcal{O})$ and any $a_1 \in D_1$, we let $V|_{\mathbf{x}_{-1}:a_1}$ denote the linear preference over D_{-1} that is compatible with the restriction of V to the set of alternatives whose \mathbf{x}_1 component is a_1 , that is, for any $\mathbf{a}_{-1}, \mathbf{b}_{-1} \in D_{-1}$, $\mathbf{a}_{-1} \succeq_{V|_{\mathbf{x}_{-1}:a_1}} \mathbf{b}_{-1}$ if and only if $(a_1, \mathbf{a}_{-1}) \succeq_V (a_1, \mathbf{b}_{-1})$. For any \mathcal{O} -legal profile P , $P|_{\mathbf{x}_{-1}:a_1}$ is composed of $V|_{\mathbf{x}_{-1}:a_1}$ for all $V \in P$. For any CP-net \mathcal{N} , we let $\mathcal{N}|_{\mathbf{x}_{-1}:a_1}$ denote the sub-CP-net of \mathcal{N} conditioned on $\mathbf{x}_1 = a_1$. It follows that if $V \sim \mathcal{N}$, then, $V|_{\mathbf{x}_{-1}:a_1} \sim \mathcal{N}|_{\mathbf{x}_{-1}:a_1}$.

Now, we claim that for any pair of profiles $P_1, P_2 \in L_{\Pi}$, $P_1 = (V_1, \dots, V_n)$ and $P_2 = (W_1, \dots, W_n)$, such that $a_1 = r_1(P_1) = r_1(P_2)$ and $P_1|_{\mathbf{x}_{-1}:a_1} = P_2|_{\mathbf{x}_{-1}:a_1}$, we must have $r(P_1) = r(P_2)$. To prove this, we construct a profile P such that $r(P_1) = r(P) = r(P_2)$. For any $1 \leq j \leq n$, we let $V_j^{a_1} \in \mathcal{S}_j|_{\mathbf{x}_1}$ be an arbitrary linear order over D_1 in which a_1 is in the top position. Let $P = (Q_1, \dots, Q_n) \in L_{\Pi}$ be the profile

in which for any $1 \leq j \leq n$, Q_j is the lexicographic extension of the CP-net \mathcal{N}_j that satisfies the following conditions.

- $\mathcal{N}_j|_{\mathbf{x}_1} = V_j^{a_1}$.
- $\mathcal{N}_j|_{\mathbf{x}_{-1}:a_1} = \hat{\mathcal{N}}_j|_{\mathbf{x}_{-1}:a_1}$, where $\hat{\mathcal{N}}_j$ is the CP-net that V_j extends.

Let $\mathbf{a} = (a_1, \mathbf{a}_{-1}) = r(P_1)$. For any $1 \leq j \leq n$ and any $\mathbf{b} \in \mathcal{X}$ with $\mathbf{b} \succ_{Q_j} \mathbf{a}$, we have that the \mathbf{x}_1 component of \mathbf{b} must be a_1 , because Q_j is lexicographic, and a_1 is in the top position of $Q_j|_{\mathbf{x}_1}$. We let $\mathbf{b} = (a_1, \mathbf{b}_{-1})$. It follows that $\mathbf{b}_{-1} \succ_{Q_j|_{\mathbf{x}_{-1}:a_1}} \mathbf{a}_{-1}$. We note that $Q_j|_{\mathbf{x}_{-1}:a_1}$ is the lexicographic extension of $\mathcal{N}_j|_{\mathbf{x}_{-1}:a_1}$, $V_j|_{\mathbf{x}_{-1}:a_1}$ is the lexicographic extension of $\hat{\mathcal{N}}_j|_{\mathbf{x}_{-1}:a_1}$, and

$\mathcal{N}_j|_{\mathbf{x}_{-1}:a_1} = \hat{\mathcal{N}}_j|_{\mathbf{x}_{-1}:a_1}$. Therefore, $Q_j|_{\mathbf{x}_{-1}:a_1} = V_j|_{\mathbf{x}_{-1}:a_1}$, which means that $\mathbf{b}_{-1} \succ_{V_j|_{\mathbf{x}_{-1}:a_1}} \mathbf{a}_{-1}$. Hence, we have $\mathbf{b} \succ_{V_j} \mathbf{a}$. By Lemma 1, we have $r(P) = r(P_1)$. By similar reasoning, $r(P) = r(P_2)$, which means that $r(P_1) = r(P) = r(P_2)$. It follows that for any $a_1 \in D_1$, there exists a voting rule $r|_{\mathbf{x}_{-1}:a_1}$ over $D_2 \times \cdots \times D_p$ such that for any $P \in L_{\Pi}$,

$$r(P) = (r_1(P|_{\mathbf{x}_1}), r|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})}(P|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})}))$$

At this point, we have shown that r can be decomposed as desired. We next show that for any $a_1 \in D_1$, $r|_{\mathbf{x}_{-1}:a_1}$ is strategy-proof over $\prod_{j=1}^n LD(\mathcal{S}_j|_{\mathbf{x}_{-1}:a_1})$. Suppose for the sake of contradiction that there exists a successful manipulation (P^{-1}, \hat{V}_l^{-1}) , where voter l is the manipulator, and $P^{-1} = (V_1^{-1}, \dots, V_n^{-1})$. Let $\mathcal{N}_1, \dots, \mathcal{N}_n, \hat{\mathcal{N}}_l$ be $n+1$ CP-nets satisfying the following conditions.

- For any $1 \leq j \leq n$, $\text{top}(\mathcal{N}_j|_{\mathbf{x}_1}) = a_1$. That is, a_1 is ranked in the top position in the restriction of \mathcal{N}_j to \mathbf{x}_1 . Also, $\text{top}(\hat{\mathcal{N}}_l|_{\mathbf{x}_1}) = a_1$.
- For any $1 \leq j \leq n$, $\mathcal{N}_j|_{\mathbf{x}_{-1}:a_1}$ is the CP-net over D_{-1} that V_j^{-1} extends; $\hat{\mathcal{N}}_l|_{\mathbf{x}_{-1}:a_1}$ is the CP-net over D_{-1} that \hat{V}_l^{-1} extends.
- For any $1 \leq j \leq n$, $\mathcal{N}_j \in \text{CPnets}(\mathcal{S}_j)$; $\hat{\mathcal{N}}_l \in \text{CPnets}(\mathcal{S}_l)$.

