Theoretical Economics 16 (2021), 1195–1220 1555-7561/20211195

Local-global equivalence in voting models: A characterization and applications

Ujjwal Kumar
Economic Research Unit, Indian Statistical Institute, Kolkata

Souvik Roy
Economic Research Unit, Indian Statistical Institute, Kolkata

Arunava Sen
Economics and Planning Unit, Indian Statistical Institute, New Delhi

Sonal Yadav
Department of Economics, Umeå University

Huaxia Zeng
School of Economics, Shanghai University of Finance and Economics and Key Laboratory of Mathematical Economics (SUFE), Ministry of Education

The paper considers a voting model where each voter's type is her preference. The type graph for a voter is a graph whose vertices are the possible types of the voter. Two vertices are connected by an edge in the graph if the associated types are "neighbors." A social choice function is locally strategy-proof if no type of a voter can gain by misrepresentation to a type that is a neighbor of her true type. A social choice function is strategy-proof if no type of a voter can gain by misrepresentation to an arbitrary type. Local-global equivalence (LGE) is satisfied if local strategy-proofness implies strategy-proofness. The paper identifies a condition

Ujjwal Kumar: ujju489@gmail.com
Souvik Roy: souvik.2004@gmail.com
Arunava Sen: asen@isid.ac.in
Sonal Yadav: sonal.crk@gmail.com
Huaxia Zeng: huaxia.zeng.2011@gmail.com
We are grateful to Shurojit Chatterji, Debasis Mishra, Eyal Winter, and several referees for valuable comments and suggestions. We would also like to thank the participants of the 2015 IDGP workshop at Universitat Autònoma de Barcelona, the conferences on economic design at Bilgi University, Istanbul (2015), the University of York (2017), and Corvinus University, Budapest (2019), the meetings of the Society for Social Choice and Welfare at Lund University (2016) and Seoul National University (2018), and the 2019 Conference on Economic Design and Algorithms at the Higher School of Economics, St. Petersburg as well as seminar participants at the Stockholm School of Economics, the City University of Hong Kong, and Ashoka University for their helpful comments. Huaxia Zeng acknowledges that his work was supported by the National Natural Science Foundation of China (No. 71803116), the Program for Professor of Special Appointment (Eastern Scholar) at Shanghai Institutions of Higher Learning (No. 2019140015), and the Fundamental Research Funds for the Central Universities (No. 2018110153). Sonal Yadav gratefully acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation.

© 2021 The Authors. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at https://econtheory.org. https://doi.org/10.3982/TE4177

on the graph that characterizes LGE. Our notion of "localness" is perfectly general. We use this feature of our model to identify notions of localness according to which various models of multidimensional voting satisfy LGE. Finally, we show that LGE for deterministic social choice functions does not imply LGE for random social choice functions.

Keywords. Local incentive constraints, strategy-proofness, mechanism design, strategic voting.

JEL classification. D71.

1. Introduction

Mechanism design theory is concerned with models where agents have private information (called a type) that has to be elicited by the mechanism designer. The cornerstone of the theory is the collection of strategy-proofness constraints that ensure that agents do not have incentives to misreport their types (or manipulate). The standard assumption in the theory is that the proposed social choice function must be immune to *all* possible misreports of agents. There is, however, considerable experimental evidence that agents do not always lie in an optimal payoff-maximizing way. For instance, Fischbacher and Föllmi-Heusi (2013) conduct an experiment where agents are paid money on the basis of a report of a privately observed roll of a die. In their results, only 20 percent of the subjects lie optimally, 39 percent are fully honest, and the remaining lie partially. Agents often choose to lie credibly by misreporting only to types that are near or close to their true types. We consider a model where an agent of a particular type can only misreport to an arbitrary set of pre-specified local types. Our main contribution is a complete answer to the following question: under what circumstances is immunity to misreporting via a local type (local strategy-proofness) equivalent to immunity to misreporting via an arbitrary type (strategy-proofness)?

The equivalence issue has important conceptual and practical implications.[1] If it is not satisfied, the mechanism designer can choose from a wider class of locally strategy-proof social choice functions. It may enable her, in principle, to avoid negative results such as the Gibbard–Satterthwaite theorem (Gibbard (1973), Satterthwaite (1975)). Alternatively, suppose that the problem at hand satisfies equivalence. So as to verify that a social choice function is strategy-proof, it suffices to check that it is locally strategy-proof. The latter is a simpler task because it involves checking fewer constraints.

We consider a model where an agent's type is a strict preference ordering over a finite set of alternatives. There are no monetary transfers. For convenience, we refer to this model as the voting model and refer to the agent as a voter, even though the model could apply to other settings such as matching. For our purpose, it is sufficient to restrict attention to the case of a single voter.[2] The set of possible preferences is called a *domain*. An *environment* is an undirected graph whose vertices are preferences in the domain. The agent whose preference is specified by a particular vertex can only misreport to another preference (or vertex) if the two vertices are connected by an edge in

[1] They are also discussed extensively in Carroll (2012) and Sato (2013).

[2] Our results can easily be interpreted in the multi-voter setting.

the environment. The set of vertices connected by an edge to a vertex are its *neighbors*. A social choice function is locally strategy-proof if no type of the agent can gain by manipulating to a neighbor; it is strategy-proof if the agent cannot gain by manipulating to any vertex in the graph. An environment satisfies *local-global equivalence* (LGE) if local strategy-proofness implies strategy-proofness.[3]

Section 2 of the paper contains some examples and observations that highlight the issues underlying LGE. It serves to motivate our main result in Section 3, Theorem 1, which is a characterization of environments that satisfy LGE. Section 4 contains a discussion of the computational complexity of Property L and its relationship with earlier results in the literature. Section 5 applies Theorem 1 to multidimensional voting environments. Finally, Section 6 uses Theorem 1 to construct an example of an environment where LGE holds but equivalence fails for random social choice functions.

The LGE property depends on the existence of certain types of paths in the environment. For every pair of preferences P and P' in the domain and alternative a, there must exist a path from P to P' satisfying a monotonicity property with respect to all alternatives that are ranked worse than a according to P. Specifically, the relative ranking of a and any alternative b ranked worse than a according to P, can change at most once along the path. We call this condition Property L. According to Theorem 1, Property L is both necessary and sufficient for LGE.

One of the strengths of our approach is that our notion of neighbors in the definition of local strategy-proofness is perfectly general. The earlier literature (discussed below) used the Kemeny distance metric to define "localness." Thus, two preferences are neighbors if there is a single pair of consecutively ranked alternatives that are switched between the two preferences. Preferences that are neighbors in this sense are referred to as being adjacent. A limitation of adjacency is that it excludes several multidimensional voting models that are of interest. In these models, an alternative is an m-tuple ($m > 1$) and preferences are typically assumed to satisfy some form of separability. Consequently, it is not always possible to switch a consecutively ranked pair of alternatives without affecting the ranking of other alternatives. We consider two such domains—separable domains and multidimensional single-peaked domains—and propose natural notions of neighbors such that the resulting environments satisfy LGE.

