

Strategy-Proof Voting Rules over Restricted Domains

We have seen in Chapter 7 and Chapter 11 that in some voting games, strategic behavior sometimes leads to extremely undesirable outcomes. Therefore, we may want to prevent voters' strategic behavior. However, the Gibbard-Satterthwaite theorem tells us that for any voting rule that satisfies some natural properties, there must exist at least one voter who has an incentive to misreport her preferences, if the voters are allowed to use any linear order to represent their preferences. To circumvent the Gibbard-Satterthwaite theorem, researchers in Computational Social Choice have investigated the possibility of using computational complexity to prevent voters' strategic behavior. Chapter 4 showed that for some common voting rules computational complexity can provide some protection from manipulation, while Chapter 5 gave some evidence that computational complexity does not seem to be a very strong barrier against manipulation.

In fact, there is another, older, line of research on circumventing the Gibbard-Satterthwaite theorem. 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 (Black, 1948). For single-peaked preferences, desirable strategy-proof rules exist, such as the *median* rule. Other strategy-proof rules are also possible in this preference domain: for example, it is possible to add some artificial (*phantom*) votes before running the median rule. In fact, this characterizes all strategy-proof rules for single-peaked preferences (Moulin, 1980). On the other hand, preferences have to be significantly restricted to obtain such positive results: Aswal et al. (2003) 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 this chapter we will focus on exploring the possibility of using domain restriction to circumvent the Gibbard-Satterthwaite theorem in multi-issue domains. The problem of characterizing strategy-proof voting rules in multi-issue domains has already received significant attention. Strategy-proof voting rules for high-dimensional single-peaked preferences (where each dimension can be seen as an issue) have been characterized (Border and Jordan, 1983; Barbera et al., 1993, 1997; Nehring and Puppe, 2007). Barbera et al. (1991) 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 (2003) studied multi-issue domains where each issue can take three values: "good", "bad", and "null", and characterized all strategy-proof voting rules that satisfy *null-independence*, that is, if a voter votes "null" on an issue i , then her preferences over other issues do not affect the value of issue i .

The prior research that is closest to ours was performed by Le Breton and Sen (1999). They proved that if the voters' preferences are separable, and the restricted preference domain of the voters satisfies a *richness* condition, then, a voting rule is strategy-proof if and only if it is an issue-by-issue voting rule, in which each issue-wise voting rule is strategy-proof over its respective domain.

Despite its elegance, the work by Le Breton and Sen is limited by the restrictiveness of separable preferences: as we have argued in Chapter 8, 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. Therefore, it seems that sequential voting (Section 8.3) is better than issue-by-issue voting. While the assumption of sequential voting that there exists an ordering \mathcal{O} over issues such that all voters’ preferences are \mathcal{O} -legal is still restrictive, it is much less restrictive than assuming that preferences are separable. Chapter 9 and Chapter 10 concerned how to design new voting rules when voters use a much more expressive voting language (i.e., possibly cyclic CP-nets), but in this chapter, we only study the setting where all voters’ preferences are \mathcal{O} -legal, and w.l.o.g. we fix $\mathcal{O} = X_1 > X_2 > \dots > X_p$.

The main theorem of this chapter is the following: 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 chapter 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.

12.1 Conditional Rule Nets (CR-Nets)

In this section, we give the motivation and formal definition of CR-nets. In a sequential voting rule, the local voting rule that is used for a given issue is always the

same, that is, the local voting *rule* does not depend on the decisions made on earlier issues (though, of course, the voters’ *preferences* for this issue do depend on those decisions).