The existence of these CP-nets is guaranteed by the richness of \mathcal{S}_j for any $1 \leq j \leq n$. For any $1 \leq j \leq n$, let V_j be the lexicographic extension of \mathcal{N}_j . Let \hat{V}_l be the lexicographic extension of $\hat{\mathcal{N}}_l$. Let $P = (V_1, \dots, V_n)$. We note that

$$\begin{aligned} r(P) &= (r_1(P|_{\mathbf{x}_1}), r|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})}(P|_{\mathbf{x}_{-1}:r_1(P|_{\mathbf{x}_1})})) \\ &= (a_1, r|_{\mathbf{x}_{-1}:a_1}(P|_{\mathbf{x}_{-1}:a_1})) = (a_1, r|_{\mathbf{x}_{-1}:a_1}(P^{-1})) \\ &\prec_{V_l} (a_1, r|_{\mathbf{x}_{-1}:a_1}(P^{-1}, \hat{V}_l)) = r(P_{-l}, \hat{V}_l) \end{aligned}$$

This contradicts the strategy-proofness of r . Hence, we have shown that for any $a_1 \in D_1$, $r|_{\mathbf{x}_{-1}:a_1}$ is strategy-proof over $\prod_{j=1}^n LD(\mathcal{S}_j|_{\mathbf{x}_{-1}:a_1})$.

Moreover, because r satisfies non-imposition, for any $a_1 \in D_1$, $r|_{\mathbf{x}_{-1}:a_1}$ satisfies non-imposition. Hence, for any $a_1 \in D_1$, we can apply the induction assumption to $r|_{\mathbf{x}_{-1}:a_1}$ and conclude that it is a locally strategy-proof CR-net over D_{-1} . It follows that r is a locally strategy-proof CR-net over \mathcal{X} , completing the first part of the proof.

We next prove the ‘‘if’’ part. If the proposition does not hold, then there exists a locally strategy-proof CR-net \mathcal{M} for which there is a successful manipulation (P, \hat{V}_l) . Let $i \leq p$ be the smallest natural number such that $\mathcal{M}(P)|_{\mathbf{x}_i} \neq \mathcal{M}(P_{-l}, \hat{V}_l)|_{\mathbf{x}_i}$. Let \mathbf{d}_i be the first $i-1$ components of $\mathcal{M}(P)$ and $\mathcal{M}(P_{-l}, \hat{V}_l)$. Because $\mathcal{M}|_{\mathbf{x}_i:\mathbf{d}_i}$ is strategy-

proof, we have the following calculation.

$$\begin{aligned} \mathcal{M}(P)|_{\mathbf{x}_i} &= \mathcal{M}|_{\mathbf{x}_i:\mathbf{d}_i}(P|_{\mathbf{x}_i:\mathbf{d}_i}) \\ &\succ_{V_i|_{\mathbf{x}_i:\mathbf{d}_i}} \mathcal{M}|_{\mathbf{x}_i:\mathbf{d}_i}(P_{-1}, \hat{V}_i|_{\mathbf{x}_i:\mathbf{d}_i}) \\ &= \mathcal{M}(P_{-1}, \hat{V}_i)|_{\mathbf{x}_i} \end{aligned}$$

Because V_i is lexicographic, for any $\mathbf{y}, \mathbf{z} \in D_{i+1} \times \cdots \times D_p$, we have

$$(\mathbf{d}_i, \mathcal{M}|_{\mathbf{x}_i:\mathbf{d}_i}(P), \mathbf{y}) \succ_{V_i} (\mathbf{d}_i, \mathcal{M}|_{\mathbf{x}_i:\mathbf{d}_i}(P_{-1}, \hat{V}_i), \mathbf{z})$$

Therefore, $\mathcal{M}(P) \succ_{V_i} \mathcal{M}(P_{-1}, \hat{V}_i)$, which contradicts the assumption that (P, \hat{V}_i) is a successful manipulation. Hence, locally strategy-proof CR-nets are strategy-proof for lexicographic preferences. \square

Proof of Theorem 2: If, for some $j \leq n$, there is a $V_j' \in LD(\mathcal{S}_j)$ that is not in L_j , then there must also be a $V_j \in L_j$ that is not in $LD(\mathcal{S}_j)$, because some vote in L_j must extend the CP-net that V_j' extends. Hence, if $L_\Pi \neq \prod_{j=1}^n LD(\mathcal{S}_j)$, there must exist some $j \leq n$, $V_j \in L_j$ such that V_j is not in $LD(\mathcal{S}_j)$. For this V_j , there must exist $i \leq p$, $\mathbf{a}_{i-1} \in D_1 \times \cdots \times D_{i-1}$, $\mathbf{a}_i, \mathbf{b}_i \in D_i$, $\mathbf{a}_{i+1}, \mathbf{b}_{i+1} \in D_{i+1} \times \cdots \times D_p$ such that $\mathbf{a}_i \succ_{V_j|_{\mathbf{x}_i:\mathbf{a}_{i-1}}} \mathbf{b}_i$, and $(\mathbf{a}_{i-1}, \mathbf{b}_i, \mathbf{b}_{i+1}) \succ_{V_j} (\mathbf{a}_{i-1}, \mathbf{a}_i, \mathbf{a}_{i+1})$. Now, let us define a CR-net \mathcal{M} as follows.

- $\mathcal{M}|_{\mathbf{x}_i:\mathbf{a}_{i-1}}$ is the plurality rule that only counts voter 1 and voter j 's votes; ties are broken in the order $\mathbf{b}_i \succ \mathbf{a}_i \succ D_i - \{\mathbf{a}_i, \mathbf{b}_i\}$.
- Any other local conditional voting rule is a dictatorship by voter 1.

Now, let $\mathcal{N}_1 \in \text{CPnets}(\mathcal{S}_1)$ be a CP-net such that $\text{top}(\mathcal{N}_1) = \mathbf{a}_{i-1}\mathbf{a}_i\mathbf{a}_{i+1}$, and for any $k \geq i+1$, $\text{top}(\mathcal{N}_1|_{\mathbf{x}_k:\mathbf{a}_{i-1}\mathbf{b}_i\mathbf{b}_{i+1}\cdots\mathbf{b}_{k-1}}) = \mathbf{b}_k$. Let $\mathcal{N}_j' \in \text{CPnets}(\mathcal{S}_j)$ be a CP-net such that $\text{top}(\mathcal{N}_j') = \mathbf{a}_{i-1}\mathbf{b}_i\mathbf{b}_{i+1}$. Let $V_1 \in L_1$ be such that $V_1 \sim \mathcal{N}_1$, and let $V_j' \in L_j$ be such that $V_j' \sim \mathcal{N}_j'$. Such V_1 and V_j' must exist, because L_1 extends \mathcal{S}_1 , and L_j extends \mathcal{S}_j . For any profile $P = (V_1, \dots, V_j, \dots, V_n) \in L_\Pi$ (that is, for any $l \neq 1, j$, V_l is chosen arbitrarily, because $\mathcal{M}(P)$ does not depend on them), it follows that $\mathcal{M}(P) = \mathbf{a}_{i-1}\mathbf{a}_i\mathbf{a}_{i+1}$, and $\mathcal{M}(P_{-j}, V_j') = \mathbf{a}_{i-1}\mathbf{b}_i\mathbf{b}_{i+1}$, which means that (P, V_j') is a successful manipulation for voter j . So, \mathcal{M} is not strategy-proof (and it satisfies non-imposition). \square

B Impossibility result for extensions of lexicographic preference domains

Section 5 settles the case of lexicographic preferences, but preferences are not always lexicographic, even for acyclic CP-nets. For example, in a simplified menu example with beef, fish, red wine, and white wine, a red-wine fanatic may prefer $br \succ fr \succ bw \succ fw$. This is consistent with the order $\mathbf{M} > \mathbf{W}$ (in fact, the voter's preferences are separable), but the preferences are not lexicographic with respect to this order. In this section, we investigate the possibility of strategy-proof voting rules for supersets of a lexicographic preference domain. For any linear order V , we let $\text{Top}(V)$ denote the alternative that is ranked in the top position in V .