The question of local-global equivalence also arises naturally in the context of random social choice functions. We follow the standard approach of comparing lotteries via stochastic dominance (see Gibbard (1977)). Earlier results (again discussed below) suggest that environments that satisfy LGE for deterministic social choice functions also do so for random social choice functions. We use our characterization result for the deterministic case to show that this is not true generally. We construct an environment that satisfies Property L and, therefore, satisfies deterministic LGE. We also find a random social choice in the same environment that satisfies local strategy-proofness but violates strategy-proofness.

[3]The converse is, of course, always true.

1.1 *Related literature*

Two important papers on LGE in voting models are Carroll (2012) and Sato (2013). Both papers use the adjacency version of localness. Carroll (2012) considers random social choice functions and shows that specific preference domains, such as the set of all strict preferences, the set of all single-peaked preferences, and particular subsets of single-crossing preferences satisfy LGE. Sato (2013) provides a necessary condition and a stronger sufficiency condition for LGE in the context of deterministic social choice functions. Section 4.2 describes the relationship between Sato's results and ours in greater detail. As already mentioned, there are two significant ways in which our main result extends and refines the earlier analysis. The first is that our notion of neighbors is completely general; the second is that we have a complete characterization. Both aspects of our result permit a wider range of applications than was earlier possible.

Cho (2016) provides sufficient conditions for LGE with random social choice functions. The notion of neighbors is once again adjacency, but several notions of preference extensions to lotteries are considered. In particular, it shows that a stronger version of the sufficient condition proposed in Sato (2013) (see Property U in Section 4.2) is sufficient for LGE if lotteries are compared via stochastic dominance. We show in Section 6 that the condition that is necessary and sufficient for LGE with deterministic social choice functions (using adjacency as the notion of localness) is *not* sufficient for LGE with random social choice functions.

There are several papers that investigate LGE in models where monetary transfers to agents are permitted and preferences are quasilinear in the usual sense (see, for instance, Carroll (2012), Archer and Kleinberg (2014), and Mishra et al. (2016)). Although the basic question is the same, the flavor of the analysis and the results in the two models are very different from each other.

In a companion paper (Kumar et al. (2021)), we consider a multi-voter model and address the question, "Under what conditions on the environment is it the case that every locally strategy-proof social choice function that also satisfies the mild condition of unanimity[4] is also strategy-proof?" We show that a condition much weaker than Property L is sufficient for LGE in this sense for both deterministic and random social choice functions.

2. The model

Let $A = \{a, b, \dots\}$ denote a finite set of alternatives with $|A| \geq 2$. Throughout the paper, we assume that there is a single voter. This assumption is without loss of generality as is soon apparent.

A *preference* P is an antisymmetric, complete, and transitive binary relation over A, i.e., given $a, b \in A$, aPb is interpreted as a is strictly preferred to b according to P. Let $\mathcal{P}$ denote the set of all preferences: the set $\mathcal{P}$ is referred to as the *universal domain*. We refer to an arbitrary set $\mathcal{D} \subseteq \mathcal{P}$ as a *domain*.

[4]A deterministic social choice function satisfies unanimity if it always picks an alternative in a profile where it is first-ranked by all voters. In the case of a random social choice function such an alternative is picked with probability 1.

An *environment* is an (undirected) graph $G = \langle \mathcal{D}, \mathcal{E} \rangle$. The set of vertices of the graph is a domain $\mathcal{D}$. The set of edges is the set $\mathcal{E}$. If $P, P' \in \mathcal{D}$ and $(P, P') \in \mathcal{E}$, the two preferences are said to be *neighbors* or are *local*.

The notion of neighbors is perfectly general. One possible specification is that used by Carroll (2012) and Sato (2013). Fix a pair of preferences $P, P' \in \mathcal{D}$. Two alternatives a and b in A are *reversed* if aPb and $bP'a$, or bPa and $aP'b$. Let $P \triangle P' = \{\{a, b\} \subseteq A : a \text{ and } b \text{ are reversed in } P \text{ and } P'\}$ be the set of all reversed pairs of alternatives between P and P'.[5] Two preferences P and P' are called *adjacent* if $|P \triangle P'| = 1$.[6] An environment where neighbors are defined by adjacency is referred to as an *adjacency environment*. Whenever the notion of neighbors is defined by adjacency, we denote the set of edges by $\mathcal{E}^{\text{adj}}$. An adjacency environment typically is denoted by $G = \langle \mathcal{D}, \mathcal{E}^{\text{adj}} \rangle$. In Section 5, we provide an example of a nonadjacency environment.

Definition 1. A social choice function (SCF) is a map $f : \mathcal{D} \to A$.

Definition 2. Consider an environment $G = \langle \mathcal{D}, \mathcal{E} \rangle$. An SCF $f : \mathcal{D} \to A$ is locally manipulable at P if there exists $P' \in \mathcal{D}$ with $(P, P') \in \mathcal{E}$ such that $f(P')Pf(P)$. The SCF f is locally strategy-proof if it is not locally manipulable at any $P \in \mathcal{D}$.

Consider a graph or an environment. An SCF labels each vertex of the graph with an alternative. It is locally strategy-proof if the voter with a preference for a particular vertex cannot gain by misrepresenting her preference to a vertex that is a neighbor of her true preference.

In contrast to local strategy-proofness, an SCF is *strategy-proof* if the voter cannot gain by an arbitrary misrepresentation.

Definition 3. An SCF $f : \mathcal{D} \to A$ is manipulable at P if there exists $P' \in \mathcal{D}$ such that $f(P')Pf(P)$. The SCF f is strategy-proof if it is not manipulable at any $P \in \mathcal{D}$.

A strategy-proof SCF is clearly locally strategy-proof. We investigate the structure of an environment when the converse is true.

Definition 4. The environment $G = \langle \mathcal{D}, \mathcal{E} \rangle$ satisfies local-global equivalence (LGE) if every locally strategy-proof SCF $f : \mathcal{D} \to A$ is strategy-proof.

The next subsection makes some important observations regarding LGE.

2.1 *Preliminary observations*

Our goal in this subsection is to illustrate the issues involved in LGE and to provide some intuition behind our result. We begin with some standard concepts from graph theory.

[5] We are guilty of abuse of notation here. Since a preference is an ordered pair, $P \triangle P'$ should include both ordered pairs (a, b) and (b, a) if a and b are reversed in P and P'. In our notation, $P \triangle P'$ includes only the unordered pair $\{a, b\}$ in this case.

[6] An alternative and equivalent statement is that the Kemeny distance between P and P' is exactly 1.

Table 1. Domain $\mathcal{D}$.