However, in many cases, it makes sense to let the local voting rules depend on the values of preceding issues. For example, let us consider again the setting in Example 8.0.2, where a group of people must make a joint decision on the menu for dinner, and the menu is composed of two issues: the main course (**M**) and the wine (**W**). 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., $r > p > w$). Still, it can happen that even conditioned on beef being selected, surprisingly, slightly more than half the voters vote for white wine ($w > p > r$), and slightly less than half vote for red ($r > p > w$). In this case, the caterer, who knows that in the general population most people prefer red to white given a meal of beef, may “overrule” the preference for white wine among the slight majority of the voters, and select red wine anyway. While this may appear somewhat snobbish on the part of the caterer, in fact she may be acting in the best interest of social welfare if we take the non-voting agents (who are likely to prefer red given beef) into account.

To model voting rules where the local rules depend on the values chosen for earlier issues, we introduce *conditional rule nets* (*CR-nets*). A CR-net is defined similarly to a CP-net—the difference is that CPTs are replaced by conditional rule tables (CRTs), which specify a local voting rule over D_i for each issue X_i and each setting of the parents of X_i .¹

Definition 12.1.1. *An (acyclic) conditional rule net (CR-net) \mathcal{M} over \mathcal{X} is composed of the following two parts.*

¹ It is not clear how a cyclic CR-net could be useful, so we only define acyclic CR-nets.

1. A directed acyclic graph G over $\{X_1, \dots, X_p\}$.
2. A set of conditional rule tables (CRTs) in which, for any variable X_i and any setting \vec{d} of $\text{Par}_G(X_i)$, there is a local conditional voting rule $\mathcal{M}|_{X:\vec{d}}$ over D_i .

A CR-net encodes a voting rule over all \mathcal{O} -legal profiles (we recall that we fix $\mathcal{O} = X_1 > \dots > X_p$ in this chapter). For any $1 \leq i \leq p$, in the i th step, the value d_i is determined by applying $\mathcal{M}|_{X_i:d_1 \dots d_{i-1}}$ (the local rule specified by the CR-net for the i th issue given that the earlier issues take the values $d_1 \dots d_{i-1}$) to $P|_{X_i:d_1 \dots d_{i-1}}$ (the profile of preferences over the i th issue, given that the earlier issues take the values $d_1 \dots d_{i-1}$). Formally, for any \mathcal{O} -legal profile P , $\mathcal{M}(P) = (d_1, \dots, d_p)$ is defined as follows: $d_1 = \mathcal{M}|_{X_1}(P|_{X_1})$, $d_2 = \mathcal{M}|_{X_2:d_1}(P|_{X_2:d_1})$, etc. Finally, $d_p = \mathcal{M}|_{X_p:d_1 \dots d_{p-1}}(P|_{X_p:d_1 \dots d_{p-1}})$.

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

12.2 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 Le Breton and Sen (1999), 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 $r > w > p$ 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 \vec{d}_i of the issues before issue X_i , there

is a set of “reasonable” (or: possible, admissible) preferences over X_i , which we call $\mathcal{L}|_{X_i:\vec{d}_i}$. Formally, *admissible conditional preference sets*, which encode all possible conditional preferences of voters, are defined as follows.

Definition 12.2.1. *An admissible conditional preference set \mathcal{L} over \mathcal{X} is composed of multiple local conditional preference sets, denoted by $\mathcal{L}|_{X_i:\vec{d}_i}$, such that for any $1 \leq i \leq p$ and any $\vec{d}_i \in D_1 \times \dots \times D_{i-1}$, $\mathcal{L}|_{X_i:\vec{d}_i}$ is a set of (not necessarily all) linear orders over D_i .*

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

An admissible conditional preference set restricts the possible CP-nets, preferences, and lexicographic preferences. We note that Le Breton and Sen (1999) 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{L} 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 $\vec{d}_i \in D_1 \times \dots \times D_{i-1}$, the restriction of \mathcal{N} on X_i given \vec{d}_i is chosen from $\mathcal{L}|_{X_i:\vec{d}_i}$; second, an extension of \mathcal{N} is chosen as the voter’s vote. By restricting the freedom in either of the two steps (or both), we obtain a set of restricted preferences for the voter. Hence, we have the following definitions.