Definition 5 A CP-net \mathcal{N} is tops-only-separable if for any $1 \leq i \leq p$, $\mathbf{a}_i, \mathbf{b}_i \in D_1 \times \dots \times D_{i-1}$, $\text{top}(\mathcal{N}|_{\mathbf{x}_i:\mathbf{a}_i}) = \text{top}(\mathcal{N}|_{\mathbf{x}_i:\mathbf{b}_i})$.

That is, in a tops-only-separable CP-net, the most preferred value for any issue is independent of the values of the other issues (though there may be dependencies in the lower-ranked values).

We now give a condition on the preference domain that indicates that any issue can be considered more important than the first issue in some vote.

Definition 6 (Condition I) L_Π satisfies Condition I if for any $1 \leq j \leq n$, any $1 \leq i \leq p$, any $\mathbf{a} = (a_1, \dots, a_p) \in \mathcal{X}$, any $V_j^1 \in \mathcal{S}_j|_{\mathbf{x}_1}$ with $\text{top}(V_j^1) = a_1$, any $V_j^i \in \mathcal{S}_j|_{\mathbf{x}_i:a_1 \dots a_{i-1}}$ with $\text{top}(V_j^i) = a_i$, any $b_1 \in D_1$ ($b_1 \neq a_1$), and any $b_i \in D_i$ ($b_i \neq a_i$), there exists a tops-only-separable CP-net $\mathcal{N}_j \in \text{CPnets}(\mathcal{S}_j)$ and a vote $V_j \in L_j$ that extends \mathcal{N}_j , such that

- $\text{top}(\mathcal{N}_j) = \mathbf{a}$.
- $\mathcal{N}_j|_{\mathbf{x}_1} = V_j^1$, $\mathcal{N}_j|_{\mathbf{x}_i:a_1 \dots a_{i-1}} = V_j^i$.
- $(b_1, \mathbf{a}_{-1}) \succ_{V_j} (\mathbf{a}_{-i}, b_i)$.

Condition I may seem unnatural and hard to read at first glance, but we argue that it is actually quite a natural approach to capturing the idea that “each issue can be more important than the first issue in some vote.” In order for issue i to be more important than issue 1 for a voter, it should be the case that (roughly speaking), for any pair of alternatives $\mathbf{a}_* = (b_1, \mathbf{a}_{-1})$ and $\mathbf{b}_* = (\mathbf{a}_{-i}, b_i)$ (so that one differs from \mathbf{a} on the first issue, and one on the i th issue), the following is true: If it is the case that a_i is always preferred to b_i in the local preferences of the voter on issue i (regardless of the values of the preceding issues), then the voter prefers \mathbf{a}_* to \mathbf{b}_* —even if she prefers a_1 to b_1 .

In our definition of Condition I, requiring \mathcal{N} to be tops-only-separable implies that a_i is always preferred to b_i in the local preferences of the voter on issue i ; and because this argument should hold for any local preferences over \mathbf{x}_1 and \mathbf{x}_i , we require that we can choose V_j^1 and V_j^i freely in Condition I.

A similar notion was adopted by Le Breton and Sen [11] (see Definition 7 B(i) in this paper), but there they focus on separable profiles, which is significantly different (and more restrictive) from the preference domain studied in this paper. We also argue that Condition I is weaker than Condition B(i) in Definition 7 in some sense; see the discussion after Definition 7.

We also note that even if L_Π satisfies Condition I, it must be significantly smaller than the universal domain in which every voter is free to choose any linear order over \mathcal{X} . For example, the largest set that can satisfy Condition I is $\text{Legal}(\mathcal{O})$, and it has already been proved that the size of $\text{Legal}(\mathcal{O})$ is exponentially (by a power of $|\mathcal{X}| = 2^p$) smaller than the number of all linear orders over \mathcal{X} [31].

We now present the following impossibility result: if the preference domain satisfies Condition I and extends an admissible conditional preference set \mathcal{S} , then any locally strategy-proof CR-net either does not satisfy non-imposition, or it is a dictatorship.

Theorem 4 For any $1 \leq j \leq n$, suppose \mathcal{S}_j is a rich admissible conditional preference set, $L_j \subseteq \text{Pref}(\mathcal{S}_j)$, L_j extends \mathcal{S}_j , and L_Π satisfies Condition I. Then, for any locally

strategy-proof CR-net \mathcal{M} satisfying non-imposition, \mathcal{M} is strategy-proof over L_Π if and only if \mathcal{M} is a dictatorship.

Proof of Theorem 4: The “if” part is obvious, so we only prove the “only if” part. For any CR-net \mathcal{M} , and any $a_1 \in D_1$, we say that voter j is an a_1 -dictator if for any $1 \leq i \leq p$, any $\mathbf{a}_2 \in D_2 \times \cdots \times D_{i-1}$, we have that $\mathcal{M}|_{\mathbf{x}_i: a_1 \mathbf{a}_2}$ is a j -dictatorship (that is, the winner is always the alternative that is ranked in the top position by voter j). We first prove the following lemma.

Lemma 3 *Under the conditions of the theorem, let $P^1 = (V_1^1, \dots, V_n^1)$ be a profile in $\prod_{j=1}^n \mathcal{S}_j|_{\mathbf{x}_1}$, and let \mathcal{M} be a non-dictatorial locally strategy-proof CR-net satisfying non-imposition, with $\mathcal{M}|_{\mathbf{x}_1}(P^1) = a_1$. If there exist $j \leq n$ and $W_j^1 \in \mathcal{S}_j|_{\mathbf{x}_1}$ such that $\mathcal{M}|_{\mathbf{x}_1}(P^1) \neq \mathcal{M}|_{\mathbf{x}_1}(P_{-j}^1, W_j^1)$, and voter j is not an a_1 -dictator, then, \mathcal{M} is not strategy-proof.*

Proof of Lemma 3: Suppose on the contrary that there exists a non-dictatorial locally strategy-proof CR-net \mathcal{M} that satisfies non-imposition and is strategy-proof over L_Π , and satisfies all conditions in the lemma. Let $V_j^{a_1} \in \mathcal{S}_j|_{\mathbf{x}_1}$ be such that $\text{top}(V_j^{a_1}) = a_1$; then, it follows from the strategy-proofness of $\mathcal{M}|_{\mathbf{x}_1}$ and Lemma 1 that $\mathcal{M}|_{\mathbf{x}_1}(P_{-j}^1, V_j^{a_1}) = a_1$. Since voter j is not an a_1 -dictator, there exist $i^* \leq p$, $\mathbf{a}_2 = (a_2, \dots, a_{i^*-1}) \in D_2 \times \cdots \times D_{i^*-1}$, and a profile $P^{i^*} \in \prod_{j=1}^n \mathcal{S}_j|_{\mathbf{x}_{i^*}: a_1 \mathbf{a}_2}$ such that $\mathcal{M}|_{\mathbf{x}_{i^*}: a_1 \mathbf{a}_2}(P^{i^*}) \neq \text{top}(V_j^{i^*})$.