P^1	P^2	P^3	P^4	P^5
c	c	c	c	c
$[a]$	$[b]$	$[b]$	$[b]$	a
b	a	a	a	$[b]$
z	z	z	z	z
v	v	v	u	u
w	w	u	v	v
u	u	w	w	w

Let $G = \langle \mathcal{D}, \mathcal{E} \rangle$ be an environment. A *path* $\pi = (P^1, \dots, P^t)$ is a sequence of *distinct* vertices in $\mathcal{D}$ satisfying the property that consecutive vertices are neighbors, i.e., $(P^k, P^{k+1}) \in \mathcal{E}$ for all $k = 1, \dots, t-1$.[7] Let $\Pi(P, P')$ denote the set of all paths from P to P' in G. For any path $\pi = (P^1, \dots, P^s, P^{s+1}, \dots, P^t)$, we let $\pi|_{[P^s, P^t]}$ denote the subpath $(P^s, P^{s+1}, \dots, P^t)$. We say G is *connected* if there exists a path between every pair of vertices in G, i.e., $\Pi(P, P') \neq \emptyset$ for all $P, P' \in \mathcal{D}$.

The example below highlights the reasons why LGE may fail.

Example 1. Let $A = \{a, b, c, z, u, v, w\}$. Consider the adjacency environment $G = \langle \mathcal{D}, \mathcal{E}^{\text{adj}} \rangle$, where $\mathcal{D} = \{P^1, P^2, P^3, P^4, P^5\}$ (Table 1). It is convenient to represent G by Figure 1.

The SCF $f : \mathcal{D} \to A$ picks a at P^1 and b at other preferences.[8] The SCF f is locally strategy-proof. However, it is not strategy-proof since the voter with preference P^5 can manipulate via P^1. ◊

The cause of the failure of strategy-proofness while maintaining local strategy-proofness can be clearly identified from Example 1. Consider the path $\pi = (P^5, P^4, P^3, P^2, P^1)$. The outcome at P^5 is b. Since b "improves" at P^4 relative to P^5, local strategy-proofness implies that the outcome at P^4 must be b; otherwise the voter would manipulate locally to P^5. Local strategy-proofness also implies that the outcomes at P^3 and P^2 must be b. Note that b "declines" at P^1 with respect to a. There are two options at P^1 that are consistent with the requirement of local strategy-proofness (with respect to P^1). The outcome can remain b or it can switch to a. In the former case, we maintain strategy-proofness since the outcome is b everywhere along the path π. However, if the outcome is a, a problem with strategy-proofness arises since a is preferred to b at P^5.

$$P^1 \overset{\{a,b\}}{\text{———}} P^2 \overset{\{w,u\}}{\text{———}} P^3 \overset{\{v,u\}}{\text{———}} P^4 \overset{\{b,a\}}{\text{———}} P^5$$

Figure 1. The environment $G = \langle \mathcal{D}, \mathcal{E}^{\text{adj}} \rangle$.[9]

[7] In other words, repetitions of vertices in a path are ruled out.

[8] This is indicated by the square brackets on the alternative chosen by f at each preference.

[9] Two vertices are connected by an edge in G if and only if the preferences represented by the vertices are adjacent. For instance, P^1 and P^2 are adjacent; in particular aP^1b and bP^2a. The edge between P^1 and P^2

The failure of LGE in $G = \langle \mathcal{D}, \mathcal{E}^{\text{adj}} \rangle$ arises from an inherent asymmetry in the "monotonicity" requirement imposed by local strategy-proofness. If the outcome of an SCF at a preference improves[10] relative to a local preference, the same outcome continues to be chosen at the new neighbor preference. However, if the outcome at a preference falls relative to a local preference, the new outcome can either remain the same or switch to an alternative that has improved (relative to the original outcome) in the new preference. Combining the latter option together with an improvement in the same path can lead to a failure of strategy-proofness without violating local strategy-proofness.

A key feature of the path π in Example 1 is that a and b switch relative ranking *more than once* in the path. Thus, aP^5b, bP^4a, and aP^1b. The preceding discussion makes it clear that such paths may be problematic for LGE.

Definition 5. Let $G = \langle \mathcal{D}, \mathcal{E} \rangle$ be an environment and let $a, b \in A$. A path $\pi = (P^1, P^2, \dots, P^t)$ satisfies no $\{a, b\}$ restoration if the relative ranking of a and b is reversed[11] at most once along π, i.e., there do not exist integers q, r, and s with $1 \leq q < r < s \leq t$ such that either (i) aP^qb, bP^ra and aP^sb or (ii) bP^qa, aP^rb and bP^sa.[12]

Let $P, P' \in \mathcal{D}$ and $a, b \in A$ be such that aPb. We say that b *overtakes* a in path $\pi \in \Pi(P, P')$ if bP^la for some preference P^l in the path π. The notion of overtaking can be used to restate the definition of an $\{a, b\}$ restoration in an obvious way. For instance, in case (i) of Definition 5, b overtakes a in the path $\pi^1 = (P^q, \dots, P^r)$ and a overtakes b in the path $\pi^2 = (P^r, \dots, P^s)$.

It is sometimes useful to consider paths without restoration for a pair of alternatives. Let $P, P' \in \mathcal{D}$ and $a, b \in A$ be such that aPb. Let $\pi = (P^1, P^2, \dots, P^t) \in \Pi(P, P')$ be a path without $\{a, b\}$ restoration. If $aP'b$, then aP^rb for all preferences P^r on the path π. Suppose $bP'a$ instead. Then there exists a unique preference P^r on π such that aP^sb for all $s = 1, \dots, r$ and bP^sa for all $s = r+1, \dots, t$.

To further clarify the relationship between the LGE property and paths without restoration, we make two modifications to Example 1.

Example 2. As in Example 1, $A = \{a, b, c, z, u, v, w\}$. We consider six additional preferences P^0, P^6, P^7, P^8, P^9, P^{10} as shown in Table 2. Let $\overline{\mathcal{D}}$ and $\mathcal{D}^*$ be the domains $\overline{\mathcal{D}} = \mathcal{D} \cup \{P^0\}$ and $\mathcal{D}^* = \mathcal{D} \cup \{P^6, P^7, P^8, P^9, P^{10}\}$. These domains are used to construct two adjacency environments $\overline{G} = \langle \overline{\mathcal{D}}, \mathcal{E}^{\text{adj}} \rangle$ and $G^* = \langle \mathcal{D}^*, \mathcal{E}^{\text{adj}} \rangle$. These environments are shown in Figures 2 and 3.

Consider $\overline{G}$ and a locally strategy-proof SCF $\bar{f} : \overline{\mathcal{D}} \to A$ such that $\bar{f}(P^5) = b$. Using the same arguments as in Example 1, along the path $\pi = (P^5, P^4, P^3, P^2, P^1)$, we can infer

is labelled $\{a, b\}$ so as to signify that the only "difference" between the two preferences is the ranking of a and b.

[10]We are intentionally informal in this description. These notions are made precise in due course.

[11]Recall that a pair of alternatives a, b is reversed in the pair of preferences P and P' if they are ranked differently in P and P'.

[12]It is worth emphasizing that in our definition of $\{a, b\}$ restoration, we are *not* referring to an ordered pair (a, b). Thus, $\{a, b\}$ restoration and $\{b, a\}$ restoration are the same in our definition. We use expressions such as "the path has no $\{a, b\}$ restoration" and "the path has no restoration for the pair $\{a, b\}$" interchangeably.