Definition 12.2.2. *Let \mathcal{L} be an admissible conditional preference set over \mathcal{X} .*

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

That is, $\text{CPnets}(\mathcal{L})$ is the set of all CP-nets over \mathcal{X} corresponding to preferences that are consistent with the admissible conditional preference set \mathcal{L} . $\text{Pref}(\mathcal{L})$ is the set of all linear orders that are consistent with the admissible conditional preference set \mathcal{L} . $LD(\mathcal{L})$, which we call the *lexicographic preference domain*, is the subset of linear orders in $\text{Pref}(\mathcal{L})$ that are lexicographic. For any $L \subseteq \text{Pref}(\mathcal{L})$, we say that L *extends* \mathcal{L} if for any CP-net in $\text{CPnets}(\mathcal{L})$, there exists at least one linear order in L consistent with that CP-net. It follows that $LD(\mathcal{L})$ extends \mathcal{L} ; in this case, for any CP-net \mathcal{N} in $\text{CPnets}(\mathcal{L})$, there exists exactly one linear order in $LD(\mathcal{L})$ 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.²

Definition 12.2.3. *An admissible conditional preference set \mathcal{L} is rich if for each $1 \leq i \leq p$, each valuation \vec{d}_i of the preceding issues, and each $a_i \in D_i$, there exists $V^i \in \mathcal{L}|_{X_i:\vec{d}_i}$ 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 X_i given \vec{d}_i , she should have the freedom to at least specify her most preferred local alternative in D_i at will. We note that $|\mathcal{L}|_{X_i:\vec{d}_i}|$ can be as small as $|D_i|$ (by letting each alternative in D_i be ranked in the top position exactly once), which is in sharp contrast to $|L(D_i)| = |D_i|!$ (when all local orders are allowed).

A CR-net \mathcal{M} is *locally strategy-proof* if all its local conditional rules are strategy-

² 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.

proof over their respective local domains (we recall that the voters' local preferences must be in the corresponding local conditional preference set). That is, for any $1 \leq i \leq p$, $\vec{d}_i \in D_1 \times \cdots \times D_{i-1}$, $\mathcal{M}|_{X_i:\vec{d}_i}$ is strategy-proof over $\prod_{j=1}^n \mathcal{L}_j|_{X_i:\vec{d}_i}$.

12.3 Strategy-Proof Voting Rules in Lexicographic Preference Domains

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

Theorem 12.3.1. *For any $1 \leq j \leq n$, suppose \mathcal{L}_j is a rich admissible conditional preference set, and voter j 's preferences are restricted to the lexicographic preference domain of \mathcal{L}_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.*

Proof of Theorem 12.3.1: In this proof, for any $1 \leq i \leq p$, we let X_{-i} denote $\mathcal{I} \setminus \{X_i\}$. 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 .

Before formally proving the theorem, let us first sketch the idea behind the proof. The "if" part is easy. The "only if" part is proved by induction on p (the number of issues). More precisely, suppose the theorem holds for p issues. For $p + 1$ issues, let r be a strategy-proof voting rule that satisfies non-imposition. We first prove that r

can be decomposed in the following way: there exists a local rule r_1 over D_1 and a voting rule $r_{X_{-1}:a_1}$ over $D_2 \times \cdots \times D_{p+1}$ for each $a_1 \in D_1$ that satisfy the following two conditions.

1. For any profile P , the first component of $r(P)$ is determined by applying r_1 to the projection of P on X_1 , and
2. the remaining components are determined by applying $r_{X_{-1}:a_1}$ to the restriction of P on the remaining issues given $X_1 = a_1$, where a_1 is the first component of $r(P)$, which is just determined by r_1 .