Let $a_{i^*} = \mathcal{M}|_{\mathbf{x}_{i^*}: a_1 \mathbf{a}_2}(P^{i^*})$. We arbitrarily choose

$$\overrightarrow{a_{i^*+1}} = (a_{i^*+1}, \dots, a_p) \in D_{i^*+1} \times \cdots \times D_p$$

Let $b_1 = \mathcal{M}|_{\mathbf{x}_1}(P_{-j}^1, W_j^1)$, $b_{i^*} = \text{top}(V_j^{i^*})$. Next, we construct a vector of CP-nets $\mathcal{N}_1, \dots, \mathcal{N}_n, \mathcal{N}'_j$ as follows.

- For any $l \neq j$, $\mathcal{N}_l|_{\mathbf{x}_1} = V_l^1$, $\mathcal{N}_l|_{\mathbf{x}_{i^*}: a_1 \mathbf{a}_2} = V_l^{i^*}$;
 $\text{top}(\mathcal{N}_l|_{\mathbf{x}_{-1}: a_1}) = \mathbf{a}_2 \text{top}(V_l^{i^*}) \overrightarrow{a_{i^*+1}}$,
 $\text{top}(\mathcal{N}_l|_{\mathbf{x}_{-1}: b_1}) = \mathbf{a}_2 b_{i^*} \overrightarrow{a_{i^*+1}}$.
- $\mathcal{N}_j|_{\mathbf{x}_1} = V_j^{a_1}$, $\mathcal{N}_j|_{\mathbf{x}_{i^*}: a_1 \mathbf{a}_2} = V_j^{i^*}$,
 $\text{top}(\mathcal{N}_j) = a_1 \mathbf{a}_2 b_{i^*} \overrightarrow{a_{i^*+1}}$. Let \mathcal{N}_j be any tops-only-separable CP-net obtained by Condition I (where b_{i^*} corresponds to a_i in Condition I, and a_{i^*} corresponds to b_i in Condition I).
- $\mathcal{N}'_j|_{\mathbf{x}_1} = W_j^1$, \mathcal{N}'_j is tops-only-separable, and $\text{top}(\mathcal{N}'_j) = \text{top}(W_j^1) \mathbf{a}_2 b_{i^*} \overrightarrow{a_{i^*+1}}$.
- $\mathcal{N}'_j \in \text{CPnets}(\mathcal{S}_j)$. For any $l \leq n$, $\mathcal{N}_l \in \text{CPnets}(\mathcal{S}_l)$. All entries that are not defined above are chosen arbitrarily.

Because \mathcal{S} is rich, such CP-nets must exist. We let V_j be the extension of \mathcal{N}_j (which satisfies Condition I). That is, $V_j \sim \mathcal{N}_j$ and

$$b_1 \mathbf{a}_2 b_{i^*} \overrightarrow{a_{i^*+1}} \succ_{V_j} a_1 \mathbf{a}_2 a_{i^*} \overrightarrow{a_{i^*+1}}$$

Let $P = (V_1, \dots, V_{j-1}, V_j, V_{j+1}, \dots, V_n)$ be such that for all $l \leq n$, $V_l \in L_l$ and $V_l \sim \mathcal{N}_l$. Let $W_j \in L_j$, $W_j \sim \mathcal{N}'_j$. We next show that (P, W_j) is a successful manipulation

for voter j . We note that $P|_{\mathbf{x}_1} = P^1$, $\mathcal{M}|_{\mathbf{x}_1}(P^1) = a_1$; for any $i < i^*$, a_i is ranked in the top position in all votes of $P|_{\mathbf{x}_i:a_1a_2\cdots a_{i-1}}$; $P|_{\mathbf{x}_i^*:a_1a_2} = P^{i^*}$, $\mathcal{M}|_{\mathbf{x}_i^*:a_1a_2}(P^{i^*}) = a_{i^*}$; for any $i > i^*$, a_i is ranked in the top position in all votes of $P|_{\mathbf{x}_i:a_1a_2a_{i^*}a_{i^*+1}\cdots a_{i-1}}$. Therefore, $\mathcal{M}(P) = a_1a_2a_{i^*}\overrightarrow{a_{i^*+1}}$. On the other hand, $\mathcal{M}|_{\mathbf{x}_1}(P_{-j}^1, W_j^1) = b_1$; for any $i < i^*$, a_i is ranked in the top position in all votes of $P_{-j}|_{\mathbf{x}_i:b_1a_2\cdots a_{i-1}}$ and $W_j|_{\mathbf{x}_i:b_1a_2\cdots a_{i-1}}$; b_{i^*} is ranked at the top position in all votes of $P_{-j}|_{\mathbf{x}_i^*:b_1a_2}$ and $W_j|_{\mathbf{x}_i^*:b_1a_2}$; for any $i > i^*$, a_i is ranked in the top position in all votes of $P|_{\mathbf{x}_i:b_1a_2b_{i^*}a_{i^*+1}\cdots a_{i-1}}$ and $W_j|_{\mathbf{x}_i:b_1a_2b_{i^*}a_{i^*+1}\cdots a_{i-1}}$. Therefore,

$$\begin{aligned} \mathcal{M}(P_{-j}, W_j) &= b_1a_2b_{i^*}\overrightarrow{a_{i^*+1}} \\ &\succ_{V_j} a_1a_2a_{i^*}\overrightarrow{a_{i^*+1}} \\ &= \mathcal{M}(P) \end{aligned}$$

This contradicts the strategy-proofness of \mathcal{M} . (**End of proof of Lemma 3.**) \square