TABLE 2. Preferences P^0 and P^6, P^7, P^8, P^9, P^{10}.

P^0	P^6	P^7	P^8	P^9	P^{10}
c	a	a	a	a	a
a	c	c	c	c	c
b	b	z	z	z	b
z	z	b	b	b	z
v	u	u	v	v	v
u	v	v	u	w	w
w	w	w	w	u	u

that local strategy-proofness implies $\bar{f}(P^k) = b$ for all $k = 5, 4, 3, 2$ and $\bar{f}(P^1)$ is either b or a. Due to the presence P^0, there is now another path $\bar{\pi} = (P^5, P^0, P^1)$ from P^5 to P^1. This path has no $\{a, b\}$ restoration. Furthermore, the path $\bar{\pi}$ has the properties that (i) a and b are identically consecutively ranked, and (ii) c always ranks above a, while z, u, v, and w are all ranked below b. Clearly, b does not switch places with any other alternative along $\bar{\pi}$. As a result, local strategy-proofness forces the outcome of $\bar{f}$ to be b everywhere along $\bar{\pi}$, which rules out the manipulability of $\bar{f}$.

Now consider G^* and a locally strategy-proof SCF $f^* : \mathcal{D}^* \to A$ such that $f^*(P^5) = b$. Once again, local strategy-proofness along the path $\pi = (P^5, P^4, P^3, P^2, P^1)$ implies that $f^*(P^k) = b$ for all $k = 5, 4, 3, 2$ and $f^*(P^1)$ is either b or a. Consider the path $\pi^* = (P^5, P^6, P^7, P^8, P^9, P^{10}, P^1)$. Observe that π^* has no restoration for a and any of the alternatives in the set $Z = \{b, z, u, v, w\}$, which are all ranked below a in P^5. Alternatives of Z switch places among themselves along π^* (see, for example, the subpath $(P^6, P^7, P^8, P^9, P^{10})$). Consequently, the local strategy-proofness of f^* does not preclude the outcomes for preferences along π^* from belonging to Z. Suppose $f^*(P^1) = a$. Since $f^*(P^5) = b$, local strategy-proofness implies that some alternative in Z must "jump above" a and then "jump below" a (so as to conform with P^1) along the path π^*.[13] However, this is explicitly ruled out by the observation that π^* has no restoration for a and any of the alternatives in Z. Therefore, it must be the case that $f^*(P^1) = b$. In fact, only two possibilities can arise: (i) $f^*(P^k) = b$ for all $k = 1, \ldots, 10$ or (ii) $f^*(P^k) = b$ for all $k = 1, 2, 3, 4, 5, 6, 10$ and $f^*(P^{k'}) = z$ for all $k' = 7, 8, 9$. In either case, f^* is strategy-proof.

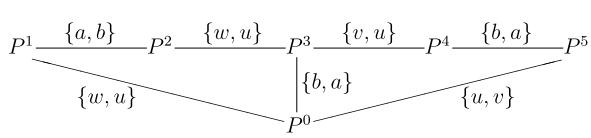

FIGURE 2. The environment $\overline{G} = \langle \overline{\mathcal{D}}, \mathcal{E}^{\text{adj}} \rangle$.

[13]We can first easily rule out the possibility that c is chosen at some preference in the subpath $(P^6, P^7, P^8, P^9, P^{10})$. In that case, local strategy-proofness forces the outcome of f^* to be c everywhere in G^*.

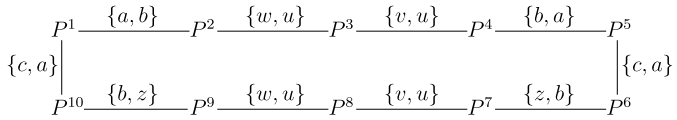

FIGURE 3. The environment $G^* = \langle \mathcal{D}^*, \mathcal{E}^{\text{adj}} \rangle$.

We conclude with an important observation. The alternative c is always ranked above a along the path $\bar{\pi}$ in $\overline{G}$. However, the path π^* in G^* does not forbid restoration between a and alternatives better than a in the initial preference P^5. $\Diamond$

We summarize the insights of Examples 1 and 2. There is "potential" for the failure of LGE whenever there is a path in an environment that has restoration for some pair of alternatives. However, LGE can be restored by the existence of certain "other" paths in the environment. As the argument relating to π^* in G^* suggests, the existence of a path that satisfies no restoration of an alternative with respect to all alternatives that are worse at a preference is sufficient to ensure strategy-proofness and, hence, LGE. In the next section, we show that this insight is general. In fact, this condition is also necessary, though the argument establishing necessity is more subtle.

3. The main result

The key condition for LGE is the *lower contour set no-restoration* property that we define below.

For any $P \in \mathcal{D}$ and $a \in A$, the lower contour set of a at P is the set of alternatives strictly worse than a according to P, i.e., $L(a, P) = \{b \in A : aPb\}$.

PROPERTY L. The environment G satisfies the lower contour set no-restoration property (Property L) if, for all $P, P' \in \mathcal{D}$ and $a \in A$, there exists a path $\pi \in \Pi(P, P')$ with no $\{a, b\}$ restoration for all $b \in L(a, P)$.

Pick an arbitrary pair of preferences $P, P' \in \mathcal{D}$ and an alternative $a \in A$ that is not ranked last in P. Suppose $L(a, P) = \{b_1, \dots, b_m\}$. If G satisfies Property L, there exists a path from P to P' that has no $\{a, b_i\}$ restoration for *all* $b_i \in \{b_1, \dots, b_m\}$. More informally, if a lies above b_i in P', then it lies above b_i everywhere along the path. Alternatively, if the rankings of a and b_i are reversed between P and P', there is a single reversal between a and b_i along the path.

The environment G^* in Example 2 satisfies Property L. In G^*, there are exactly two paths between any pair of vertices, one clockwise path and the other counterclockwise. For instance, between P^1 and P^5, the paths $(P^1, P^2, P^3, P^4, P^5)$ and $(P^1, P^{10}, P^9, P^8, P^7, P^6, P^5)$ are the clockwise and counterclockwise paths, respectively. These paths satisfy an important property. Fix an arbitrary pair of distinct preferences P and P'. If a path between P and P' possesses a restoration, say an $\{x, y\}$ restoration, and x is better than y in P, then the other path between P and P' must have no restoration for x and any alternative of $L(x, P)$. For example, consider P^1 and P^5. The

clockwise path $(P^1, P^2, P^3, P^4, P^5)$ has $\{a, b\}$ restoration and aP^1b. The counterclockwise path $(P^1, P^{10}, P^9, P^8, P^7, P^6, P^5)$ has no $\{a, x\}$ restoration for all $x \in L(a, P^1)$. The counterclockwise path $(P^1, P^{10}, P^9, P^8, P^7, P^6, P^5)$ has both $\{c, a\}$ restoration and $\{b, z\}$ restoration, cP^1a and bP^1z. Alternatively, the clockwise path $(P^1, P^2, P^3, P^4, P^5)$ has no $\{c, x\}$ restoration for all $x \in L(c, P^1)$ and no $\{b, x\}$ restoration for all $x \in L(b, P^1)$. This property ensures that G^* satisfies Property L.