Moreover, we prove that r_1 and $r_{X_{-1}:a_1}$ (for all $a_1 \in D_1$) satisfy non-imposition and strategy-proofness. Therefore, by the induction hypothesis, for each $a_1 \in D_1$, $r_{X_{-1}:a_1}$ is a locally strategy-proof CR-net over $D_2 \times \cdots \times D_{p+1}$. It follows that r is a locally strategy-proof CR-net over $D_1 \times \cdots \times D_{p+1}$, in which the (unconditional) rule for X_1 is r_1 , and given any $a_1 \in D_1$, the sub-CR-net conditioned on $X_1 = a_1$ is $r_{X_{-1}:a_1}$.

We now formally prove the theorem. We will use Lemma 12.3.2, which 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 12.3.2 (Known). *Any strategy-proof voting rule satisfies monotonicity.*

Proof of Lemma 12.3.2: 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.

□

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 $\mathcal{X}_{k+1} = D_1 \times \cdots \times D_{k+1}$, we prove that this rule can be decomposed into two parts: first, it applies a local voting rule r_1 for X_1 , and subsequently, it applies a rule $r|_{X_{-1}:a_1}$ for X_{-1} , which depends on the outcome of r_1 . Thus, we have the property that for any $P \in \prod_{j=1}^n LD(\mathcal{L}_j)$, we have $r(P) = (r_1(P|_{X_1}), r|_{X_{-1}:r_1(P|_{X_1})}(P|_{X_{-1}:r_1(P|_{X_1})}))$. Then, we will show that the induction assumption can be applied to the second part.

To prove these, we claim that for any strategy-proof voting rule r satisfying non-imposition, and any $P \in \prod_{j=1}^n LD(\mathcal{L}_j)$, the value of issue X_1 for the winning alternative only depends on the restriction of the profile to X_1 . That is, we show that for any pair of profiles $P, Q \in \prod_{j=1}^n LD(\mathcal{L}_j)$, where $P = (V_1, \dots, V_n)$, $Q = (W_1, \dots, W_n)$ and $P|_{X_1} = Q|_{X_1}$, we must have $r(P)|_{X_1} = r(Q)|_{X_1}$. Suppose on the contrary that $r(P)|_{X_1} \neq r(Q)|_{X_1}$. 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 $0 \leq j \leq n - 1$, $r(P_j)|_{X_1} = r(P_{j+1})|_{X_1}$. For the sake of contradiction, suppose $r(P_j)|_{X_1} \neq r(P_{j+1})|_{X_1}$ for some $j \leq n - 1$. Let $a_1 = r(P_j)|_{X_1}$ and $b_1 = r(P_{j+1})|_{X_1}$. If $a_1 \succ_{V_{j+1}|_{X_1}} b_1$, then, because $V_{j+1}|_{X_1} = W_{j+1}|_{X_1}$, (P_{j+1}, V_{j+1}) is a successful manipulation; on the other hand, if $b_1 \succ_{V_{j+1}|_{X_1}} a_1$, 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 X_1 for the winning alternative only depends on the restriction of the profile to X_1 .

Therefore, we can define a voting rule r_1 over D_1 as follows. For any $P^1 \in \prod_{j=1}^n \mathcal{L}_j|_{X_1}$, $r_1(P^1) = r(P)|_{X_1}$, where $P \in \prod_{j=1}^n LD(\mathcal{L}_j)$ and $P|_{X_1} = P^1$. Such a P exists because $LD(\mathcal{L}_j)$ extends \mathcal{L}_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 $1 \leq j \leq n$, $\mathcal{N}_j|_{X_1} = V_j^1$; $\hat{\mathcal{N}}_l|_{X_1} = \hat{V}_l^1$.
- For any $1 \leq j \leq n$, $\mathcal{N}_j \in \text{CPnets}(\mathcal{L}_j)$, $\hat{\mathcal{N}}_l \in \text{CPnets}(\mathcal{L}_l)$.