We prove the theorem by contradiction. Suppose there exists a non-dictatorial locally strategy-proof CR-net \mathcal{M} that satisfies non-imposition and is strategy-proof over L_{II} . For any $a_1 \in D_1$, we let $P^{a_1} = (V_1^{a_1}, \dots, V_n^{a_1})$ be a profile in $\prod_{j=1}^n \mathcal{S}_j|_{\mathbf{x}_1}$ such that each voter ranks a_1 in the top position. Because $\mathcal{M}|_{\mathbf{x}_1}$ is strategy-proof and satisfies non-imposition, $\mathcal{M}|_{\mathbf{x}_1}$ satisfies unanimity by Lemma 2, which means that $\mathcal{M}|_{\mathbf{x}_1}(P^{a_1}) = a_1$. For any $b_1 \neq a_1$, because $\mathcal{M}|_{\mathbf{x}_1}(P^{a_1}) \neq \mathcal{M}|_{\mathbf{x}_1}(P^{b_1})$, there exists a minimum $j \leq n$ such that

$$\begin{aligned} \mathcal{M}|_{\mathbf{x}_1}(V_1^{b_1}, \dots, V_{j-1}^{b_1}, V_j^{a_1}, V_{j+1}^{a_1}, \dots, V_n^{a_1}) &= a_1 \\ \mathcal{M}|_{\mathbf{x}_1}(V_1^{b_1}, \dots, V_{j-1}^{b_1}, V_j^{b_1}, V_{j+1}^{a_1}, \dots, V_n^{a_1}) &\neq a_1 \end{aligned}$$

That is, by replacing the $V_l^{a_1}$ by $V_l^{b_1}$ one after another for $l = 1, \dots, n$, before step $j - 1$, the winner of the profile is a_1 , and in step j the winner is not a_1 . By Lemma 3, voter j must be an a_1 -dictator.

Therefore, for any $a_1 \in D_1$, there exists $j \leq n$ such that for any $i \geq 2$, any $a_2 \in D_2 \times \cdots \times D_{i-1}$, $\mathcal{M}|_{\mathbf{x}_i:a_1a_2}$ is a j -dictatorship. We consider the following two cases.

Case 1: there exists $j \leq n$ such that for all $a_1 \in D_1$, voter j is an a_1 -dictator. Because \mathcal{M} is non-dictatorial, \mathcal{M} is not a j -dictatorship, which means that $\mathcal{M}|_{\mathbf{x}_1}$ is not a j -dictatorship. Therefore, there exists a profile P^1 in $\prod_{j=1}^n \mathcal{S}_j|_{\mathbf{x}_1}$ such that $\mathcal{M}|_{\mathbf{x}_1}(P^1) \neq \text{top}(V_j^1)$. Without loss of generality we let $j = 1$. We let $a_1 = \mathcal{M}|_{\mathbf{x}_1}(P^1)$, $b_1 = \text{top}(V_1^1)$. Because $\mathcal{M}|_{\mathbf{x}_1}$ is strategy-proof and satisfies non-imposition, $\mathcal{M}|_{\mathbf{x}_1}(V_1^1, V_2^{b_1}, \dots, V_n^{b_1}) = b_1$ (we recall that $\text{top}(V_1^1) = b_1$, and for all $2 \leq l \leq n$, $\text{top}(V_l^{b_1}) = b_1$). Therefore, there exists $2 \leq k \leq n$ such that

$$\begin{aligned} \mathcal{M}|_{\mathbf{x}_1}(V_1^1, V_2^{b_1}, \dots, V_{k-1}^{b_1}, V_k^1, V_{k+1}^1, \dots, V_n^1) &= a_1 \\ \mathcal{M}|_{\mathbf{x}_1}(V_1^1, V_2^{b_1}, \dots, V_{k-1}^{b_1}, V_k^{b_1}, V_{k+1}^1, \dots, V_n^1) &\neq a_1 \end{aligned}$$

Because voter 1 is an a_1 -dictator, voter k is not an a_1 -dictator. But this contradicts Lemma 3.

Case 2: there exists $j_1 \neq j_2$ and $a_1 \neq b_1$ such that voter j_1 (j_2) is an a_1 (b_1)-dictator. Without loss of generality, we let $j_1 = 1, j_2 = 2$. Let

$$P^1 = (V_1^{a_1}, V_2^{b_1}, V_3^{a_1}, \dots, V_n^{a_1})$$

$$Q^1 = (V_1^{a_1}, V_2^{b_1}, V_3^{b_1}, \dots, V_n^{b_1})$$

If $\mathcal{M}|_{x_1}(P^1) \neq a_1$, then, because $\mathcal{M}|_{x_1}(V_1^{a_1}, \dots, V_n^{a_1}) = a_1$, Lemma 3 implies that voter 2 is an a_1 -dictator, which is not possible because voter 1 is an a_1 -dictator. Therefore, $\mathcal{M}|_{x_1}(P^1) = a_1$. Similarly, $\mathcal{M}|_{x_1}(Q^1) = b_1$. Next, we consider the following steps: we change voter j 's vote from $V_j^{a_1}$ to $V_j^{b_1}$, one after another, for $3 \leq j \leq n$. It follows that there exists $3 \leq j \leq n$ such that

$$\mathcal{M}|_{x_1}(V_1^{a_1}, V_2^{b_1}, \dots, V_{j-1}^{b_1}, V_j^{a_1}, V_{j+1}^{a_1}, \dots, V_n^{a_1}) = a_1$$

$$\mathcal{M}|_{x_1}(V_1^{a_1}, V_2^{b_1}, \dots, V_{j-1}^{b_1}, V_j^{b_1}, V_{j+1}^{a_1}, \dots, V_n^{a_1}) \neq a_1$$

Lemma 3 implies that voter j is an a_1 -dictator, which is not possible because voter 1 is an a_1 -dictator.

Hence, we have obtained the desired contradiction, and can conclude that \mathcal{M} is dictatorial. **(End of proof of Theorem 4.)** \square

The following corollary is easily obtained from Theorem 4.

Corollary 2 *For any $1 \leq j \leq n$, suppose \mathcal{S}_j is a rich admissible conditional preference set, $LD(\mathcal{S}_j) \subseteq L_j \subseteq \text{Pref}(\mathcal{S}_j)$, and L_j satisfies Condition I. Then, a CR-net \mathcal{M} that satisfies non-imposition is strategy-proof over L_Π if and only if \mathcal{M} is a dictatorship.*

Proof of Corollary 2: Let \mathcal{M} be a strategy-proof CR-net over L_Π . Because $LD(\mathcal{S}_j) \subseteq L_j$ for every $1 \leq j \leq n$, \mathcal{M} is strategy-proof over $\prod_{j=1}^n LD(\mathcal{S}_j)$, which implies that \mathcal{M} is locally strategy-proof by Theorem 1. We note that $LD(\mathcal{S}_j)$ extends \mathcal{S}_j for all j , which means that L_j extends \mathcal{S}_j for every $1 \leq j \leq n$. Hence, by Theorem 4, \mathcal{M} is dictatorial. \square

The next theorem states that over any superset of the lexicographic preference domain, the only strategy-proof voting rule that satisfies non-imposition is a locally strategy-proof CR-net. We note that this result does not directly follows from Theorem 1, because from Theorem 1 we only know that this rule must be a CR-net when all votes are lexicographic, which does not mean that it is still a CR-net beyond the lexicographic preference domain.