Theorem 1. *An environment satisfies LGE if and only if it satisfies Property L.*

Proof. Sufficiency: Suppose $G = \langle \mathcal{D}, \mathcal{E} \rangle$ satisfies Property L but fails LGE, i.e., there exists a locally strategy-proof SCF $f : \mathcal{D} \to A$ that is not strategy-proof. Suppose f is manipulable at P. Define the alternative $x^1 = \max_P\{a \in A : f(\bar{P}) = a \text{ for some } \bar{P} \in \mathcal{D}\}$. In other words, x^1 is the highest-ranked alternative in the range of f according to P.[14] Let P' be such that $f(P') = x^1$. Since f is manipulable at P, we have $x^1 \neq f(P)$.

By Property L, there exists a path $\pi = (P^1, P^2, \dots, P^t) \in \Pi(P, P')$ that has no $\{x^1, z\}$ restoration for all $z \in L(x^1, P)$. Searching the path π backward from P^t to P^1, let P^s be the first vertex such that $f(P^s) = x^2 \neq x^1$, i.e., $f(P^k) = x^1$ for all $s < k \leq t$. Note that P^s always exists since $f(P^t) \neq f(P^1)$. It follows from the definition of x^1 that $x^1P^1x^2$. Since $(P^s, P^{s+1}) \in \mathcal{E}$, local strategy-proofness implies $x^2P^sx^1$ and $x^1P^{s+1}x^2$. We therefore have an $\{x^1, x^2\}$ restoration on the path π, contradicting our hypothesis. Therefore, $G = \langle \mathcal{D}, \mathcal{E} \rangle$ satisfies LGE and completes the proof of the sufficiency part of Theorem 1.

Necessity: We define a class of SCFs that we employ repeatedly in the proof.

Definition 6. Fix an environment $G = \langle \mathcal{D}, \mathcal{E} \rangle$. Let $a \in A$, let $\hat{P} \in \mathcal{D}$, and let B be a nonempty set with $B \subseteq L(a, \hat{P})$. An SCF $f : \mathcal{D} \to A$ is monotonic with respect to $(a, B, \hat{P})$ if the following statements hold:

(i) We have $f(P) = a$ if there is a path $\pi \in \Pi(\hat{P}, P)$ such that $B \subseteq L(a, \bar{P})$ for all $\bar{P} \in \pi$.

(ii) We have $f(P) = \max_P(B)$ otherwise.

Thus, $f(P) = a$ if there exists a path from $\hat{P}$ to P such that no alternative $x \in B$ overtakes a along the path (note that $a\hat{P}x$). Clearly $f(\hat{P}) = a$. The next lemma shows that the SCF f of Definition 6 is locally strategy-proof.

Lemma 1. *Suppose $f : \mathcal{D} \to A$ is monotonic with respect to $(a, B, \hat{P})$. Then f is locally strategy-proof.*

Proof. Pick an arbitrary pair $P, P' \in \mathcal{D}$ with $(P, P') \in \mathcal{E}$. We show either $f(P) = f(P')$ or $f(P)Pf(P')$ and $f(P')P'f(P)$, establishing local strategy-proofness.

Let $\mathcal{D}_a = \{\bar{P} \in \mathcal{D} : f(\bar{P}) = a\}$ denote the set of preferences that are associated to a at SCF f. There are four cases to consider.

Case 1: $P, P' \in \mathcal{D}_a$. Then $f(P) = f(P') = a$.

[14]For later reference, $\max_P(B)$ refers to the P-maximal alternative in the set $B \subseteq A$.

Case 2: $P, P' \notin \mathcal{D}_a$. Then $f(P) = \max_P(B)$ and $f(P') = \max_{P'}(B)$. Hence, either $f(P) = f(P')$ or $f(P)Pf(P')$ and $f(P')P'f(P)$ must hold.

Case 3: $P \in \mathcal{D}_a$ and $P' \notin \mathcal{D}_a$. Thus, $f(P) = a \neq b = \max_{P'}(B) = f(P')$. Since $P \in \mathcal{D}_a$, there exists a path $\pi = (P^1, \dots, P^t) \in \Pi(\hat{P}, P)$ such that $B \subseteq L(a, P^k)$ for all $1 \leq k \leq t$ (recall Definition 6). Since $b \in B$, we have aPb. Next, suppose $aP'b$. Since $b = \max_{P'}(B)$, it follows that $B \subseteq L(a, P')$. Observe that P' must be distinct from the vertices in the path π; otherwise, we would contradict the hypothesis that $P' \notin \mathcal{D}_a$. Since $(P, P') \in \mathcal{E}$, we now have a new path $\bar{\pi} = (P^1, \dots, P^t, P') \in \Pi(\hat{P}, P')$ such that $B \subseteq L(a, \bar{P})$ for all $\bar{P} \in \bar{\pi}$. Consequently, Definition 6 implies $f(P') = a$. This contradicts our initial assumption that $f(P') = b$. Therefore, $bP'a$.

Case 4: $P \notin \mathcal{D}_a$ and $P' \in \mathcal{D}_a$. This case is symmetric to Case 3 above and is omitted.

This completes the proof of the lemma. □

Lemma 1 and the LGE property imply that monotonic SCFs are also strategy-proof. This, in turn, imposes certain no-restoration conditions on the environment. The rest of the proof essentially shows that Property L is the consequence of the strategy-proofness of monotonic SCFs.

Let $G = \langle \mathcal{D}, \mathcal{E} \rangle$ be an environment satisfying LGE. We show that G satisfies Property L. We begin with an observation.

Claim 1. *The environment G is connected.*

Proof. Suppose the claim is false. Then there exists a component G' of G such that $G' \neq \emptyset$ and G' is a strict subset of G,[15] i.e., there does not exist a path from any vertex in G' to any vertex not in G'. Denote the set of vertices in G' by $\mathcal{D}'$. Pick an arbitrary vertex P^* in $\mathcal{D}'$ and let $a, b \in A$ be such that aP^*b. Define the SCF f as $f(P) = b$ for all vertices $P \in \mathcal{D}'$ and $f(P) = a$ for all $P \notin \mathcal{D}'$.

Clearly f is not strategy-proof because $f(P^*) = b$ while $f(P') = a$ for any $P' \notin \mathcal{D}'$. However, f is locally strategy-proof because the outcome does not change if the voter misrepresents via a neighboring preference.Thus, LGE is violated. □

Suppose G violates Property L, i.e., there exist $P^0, P^1 \in \mathcal{D}$ and $a \in A$ such that every path of $\Pi(P^0, P^1)$ has an $\{a, x\}$ restoration for some $x \in L(a, P^0)$. In view of Claim 1, this statement cannot hold vacuously.