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 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 X_1 component of $r(P)$. Because V_l is the lexicographic extension of \mathcal{N}_l , and $\mathcal{N}_l|_{X_1} = V_l^1$, 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 $r|_{X_{-1}:r_1(P|_{X_1})}(P|_{X_{-1}:r_1(P|_{X_1})})$. That is, the rule for the remaining issues X_{-1} only depends on the outcome for X_1 . For any $V \in \text{Legal}(\mathcal{O})$ and any $a_1 \in D_1$, we let $V|_{X_{-1}:a_1}$ denote the linear preference over D_{-1} that is compatible with the restriction of V to the set of alternatives whose X_1 component is a_1 , that is, for any $\vec{a}_{-1}, \vec{b}_{-1} \in D_{-1}$, $\vec{a}_{-1} \geq_{V|_{X_{-1}:a_1}} \vec{b}_{-1}$ if and only if $(a_1, \vec{a}_{-1}) \geq_V (a_1, \vec{b}_{-1})$. For any \mathcal{O} -legal profile P , $P|_{X_{-1}:a_1}$ is composed of $V|_{X_{-1}:a_1}$ for all $V \in P$. For any CP-net \mathcal{N} , we let $\mathcal{N}|_{X_{-1}:a_1}$ denote the sub-CP-net of \mathcal{N} conditioned on $X_1 = a_1$. It follows that if $V \sim \mathcal{N}$, then, $V|_{X_{-1}:a_1} \sim \mathcal{N}|_{X_{-1}:a_1}$.

Now, we claim that for any pair of profiles $P_1, P_2 \in \prod_{j=1}^n LD(\mathcal{L}_j)$, $P_1 = (V_1, \dots, V_n)$ and $P_2 = (W_1, \dots, W_n)$, such that $a_1 = r_1(P_1) = r_1(P_2)$ and $P_1|_{X_{-1}:a_1} = P_2|_{X_{-1}:a_1}$, 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 $V_j^{a_1} \in \mathcal{L}_j|_{X_1}$ be an arbitrary linear order over D_1 in which a_1 is in the top position. Let $P = (Q_1, \dots, Q_n) \in \prod_{j=1}^n LD(\mathcal{L}_j)$

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.

- $\mathcal{N}_j|_{X_1} = V_j^{a_1}$.
- $\mathcal{N}_j|_{X_{-1}:a_1} = \hat{\mathcal{N}}_j|_{X_{-1}:a_1}$, where $\hat{\mathcal{N}}_j$ is the CP-net that V_j extends.

Let $\vec{a} = (a_1, \vec{a}_{-1}) = r(P_1)$. For any $1 \leq j \leq n$ and any $\vec{b} \in \mathcal{X}$ with $\vec{b} \succ_{Q_j} \vec{a}$, we have that the X_1 component of \vec{b} must be a_1 , because Q_j is lexicographic, and a_1 is in the top position of $Q_j|_{X_1}$. We let $\vec{b} = (a_1, \vec{b}_{-1})$. It follows that $\vec{b}_{-1} \succ_{Q_j|_{X_{-1}:a_1}} \vec{a}_{-1}$. We note that $Q_j|_{X_{-1}:a_1}$ is the lexicographic extension of $\mathcal{N}_j|_{X_{-1}:a_1}$, $V_j|_{X_{-1}:a_1}$ is the lexicographic extension of $\hat{\mathcal{N}}_j|_{X_{-1}:a_1}$, and $\mathcal{N}_j|_{X_{-1}:a_1} = \hat{\mathcal{N}}_j|_{X_{-1}:a_1}$. Therefore, $Q_j|_{X_{-1}:a_1} = V_j|_{X_{-1}:a_1}$, which means that $\vec{b}_{-1} \succ_{V_j|_{X_{-1}:a_1}} \vec{a}_{-1}$. Hence, we have $\vec{b} \succ_{V_j} \vec{a}$. By Lemma 12.3.2, 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 $r|_{X_{-1}:a_1}$ over $D_2 \times \cdots \times D_p$ such that for any $P \in \prod_{j=1}^n LD(\mathcal{L}_j)$,