Theorem 5 *For any $1 \leq j \leq n$, suppose \mathcal{S}_j is a rich admissible conditional preference set, and $LD(\mathcal{S}_j) \subseteq L_j \subseteq \text{Pref}(\mathcal{S}_j)$. If a voting rule r that satisfies non-imposition is strategy-proof over L_Π , then r is a locally strategy-proof CR-net.*

Proof of Theorem 5: Because r is strategy-proof over L_Π , the restriction of r to $\prod_{j=1}^n LD(\mathcal{S}_j)$, denoted by $r_{LD(\mathcal{S}_\Pi)}$, is strategy-proof over $\prod_{j=1}^n LD(\mathcal{S}_j)$. It follows from Theorem 1 that $r_{LD(\mathcal{S}_\Pi)}$ is a locally strategy-proof CR-net, denoted by \mathcal{M} . Because for any $1 \leq j \leq n$, $LD(\mathcal{S}_j)$ extends \mathcal{S}_j , \mathcal{M} can be naturally extended to L_Π . All that remains to show is that r and \mathcal{M} are the same rule.

Lemma 4 For any profile $P \in L_\Pi$, if at most one of the votes in P is not lexicographic, then $r(P) = \mathcal{M}(P)$.

Proof of Lemma 4: Suppose that the lemma does not hold. Then, there exists $P = (V_1, \dots, V_n) \in L_\Pi$ such that $r(P) \neq \mathcal{M}(P)$, (without loss of generality) $V_1 \notin \text{Lex}(\mathcal{S}_1)$, and, for any $j \geq 2$, V_j is lexicographic. Let i^* be the index of the first component of $r(P)$ that is different from the same component of $\mathcal{M}(P)$. That is, the value of issue \mathbf{x}_{i^*} in $r(P)$ (denoted by a_{i^*}) is different from the value of issue \mathbf{x}_{i^*} in $\mathcal{M}(P^*)$ (denoted by b_{i^*}); and for any $l < i^*$, the value of issue \mathbf{x}_l in $r(P)$ is the same as the value of issue \mathbf{x}_l in $\mathcal{M}(P)$. Let $\mathbf{a} = (a_1, \dots, a_p) = r(P)$. For any $1 \leq j \leq n$, we define a CP-net \mathcal{N}'_j as follows.

- $\mathcal{N}'_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}} = V_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}$.
- \mathcal{N}'_j is tops-only-separable, and $\text{top}(\mathcal{N}'_j) = (a_1, \dots, a_{i^*-1}, \text{top}(V_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}), a_{i^*+1}, \dots, a_p)$.

For any $1 \leq j \leq n$, let V'_j be the lexicographic extension of \mathcal{N}'_j . Because V'_j is lexicographic, for any $j \geq 2$, any $\mathbf{d} \in \mathcal{X}$, if $\mathbf{d} \succ_{V'_j} \mathbf{a}$, then, $d_{i^*} \succ_{V'_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}} a_{i^*}$. We note that $V'_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}} = V_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}$, which means that $d_{i^*} \succ_{V_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}} a_{i^*}$. Therefore, $\mathbf{d} \succ_{V_j} \mathbf{a}$. It follows from Lemma 1 that $r(V_1, V'_2, \dots, V'_n) = \mathbf{a}$. We note that $r(V'_1, V'_2, \dots, V'_n) = \mathcal{M}(V'_1, V'_2, \dots, V'_n) = (\mathbf{a}_{-i^*}, b_{i^*})$, where $b_{i^*} \neq a_{i^*}$, because this is a lexicographic profile. If $b_{i^*} \succ_{V_1|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}} a_{i^*}$, then, $(\mathbf{a}_{-i^*}, b_{i^*}) \succ_{V_1} \mathbf{a}$, which means that $((V_1, V'_2, \dots, V'_n), V'_1)$ is a successful manipulation for voter 1; on the other hand, if $a_{i^*} \succ_{V_1|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}} b_{i^*}$, then, because $V'_1|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}} = V_1|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}$, we have $\mathbf{a} \succ_{V'_1} (\mathbf{a}_{-i^*}, b_{i^*})$, which means that $((V'_1, V'_2, \dots, V'_n), V_1)$ is a successful manipulation for voter 1. This contradicts the strategy-proofness of r . **(End of proof of Lemma 4.)** \square

Next, we prove the more general proposition that for any $P \in L_\Pi$, $r(P) = \mathcal{M}(P)$, which will complete the proof of the theorem. Suppose that the claim does not hold. Then, we let \mathcal{P} be the set of profiles in L_Π whose winner under r is different from the winner under \mathcal{M} , that is, $\mathcal{P} = \{P \in L_\Pi : r(P) \neq \mathcal{M}(P)\}$. We have $\mathcal{P} \neq \emptyset$. Let $P^* \in \mathcal{P}$ denote a profile in which the number of non-lexicographic votes is minimized (equivalently, the number of lexicographic voters is maximized). That is, for any $P \in \mathcal{P}$, the number of non-lexicographic votes in P is at least the number of non-lexicographic votes in P^* . Let l be the number of non-lexicographic votes in P^* (by Lemma 4, $l \geq 2$). It follows that for any $P \in L_\Pi$, if the number of non-lexicographic votes in P is at most $l - 1$, then $r(P) = \mathcal{M}(P)$.

Without loss of generality, we let $P^* = (V_1, \dots, V_n)$, where V_1, \dots, V_l are non-lexicographic, and V_{l+1}, \dots, V_n are lexicographic. For any $1 \leq j \leq n$, we let $\mathcal{N}_j \in \text{CPnets}(\mathcal{S}_j)$ be the CP-net that V_j extends. Let $\mathcal{M}(P) = \mathbf{a}$, $r(P) = \mathbf{b}$. By the minimality of l , $r(\text{Lex}(\mathcal{N}_1), V_2, \dots, V_n) = \mathcal{M}(\text{Lex}(\mathcal{N}_1), V_2, \dots, V_n) = \mathbf{a}$, because the number of non-lexicographic votes in the modified profile is $l - 1$. Because r is strategy-proof, we must have that $\mathbf{b} \succ_{V_1} \mathbf{a}$: otherwise, $(P^*, \text{Lex}(\mathcal{N}_1))$ is a successful manipulation for voter 1.