Let Γ be the set of alternatives in $L(a, P^0)$ that appear in some restoration with a on some path of $\Pi(P^0, P^1)$:

$$\Gamma = \{x \in L(a, P^0) : \text{there exists } \pi \in \Pi(P^0, P^1) \text{ with } \{a, x\} \text{ restoration}\}.$$

Then the hypothesis for the contradiction can be restated as follows: each path of $\Pi(P^0, P^1)$ has an $\{a, x\}$ restoration for some $x \in \Gamma$.

For a specific path $\pi \in \Pi(P^0, P^1)$, let Γ_1^π denote the set of alternatives in $L(a, P^0)$ that appear in some restoration with a on the path π, i.e.,

$$\Gamma_1^\pi = \{x \in L(a, P^0) : \pi \text{ has } \{a, x\} \text{ restoration}\}.$$

[15]We say that G' is a component of G if G' is a maximal connected subgraph of G.

Let $\Gamma^1 \subseteq [\Gamma \cap L(a, P^1)]$ be the set of alternatives such that *every* path $\pi \in \Pi(P^0, P^1)$ has $\{a, x\}$ restoration for some $x \in \Gamma^1$. Note that either $\Gamma^1 \neq \emptyset$ or $\Gamma^1 = \emptyset$ holds, and every alternative in Γ^1 (if Γ^1 is nonempty) is ranked below a in both preferences P^0 and P^1. We show that each one of the two possible cases $\Gamma^1 \neq \emptyset$ and $\Gamma^1 = \emptyset$ leads to a contradiction.

Case A: $\Gamma^1 \neq \emptyset$. Let $f : \mathcal{D} \to A$ be the SCF that is monotonic with respect to (a, Γ^1, P^0). Note that f is well defined since $\emptyset \neq \Gamma^1 \subseteq L(a, P^0)$. According to Lemma 1, f is locally strategy-proof. We show that f is not strategy-proof.

According to Definition 6, $f(P^0) = a$. Pick an arbitrary path $\pi \in \Pi(P^0, P^1)$. By definition, there exists $z \in \Gamma^1$ such that π has $\{a, z\}$ restoration, i.e., there exists $P^r \in \pi$ such that $zP^r a$. Hence $\Gamma^1 \nsubseteq L(a, P^r)$. Since π was chosen arbitrarily, there does not exist $\bar{\pi} \in \Pi(P^0, P^1)$ such that $\Gamma^1 \subseteq L(a, P^s)$ for all $P^s \in \bar{\pi}$. Consequently, Definition 6 implies $f(P^1) = \max_{P^1}(\Gamma^1) \equiv b$. Since $\Gamma^1 \subseteq L(a, P^1)$, we have $f(P^0) = aP^1 b = f(P^1)$. Therefore, f is not strategy-proof and we have a contradiction to the assumption that G satisfies LGE.

This argument establishes that Case A cannot occur.

Case B: $\Gamma^1 = \emptyset$. This case is more complicated than the earlier one. We begin with a series of claims.

Claim 2. *There exists a path $\pi \in \Pi(P^0, P^1)$ such that $\Gamma_1^\pi \cap L(a, P^1) = \emptyset$.*

Proof. Suppose Claim 2 is false. This implies that in each path of $\Pi(P^0, P^1)$, at least one alternative involved in a restoration with a is ranked below a in P^1, i.e., $\Gamma_1^\pi \cap L(a, P^1) \neq \emptyset$ for all $\pi \in \Pi(P^0, P^1)$. Let $\hat{\Gamma} = \bigcup_{\pi \in \Pi(P^0, P^1)} [\Gamma_1^\pi \cap L(a, P^1)]$. Then $\emptyset \neq \hat{\Gamma} \subseteq L(a, P^1)$ and Case A holds with $\Gamma^1 = \hat{\Gamma}$. □

Following Claim 2, let $\pi^1 \in \Pi(P^0, P^1)$ be the path such that $\Gamma_1^{\pi^1} \cap L(a, P^1) = \emptyset$. Thus, $xP^1 a$ for all $x \in \Gamma_1^{\pi^1}$. Note that path π^1 has $\{a, x\}$ restoration only for all $x \in \Gamma_1^{\pi^1}$, and that $aP^0 x$ for all $x \in \Gamma_1^{\pi^1}$. Searching the path π^1 from P^1 back to P^0, let $P^2 \in \pi^1 \backslash \{P^1\}$ be the the first vertex such that a overtakes some alternative of $\Gamma_1^{\pi^1}$. Note that preference P^2 always exists since $xP^1 a$ and $aP^0 x$ for all $x \in \Gamma_1^{\pi^1}$. Let Z be the (nonempty) subset of alternatives in $\Gamma_1^{\pi^1}$ that are overtaken by a in the reverse path from P^1 to P^2, i.e., $Z \subseteq \Gamma_1^{\pi^1}$ such that (i) $aP^2 z$ for all $z \in Z$, (ii) $yP^2 a$ for all $y \in \Gamma_1^{\pi^1} \backslash Z$ (if $Z \neq \Gamma_1^{\pi^1}$), and (iii) $x\bar{P}a$ for all $x \in \Gamma_1^{\pi^1}$ and all $\bar{P} \in \pi^1|_{[P^2, P^1]} \backslash \{P^2\}$. Thus, subpath $\pi^1|_{[P^2, P^1]}$ has no $\{a, x\}$ restoration for any $x \in \Gamma_1^{\pi^1}$ and, hence, $P^2 \neq P^0$. Since π^1 has $\{a, x\}$ restoration only for all $x \in \Gamma_1^{\pi^1}$, path π^1 must have no $\{a, y\}$ restoration for any $y \in \Gamma \backslash \Gamma_1^{\pi^1}$ (if $\Gamma_1^{\pi^1} \neq \Gamma$). Therefore, subpath $\pi^1|_{[P^2, P^1]}$ has no $\{a, x\}$ restoration for any $x \in \Gamma$.