$$r(P) = (r_1(P|_{X_1}), r|_{X_{-1}:r_1(P|_{X_1})}(P|_{X_{-1}:r_1(P|_{X_1})}))$$

At this point, we have shown that r can be decomposed as desired. We next show that for any $a_1 \in D_1$, $r|_{X_{-1}:a_1}$ is strategy-proof over $\prod_{j=1}^n LD(\mathcal{L}_j|_{X_{-1}:a_1})$. 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 $1 \leq j \leq n$, $top(\mathcal{N}_j|_{X_1}) = a_1$. That is, a_1 is ranked in the top position in the restriction of \mathcal{N}_j to X_1 . Also, $top(\hat{\mathcal{N}}_l|_{X_1}) = a_1$.
- For any $1 \leq j \leq n$, $\mathcal{N}_j|_{X_{-1}:a_1}$ is the CP-net over D_{-1} that V_j^{-1} extends; $\hat{\mathcal{N}}_l|_{X_{-1}:a_1}$ is the CP-net over D_{-1} that \hat{V}_l^{-1} extends.
- For any $1 \leq j \leq n$, $\mathcal{N}_j \in \text{CPnets}(\mathcal{L}_j)$; $\hat{\mathcal{N}}_l \in \text{CPnets}(\mathcal{L}_l)$.

The existence of these CP-nets is guaranteed by the richness of \mathcal{L}_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

$$\begin{aligned} r(P) &= (r_1(P|_{X_1}), r|_{X_{-1}:r_1(P|_{X_1})}(P|_{X_{-1}:r_1(P|_{X_1})})) \\ &= (a_1, r|_{X_{-1}:a_1}(P|_{X_{-1}:a_1})) = (a_1, r|_{X_{-1}:a_1}(P^{-1})) \\ &<_{V_l} (a_1, r|_{X_{-1}:a_1}(P_{-l}^{-1}, \hat{V}_l)) = r(P_{-l}, \hat{V}_l) \end{aligned}$$

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

Moreover, because r satisfies non-imposition, for any $a_1 \in D_1$, $r|_{X_{-1}:a_1}$ satisfies non-imposition. Hence, for any $a_1 \in D_1$, we can apply the induction assumption to $r|_{X_{-1}:a_1}$ 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 $\mathcal{M}(P)|_{X_i} \neq \mathcal{M}(P_{-l}, \hat{V}_l)|_{X_i}$. Let \vec{d}_i be the first $i - 1$ components of $\mathcal{M}(P)$ and $\mathcal{M}(P_{-l}, \hat{V}_l)$. Because $\mathcal{M}|_{X_i:\vec{d}_i}$ is strategy-proof, we have the following calculation.

$$\begin{aligned} \mathcal{M}(P)|_{X_i} &= \mathcal{M}|_{X_i:\vec{d}_i}(P|_{X_i:\vec{d}_i}) \\ &>_{V_l|_{X_i:\vec{d}_i}} \mathcal{M}|_{X_i:\vec{d}_i}(P_{-l}, \hat{V}_l|_{X_i:\vec{d}_i}) \\ &= \mathcal{M}(P_{-l}, \hat{V}_l)|_{X_i} \end{aligned}$$

Because V_l is lexicographic, for any $\vec{y}, \vec{z} \in D_{i+1} \times \dots \times D_p$, we have

$$(\vec{d}_i, \mathcal{M}|_{X_i:\vec{d}_i}(P), \vec{y}) >_{V_l} (\vec{d}_i, \mathcal{M}|_{X_i:\vec{d}_i}(P_{-l}, \hat{V}_l), \vec{z})$$

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

It follows from Theorem 12.3.1 that any sequential voting rule that is composed of locally strategy-proof voting rules is strategy-proof over lexicographic preference domains, because a sequential voting rule is a separable CR-net. Specifically, when the multi-issue domain is binary (that is, for any $1 \leq i \leq p$, $|D_i| = 2$), the sequential composition of majority rules is strategy-proof when the profiles are lexicographic. It is interesting to view this in the context of previous works on the strategy-proofness of sequential composition of majority rules: Lacy and Niou (2000) and Le Breton and Sen (1999) 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 (2009) showed that this rule is not strategy-proof when the profile is restricted to the set of all \mathcal{O} -legal profiles.