Let \mathcal{N}_1^* be a CP-net in which \mathbf{b} is ranked at the top. It follows from Lemma 1 and the strategy-proofness of r that

$r(\text{Lex}(\mathcal{N}_1^*), V_2, \dots, V_n) = \mathbf{b}$. Then, because the number of non-lexicographic votes in $(\text{Lex}(\mathcal{N}_1^*), V_2, \dots, V_n)$ is $l - 1$, we have the following equations.

$$\begin{aligned} \mathbf{b} &= r(\text{Lex}(\mathcal{N}_1^*), V_2, \dots, V_n) \\ &= \mathcal{M}(\text{Lex}(\mathcal{N}_1^*), V_2, \dots, V_n) \\ &= \mathcal{M}(\text{Lex}(\mathcal{N}_1^*), \text{Lex}(\mathcal{N}_2), \dots, \text{Lex}(\mathcal{N}_n)) \end{aligned}$$

The second equation holds because the number of non-lexicographic votes in $(\text{Lex}(\mathcal{N}_1^*), V_2, \dots, V_n)$ is $l - 1$. By Lemma 4, we have the following equations.

$$\begin{aligned} &r(V_1, \text{Lex}(\mathcal{N}_2), \dots, \text{Lex}(\mathcal{N}_n)) \\ &= \mathcal{M}(V_1, \text{Lex}(\mathcal{N}_2), \dots, \text{Lex}(\mathcal{N}_n)) \\ &= \mathcal{M}(V_1, V_2, \dots, V_n) = \mathbf{a} \end{aligned}$$

We recall that $\mathbf{b} \succ_{V_1} \mathbf{a}$, which means that

$((V_1, \text{Lex}(\mathcal{N}_2), \dots, \text{Lex}(\mathcal{N}_n)), \text{Lex}(\mathcal{N}_1^*))$ is a successful manipulation for voter 1. This contradicts the strategy-proofness of r . Therefore, $r = \mathcal{M}$. **(End of proof of Theorem 5.)** \square

Combining Corollary 2 and Theorem 5, we obtain the following impossibility theorem on supersets of any lexicographic preference domain.

Theorem 6 *For any $1 \leq j \leq n$, suppose \mathcal{S}_j is a rich conditional preference set, $LD(\mathcal{S}_j) \subseteq L_j \subseteq \text{Pref}(\mathcal{S}_j)$, and L_j satisfies Condition I. Then, the only strategy-proof voting rule over L_Π that satisfies non-imposition is a dictatorship.*

We recall that if L_j satisfies Condition I, which informally means that any issue i is more important than issue 1 in at least one admissible vote. Theorem 3 follows from Theorem 6 by letting $\mathcal{S}_j|_{\mathbf{x}_i: \mathbf{d}_i} = L(D_i)$ and $L_j = \text{Pref}(\mathcal{S}_j)$ (the same corollary also follows from Theorem 8 in the next section).

C Impossibility result for extensions of rich preference domains

Le Breton and Sen [11] characterized strategy-proof voting rules when preferences are separable, that is, each vote extends a CP-net with no edges. An admissible conditional preference set \mathcal{S} is *separable* if for any \mathbf{x}_i , any $\mathbf{a}_i, \mathbf{b}_i \in D_1 \times \dots \times D_{i-1}$, we have $\mathcal{S}|_{\mathbf{x}_i: \mathbf{a}_i} = \mathcal{S}|_{\mathbf{x}_i: \mathbf{b}_i}$. In this case, we write $\mathcal{S}|_{\mathbf{x}_i} = \mathcal{S}|_{\mathbf{x}_i: \mathbf{a}_i}$. For example, Example 3 has a separable admissible conditional preference set (because the allowed preferences for wine do not depend on the choice of the main course). For any separable admissible conditional preference set \mathcal{S} , we let $\text{SCPnets}(\mathcal{S}) = \{\mathcal{N} : \mathcal{N} \text{ is a CP-net with no edge, and for any } i \leq p, \mathcal{N}|_{\mathbf{x}_i} \in \mathcal{S}|_{\mathbf{x}_i}\}$.

That is, $\text{SCPnets}(\mathcal{S})$ is the set of all CP-nets \mathcal{N} with no edges, such that the projection of \mathcal{N} to any issue \mathbf{x}_i is in $\mathcal{S}|_{\mathbf{x}_i}$. Let $\text{SPref}(\mathcal{S})$ denote the set of all separable votes that extend some CP-net in $\text{SCPnets}(\mathcal{S})$. We now present the richness definition by Le Breton and Sen (in our notation).

Definition 7 (Le Breton and Sen [11]) $R_\Pi = \prod_{j=1}^n R_j$ is a rich preference domain, if for any $1 \leq j \leq n$, there exists a separable admissible conditional preference set \mathcal{S}_j such that $R_j \subseteq \text{SPref}(\mathcal{S}_j)$ and

(A) for any $1 \leq j \leq n$, any $1 \leq i \leq p$, any $a_i \in D_i$, there exists $V^i \in \mathcal{S}_j|_{\mathbf{x}_i}$ such that $\text{top}(V^i) = a_i$.

(B) for any $1 \leq j \leq n$, any $\mathcal{N}_j \in \text{SPref}(\mathcal{S}_j)$, and any $1 \leq i \leq p$, there exist $V_j, V'_j \in R_j, V_j \sim \mathcal{N}_j, V'_j \sim \mathcal{N}_j$ such that

(i) for any $\mathbf{a}, \mathbf{b} \in \mathcal{X}$, if $a_i \succ_{\mathcal{N}_j|_{\mathbf{x}_i}} b_i$, then $\mathbf{a} \succ_{V_j} \mathbf{b}$. That is, issue i dominates all other issues for V_j .

(ii) for any $\mathbf{a}, \mathbf{b} \in \mathcal{X}$, if for all $i' \neq i$, $a_{i'} \succeq_{\mathcal{N}_j|_{\mathbf{x}_{i'}}} b_{i'}$ and there exists $i' \neq i$ such that $a_{i'} \succ_{\mathcal{N}_j|_{\mathbf{x}_{i'}}} b_{i'}$ (that is, \mathbf{a}_{-i} weakly dominates \mathbf{b}_{-i}), then, $\mathbf{a} \succ_{V'_j} \mathbf{b}$. That is, issue i is dominated by the (union of) other issues for V'_j .

R_Π satisfies condition (A) if and only if \mathcal{S} is rich (according to our earlier definition of richness). We note that Condition I (in Definition 6) is weaker than condition B(i) in the following sense: if $R_j \subseteq \text{SPref}(\mathcal{S}_j)$ satisfies condition B(i), then, it also satisfies Condition I, because the vote guaranteed to exist by condition B(i) satisfies all the premises of Condition I.

The following is the main theorem by Le Breton and Sen (in our notation).

Theorem 7 (Le Breton and Sen [11]) Let $R_\Pi = \prod_{j=1}^n R_j$ be a rich preference domain. A voting rule r that satisfies non-imposition is strategy-proof over R_Π if and only if it is a separable locally strategy-proof CR-net.

Theorem 7 works (only) for any rich preference domain $R_\Pi \subseteq \prod_{j=1}^n \text{SPref}(\mathcal{S}_j)$, where \mathcal{S}_j is the separable admissible conditional preference set that R_j corresponds to. We note that for any $1 \leq j \leq n$, $\text{SPref}(\mathcal{S}_j)$ is a strict subset of $\text{Pref}(\mathcal{S}_j)$, and $\text{SPref}(\mathcal{S}_j)$ is exponentially smaller than $\text{Pref}(\mathcal{S}_j)$. Next, we consider the case that for any $1 \leq j \leq n$, the preference domain of voter j , denoted by L_j , is both a superset of R_j , and a subset of $\text{Pref}(\mathcal{S}_j)$. We first obtain a corollary from Theorem 7.