Claim 3. *The relationship $\Gamma \cap L(a, P^1)$ is a strict subset of $\Gamma \cap L(a, P^2)$.*

Proof. It follows from the definition of Z that if $\Gamma \cap L(a, P^1) \subseteq \Gamma \cap L(a, P^2)$, then $\Gamma \cap L(a, P^1)$ must be a strict subset of $\Gamma \cap L(a, P^2)$. Suppose it is not the case that $\Gamma \cap L(a, P^1) \subseteq \Gamma \cap L(a, P^2)$, i.e., there exists $x \in \Gamma \cap L(a, P^1)$ such that $xP^2 a$. Then we have

aP^0x, xP^2a, and aP^1x, which imply the $\{a, x\}$ restoration on π^1 and $x \in \Gamma_1^{\pi^1} \cap L(a, P^1)$. This contradicts the hypothesis $\Gamma_1^{\pi^1} \cap L(a, P^1) = \emptyset$. □

Claim 4. *For every $\hat{\pi} \in \Pi(P^0, P^2)$, there exists $x \in \Gamma$ such that $\hat{\pi}$ has $\{a, x\}$ restoration.*

Proof. Suppose there exists $\hat{\pi} \in \Pi(P^0, P^2)$ and $\hat{\pi}$ has no $\{a, x\}$ restoration for any $x \in \Gamma$. Clearly P^2 is a vertex common to both $\hat{\pi}$ and $\pi^1|_{[P^2, P^1]}$. Starting from P^1, proceed along the path that is the reverse of $\pi^1|_{[P^2, P^1]}$. Let $\tilde{P}$ be the first vertex in this reverse path that also belongs to $\hat{\pi}$. From our earlier remark, such a vertex must exist (it could be P^2). Now combine the sequences of vertices $\hat{\pi}|_{[P^0, \tilde{P}]}$ and $\pi^1|_{[\tilde{P}, P^1]}$ to form the vertex sequence $\bar{\pi}$. By construction, $\bar{\pi}$ contains no repetition of vertices so that it is a path and $\bar{\pi} \in \Pi(P^0, P^1)$.

For convenience, let $\bar{\pi} = (\bar{P}^1, \dots, \bar{P}^k, \dots, \bar{P}^t)$, where $\bar{P}^k = \tilde{P}$, $\hat{\pi}|_{[P^0, \tilde{P}]} = (\bar{P}^1, \dots, \bar{P}^k)$, and $\pi^1|_{[\tilde{P}, P^1]} = (\bar{P}^k, \dots, \bar{P}^t)$. Since $\bar{\pi} \in \Pi(P^0, P^1)$, the hypothesis for the contradiction of the necessity part of Theorem 1 implies $\Gamma_1^{\bar{\pi}} \neq \emptyset$. Therefore, there exists $b \in \Gamma$ such that $\bar{\pi}$ has $\{a, b\}$ restoration. Since neither $\hat{\pi}$ nor $\pi^1|_{[P^2, P^1]}$ has $\{a, b\}$ restoration and aP^0b, it must be the case that b overtakes a on the path $(\bar{P}^1, \dots, \bar{P}^k)$ and then a overtakes b on the path $(\bar{P}^k, \dots, \bar{P}^t)$. Thus, we have $b\bar{P}^k a$ and $a\bar{P}^t b$. Now refer back to the path π^1. Since aP^0b, $b\tilde{P}a$ and aP^1b, path π^1 has $\{a, b\}$ restoration and, hence, $b \in \Gamma_1^{\pi^1} \cap L(a, P^1)$. This contradicts the hypothesis $\Gamma_1^{\pi^1} \cap L(a, P^1) = \emptyset$. □

We can now replace P^1 by P^2 in our earlier arguments and define Γ^2 in the same way as we defined Γ^1. Once again, there are two possibilities, $\Gamma^2 \neq \emptyset$ and $\Gamma^2 = \emptyset$. The former case leads to an immediate contradiction using the arguments in Case A. In the latter case, we can apply Claims 2, 3, and 4 to infer the existence of P^3 such that (i) $\Gamma \cap L(a, P^2)$ is a strict subset of $\Gamma \cap L(a, P^3)$ and (ii) every path $\pi \in \Pi(P^0, P^3)$ has $\{a, x\}$ restoration for some $x \in \Gamma$. Repeating the argument, it follows that the only way to avoid a contradiction via Case A is to find an infinite sequence of vertices $P^1, P^2, \dots P^n, \dots$ such that

$$\big[\Gamma \cap L(a, P^1)\big] \subset \big[\Gamma \cap L(a, P^2)\big] \subset \cdots \subset \big[\Gamma \cap L(a, P^n)\big] \cdots .$$ [16]

However this is impossible in view of the finiteness of G. Thus, Case B cannot occur either and the proof is complete. □

Property L can be simplified if an additional restriction is imposed on the domain.

For any preference P, $r_1(P)$ denotes the first-ranked alternative in P. A domain $\mathcal{D}$ satisfies *minimal richness* if for all $a \in A$, there exists $P \in \mathcal{D}$ such that $r_1(P) = a$.

Property L′. The environment $G = \langle \mathcal{D}, \mathcal{E} \rangle$ satisfies Property L′ if the following two conditions hold:

(i) For all $P, P' \in \mathcal{D}$ with $r_1(P) = r_1(P') = a$, there exists a path $\pi = (P^1, \dots, P^t) \in \Pi(P, P')$ such that $r_1(P^k) = a$ for all $k = 1, \dots, t$.

[16] Each of the subset relations is strict.

(ii) For all $a \in A$ and $P' \in \mathcal{D}$ with $r_1(P') \neq a$, there exists $P \in \mathcal{D}$ with $r_1(P) = a$ and a path $\pi = (P^1, \dots, P^t) \in \Pi(P, P')$ that has no $\{a, b\}$ restoration for all $b \in A \backslash \{a\}$.

Property L′ is easier to verify than Property L. So as to verify the latter, we have to find the existence of a suitable path for all pairs of preferences and all alternatives not ranked last in one of the preferences. For part (i) of Property L′, we need only to check for the existence of a path with a simple property for all pairs of preferences with the *same* first-ranked alternative. For part (ii) of Property L′, we need only to verify the existence of appropriate paths for special pairs of preferences.

Proposition 1. *Properties L and L′ are equivalent on all environments $G = \langle \mathcal{D}, \mathcal{E} \rangle$ where $\mathcal{D}$ is minimally rich.*

Proof. Let $G = \langle \mathcal{D}, \mathcal{E} \rangle$ be an environment where $\mathcal{D}$ is minimally rich. We first show that Property L implies Property L′.

Pick $P, P' \in \mathcal{D}$ such that $r_1(P) = r_1(P') = a$. Since G satisfies Property L, there exists a path from P to P' with no $\{a, b\}$ restoration for all $b \in L(a, P) = A \setminus \{a\}$, Clearly, all preferences on this path must have a as the first-ranked alternative. To show part (ii) of Property L′, consider $a \in A$ and $P' \in \mathcal{D}$, where $r_1(P') \neq a$. By minimal richness, we can find $P \in \mathcal{D}$ with $r_1(P) = a$. Property L implies the existence of a path in $\Pi(P, P')$ that has no $\{a, b\}$ restoration for all $b \in L(a, P) = A \setminus \{a\}$. This is precisely the path required to satisfy part (ii) of Property L′.

We now show that Property L′ implies Property L. Pick $P, P' \in \mathcal{D}$ and $a \in A$. We have to show the existence of a path in $\Pi(P, P')$ that has no $\{a, b\}$ restoration for all $b \in L(a, P)$. There are four cases to consider.

Case 1: $r_1(P) = r_1(P') = a$. Part (i) of Property L′ guarantees the existence of a path that satisfies the required condition.