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

Theorem 12.3.3. *For any $1 \leq j \leq n$, suppose \mathcal{L}_j is a rich admissible conditional preference set, $L_j \subseteq \text{Pref}(\mathcal{L}_j)$, and L_j extends \mathcal{L}_j . If there exists $1 \leq j \leq n$ such that L_j is not the lexicographic preference domain of \mathcal{L}_j , then there exists a locally strategy-proof CR-net \mathcal{M} that satisfies non-imposition and is not strategy-proof over $\prod_{j=1}^n L_j$.*

Proof of Theorem 12.3.3: If, for some $j \leq n$, there is a $V_j' \in LD(\mathcal{L}_j)$ that is not in L_j , then there must also be a $V_j \in L_j$ that is not in $LD(\mathcal{L}_j)$, because some vote in L_j

must extend the CP-net that V_j' extends. Hence, if $\prod_{j=1}^n LD(\mathcal{L}_j) \neq \prod_{j=1}^n LD(\mathcal{L}_j)$, there must exist some $j \leq n$, $V_j \in L_j$ such that V_j is not in $LD(\mathcal{L}_j)$. For this V_j , there must exist $i \leq p$, $\vec{a}_{i-1} \in D_1 \times \cdots \times D_{i-1}$, $a_i, b_i \in D_i$, $\vec{a}_{i+1}, \vec{b}_{i+1} \in D_{i+1} \times \cdots \times D_p$ such that $a_i \succ_{V_j|_{X_i: \vec{a}_{i-1}}} b_i$, and $(\vec{a}_{i-1}, b_i, \vec{b}_{i+1}) \succ_{V_j} (\vec{a}_{i-1}, a_i, \vec{a}_{i+1})$. Now, let us define a CR-net \mathcal{M} as follows.

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

Now, let $\mathcal{N}_1 \in \text{CPnets}(\mathcal{L}_1)$ be a CP-net such that $\text{top}(\mathcal{N}_1) = (\vec{a}_{i-1}, a_i, \vec{a}_{i+1})$, and for any $k \geq i + 1$, $\text{top}(\mathcal{N}_1|_{X_k: \vec{a}_{i-1} b_i b_{i+1} \cdots b_{k-1}}) = b_k$. Here $\text{top}(\mathcal{N}_1)$ is the top-ranked alternative in \mathcal{N}_1 . Let $\mathcal{N}'_j \in \text{CPnets}(\mathcal{L}_j)$ be a CP-net such that $\text{top}(\mathcal{N}'_j) = (\vec{a}_{i-1}, b_i, \vec{b}_{i+1})$. 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{L}_1 , and L_j extends \mathcal{L}_j . For any profile $P = (V_1, \dots, V_j, \dots, V_n) \in \prod_{j=1}^n LD(\mathcal{L}_j)$ (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 $\mathcal{M}(P) = (\vec{a}_{i-1}, a_i, \vec{a}_{i+1})$, and $\mathcal{M}(P_{-j}, V'_j) = (\vec{a}_{i-1}, b_i, \vec{b}_{i+1})$, 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

12.4 Summary

In this chapter, we studied 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 this characterization only works for lexicographic domains.

Our characterization is quite positive; however, beyond that, it is still not clear how much we can hope for desirable strategy-proof voting rules in multi-issue domains.³ 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.

³ In fact, we also proved two impossibility theorems, which (informally) both state that as soon as we go beyond lexicographic domains, there are no strategy-proof voting rules, except CR-nets where local rules are dictatorships. These results are omitted due to their heavy technicality and notation. They can be found on my homepage.