Corollary 3 Let R_Π be a rich preference domain. For any $1 \leq j \leq n$, suppose $R_j \subseteq L_j \subseteq \text{Pref}(\mathcal{S}_j)$ and L_j extends \mathcal{S}_j . If a sequential voting rule \mathcal{M} that satisfies non-imposition is strategy-proof over L_Π , then, \mathcal{M} is a dictatorship.

Proof of Corollary 3: For any $1 \leq j \leq n$, any $\mathbf{a} = (a_1, \dots, a_p) \in \mathcal{X}$, any $V_j^{a_1} \in \mathcal{S}_j|_{\mathbf{x}_1}$ such that $\text{top}(V_j^{a_1}) = a_1$, any $V_j^{a_i} \in \mathcal{S}_j|_{\mathbf{x}_i}$ such that $\text{top}(V_j^{a_i}) = a_i$, we let $\mathcal{N}_j \in \text{SCPnets}(\mathcal{S}_j)$ be such that $\mathcal{N}_j|_{\mathbf{x}_1} = V_j^{a_1}, \mathcal{N}_j|_{\mathbf{x}_i} = V_j^{a_i}$, and $\text{top}(\mathcal{N}_j) = \mathbf{a}$; let V_j be an extension of \mathcal{N}_j satisfying the condition B(i) for issue i in Definition 7. We note that for any $b_1 \in D_1, b_1 \neq a_1$, any $b_i \in D_i, b_i \neq a_i$, $(b_1, \mathbf{a}_{-1}) \succ_{V_j} (b_i, \mathbf{a}_{-i})$, because $a_i \succ_{V_j|_{\mathbf{x}_i}} b_i$. Because $R_j \subseteq L_j$, we have $V_j \in L_j$, which means that L_j satisfies Condition I.

By Theorem 7, \mathcal{M} is locally strategy-proof over $\prod_{j=1}^n \mathcal{R}_j$. Because $L_\Pi \subseteq \text{Pref}(\mathcal{S})$, \mathcal{M} is locally strategy-proof over L_Π . Therefore, by Theorem 4, \mathcal{M} is dictatorial. \square

Our next theorem states that if for any $1 \leq j \leq n$, L_j is a superset of R_j , then the only strategy-proof voter rule over L_Π is the sequential composition of locally strategy-proof rules, one for each issue.

Theorem 8 *Let R_Π be a rich preference domain. For any $1 \leq j \leq n$, let $R_j \subseteq L_j \subseteq \text{Pref}(\mathcal{S}_j)$. If voting rule r that satisfies non-imposition is strategy-proof over L_Π , then r is a locally strategy-proof sequential voting rule (separable CR-net).*

Proof of Theorem 8: Because r is strategy-proof over R_Π , by Theorem 7, there exists a separable CR-net \mathcal{M} such that for any $P \in R_\Pi$, $r(P) = \mathcal{M}(P)$. We note that the domain of \mathcal{M} can be extended to $\prod_{j=1}^n \text{Pref}(\mathcal{S}_j)$ in a natural way, as follows. For any $P \in \prod_{j=1}^n \text{Pref}(\mathcal{S}_j)$, let $\mathcal{M}(P) = (d_1, \dots, d_p)$ in which $d_i = \mathcal{M}|_{\mathbf{x}_i}(P|_{\mathbf{x}_i:d_1 \dots d_{i-1}})$. In this case, \mathcal{M} is equivalent to the sequential voting rule $\text{Seq}(\mathcal{M}|_{\mathbf{x}_1}, \dots, \mathcal{M}|_{\mathbf{x}_p})$. We next show that for any $P \in \prod_{j=1}^n \text{Pref}(\mathcal{S}_j)$, $r(P) = \mathcal{M}(P)$. Suppose for the sake of contradiction that there exists $P \in \prod_{j=1}^n \text{Pref}(\mathcal{S}_j)$ such that $r(P) \neq \mathcal{M}(P)$. Let $\mathbf{a} = r(P)$, $\mathbf{b} = \mathcal{M}(P)$, and let i^* be the smallest number that satisfies $a_{i^*} \neq b_{i^*}$. Let $\mathcal{N}_1, \dots, \mathcal{N}_n$ be a set of CP-nets with no edges such that for any $1 \leq i \leq p$, $i \neq i^*$, $\text{top}(\mathcal{N}_i|_{\mathbf{x}_i}) = a_i$, and $\mathcal{N}_j|_{\mathbf{x}_{i^*}} = V_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}$. Let $P' = (V'_1, \dots, V'_p)$ be the profile in which for all $1 \leq j \leq n$, V'_j is the extension of \mathcal{N}_j that satisfies condition B(ii) from Definition 7 w.r.t. i^* . That is, for any $1 \leq j \leq n$, any $\mathbf{y}, \mathbf{z} \in \mathcal{X}$, if \mathbf{y}_{-i^*} weakly dominates \mathbf{z}_{-i^*} in \mathcal{N}_j , then $\mathbf{y} \succ_{V'_j} \mathbf{z}$. For any $\mathbf{d} \in \mathcal{X}$, any $1 \leq j \leq n$, $\mathbf{d} \succ_{V'_j} \mathbf{a}$ if and only if for any $i \neq i^*$, $d_i = a_i$, and $d_{i^*} \succ_{V'_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}} a_{i^*}$. We note that $V'_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}} = V_j|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}$. It follows that $\mathbf{d} \succ_{V'_j} \mathbf{a}$ implies $\mathbf{d} \succ_{V_j} \mathbf{a}$. Therefore, by Lemma 1, $r(P') = \mathbf{a}$. Since $P' \in R_\Pi$, $\mathcal{M}(P') = r(P') = \mathbf{a}$. We note that $P'|_{\mathbf{x}_{i^*}} = P|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}$, which means that

$$\begin{aligned} a_{i^*} &= \mathcal{M}(P')|_{\mathbf{x}_{i^*}} = \mathcal{M}|_{\mathbf{x}_i}(P'|_{\mathbf{x}_{i^*}}) \\ &= \mathcal{M}|_{\mathbf{x}_i}(P|_{\mathbf{x}_{i^*}:a_1 \dots a_{i^*-1}}) \\ &= b_{i^*} \end{aligned}$$

This contradicts the assumption that $a_{i^*} \neq b_{i^*}$. \square

Finally, by combining Theorem 8 and Corollary 3, we obtain the following impossibility result. This theorem states that if take a rich preference domain that corresponds to a separable admissible conditional preference set, and extend it so that for any acyclic CP-net that uses the same admissible conditional preference set, we include some preferences extending that CP-net, then we must give up one of strategy-proofness, non-dictatorship, and non-imposition.

Theorem 9 *Let R_Π be a rich preference domain. For any $1 \leq j \leq n$, suppose that $R_j \subseteq L_j \subseteq \text{Pref}(\mathcal{S}_j)$ and L_j extends \mathcal{S}_j . A voting rule that satisfies non-imposition is strategy-proof over L_Π if and only if it is a dictatorship.*

We note that Theorem 3 also follows from Theorem 9.