Case 2: $r_1(P) = a$ and $r_1(P') \neq a$. According to part (ii) of Property L′, there exist $P'' \in \mathcal{D}$ with $r_1(P'') = a$ and a path $\pi' \in \Pi(P'', P')$ such that π' has no $\{a, b\}$ restoration for any $b \neq a$. Let $\tilde{\pi} \in \Pi(P, P'')$ be the path whose existence is guaranteed by part (i) of Property L′. Let $\tilde{P}$ be the first vertex in the path $\tilde{\pi}$ (proceeding from P toward P'') that lies on π'. Such a vertex must exist since P'' belongs to both $\tilde{\pi}$ and π'. Let π be the sequence of vertices obtained by concatenating the subpaths $\tilde{\pi}|_{[P,\tilde{P}]}$ and $\pi'|_{[\tilde{P},P']}$. By construction, π does not contain any repetition of vertices. Therefore, $\pi \in \Pi(P, P')$. Since there is no $\{a, b\}$ restoration in π' for any $b \neq a$, there is no such restoration on its subpath $\pi'|_{[\tilde{P},P']}$ either. Also, a is first-ranked everywhere on the subpath $\tilde{\pi}|_{[P,\tilde{P}]}$. Therefore, π has no $\{a, b\}$ restoration for all $b \in A \setminus \{a\} = L(a, P)$.

Case 3: $r_1(P) \neq a$ and $r_1(P') = a$. According to Case 2, there exists a path $\pi' \in \Pi(P', P)$ that has no $\{a, b\}$ restoration for any $b \neq a$. Let π be the reverse of path π'. Then $\pi \in \Pi(P, P')$ and π has no $\{a, b\}$ restoration for all $b \in L(a, P)$.

Case 4: $r_1(P) \neq a$ and $r_1(P') \neq a$. By minimal richness, there exists $\bar{P} \in \mathcal{D}$ with $r_1(\bar{P}) = a$. Applying the argument in Case 3, there exists a path $\tilde{\pi} \in \Pi(P, \bar{P})$ with no $\{a, b\}$ restoration for any $b \in L(a, P)$. Applying Case 2, there exists a path $\hat{\pi} \in \Pi(\bar{P}, P')$ with no $\{a, b\}$ restoration for all $b \in A \setminus \{a\}$. Arguments similar to those in Case 2 can

now be used to construct an appropriate path from P to P'. Let $\tilde{P}$ be the first vertex in the path $\tilde{\pi}$ (proceeding from P to $\bar{P}$) that also lies on $\hat{\pi}$. Let π be the sequence of vertices obtained by the concatenation of the subpaths $\tilde{\pi}|_{[P,\tilde{P}]}$ and $\hat{\pi}|_{[\tilde{P},P']}$. Clearly $\pi \in \Pi(P, P')$. Since $\tilde{\pi}$ satisfies no $\{a, b\}$ restoration for all $b \in L(a, P)$ and $a = r_1(\bar{P})$, it follows that no alternative in $L(a, P)$ overtakes a in $\tilde{\pi}|_{[P,\tilde{P}]}$, i.e., $L(a, P) \subset L(a, \tilde{P})$. The subpath $\hat{\pi}$ satisfies no $\{a, b\}$ restoration for all $b \neq a$; therefore, the subpath $\hat{\pi}|_{[\tilde{P},P']}$ satisfies no $\{a, b\}$ restoration for all $b \in L(a, P)$. We can summarize the argument thus far as follows. Pick an arbitrary $b \in L(a, P)$ and consider the path π. If $aP'b$, then b lies everywhere less preferred to a along π. If $bP'a$, then b is less preferred to a in π until $\tilde{P}$ and overtakes a once from $\tilde{P}$ to P'. In other words, π satisfies no $\{a, b\}$ restoration for all $b \in L(a, P)$. □

In Section 5, we apply Property L′ to various environments so as to show LGE.

4. Discussion

We comment on some aspects of our results.

4.1 *Computational complexity*

The problem of determining whether an environment satisfies Property L is not computationally hard. The *depth first search algorithm*[17] for efficiently traversing graphs can be modified easily to construct an algorithm that decides whether an environment satisfies Property L. The worst case time complexity of the algorithm is $O(|A|^2|\mathcal{D}|(|\mathcal{D}| + |\mathcal{E}|))$, which is polynomial in the parameters of the problem. The details of the argument can be found in Chatterjee (2020).

4.2 *Relationship with earlier results*

Carroll (2012) proved that the the environments $\langle \mathcal{P}, \mathcal{E}^{\text{adj}} \rangle$ and $\langle \mathcal{D}^{\text{SP}}, \mathcal{E}^{\text{adj}} \rangle$ satisfy LGE.[18] Both these environments satisfy a stronger version of Property L that we refer to as Property U.

Property U. The environment $G = \langle \mathcal{D}, \mathcal{E} \rangle$ satisfies the universal pairwise no-restoration property (Property U) if for all $P, P' \in \mathcal{D}$, there exists a path in $\Pi(P, P')$ that satisfies no restoration for all pairs $\{a, b\}$.

Let $\pi \in \Pi(P, P')$ be the path that satisfies no restoration for all pairs of alternatives as required by Property U. Then π also satisfies no $\{a, b\}$ restoration for any $a \in A$ and $b \in L(a, P)$. Clearly, Property L is satisfied. Alternatively, Property L does not imply Property U. To see this, consider the environment G^* in Example 2, which satisfies Property L. For the pair (P^1, P^5), the clockwise path has $\{a, b\}$ restoration while the counterclockwise path has $\{c, a\}$ restoration. Clearly, Property U is violated.

[17] See Cormen et al. (2001).

[18] Recall that $\mathcal{P}$ is the set of all strict preferences. Also $\mathcal{D}^{\text{SP}}$ is the domain of single-peaked preferences. A formal definition of single-peaked preferences can be found in Section 5.

Sato (2013) showed that Property P below is necessary for LGE in adjacency environments.

Property P. The environment $G = \langle \mathcal{D}, \mathcal{E} \rangle$ satisfies the pairwise no-restoration property (Property P) if for all $P, P' \in \mathcal{D}$ and $a, b \in A$, there exists a path in $\Pi(P, P')$ that satisfies no $\{a, b\}$ restoration.

Example 3.2 in Sato (2013) shows that Property P is not sufficient for LGE. The difficulty is that Property P does not specify the relationship between the no-restoration paths for *different* pairs of alternatives: the path satisfying no restoration between P and P' for $\{a, b\}$ could be distinct from the no-restoration path between the same vertices for another pair $\{c, d\}$. Property L is clearly a strengthening of Property P.

Sato (2013) also introduced a sufficient condition for LGE in adjacency environments (we refer to this condition as Property S for convenience) that is weaker than Property U.

Property S. Let $G = \langle \mathcal{D}, \mathcal{E}^{\text{adj}} \rangle$ be an environment. Consider $P, P' \in \mathcal{D}$. A path $\pi = (P^1, P^2, \dots, P^t) \in \Pi(P, P')$ satisfies the antidote property with respect to the pair (P, P') if, for all pairs $a, b \in A$ such that π is with $\{a, b\}$ restoration and aP^1b, then for each $h \in \{1, \dots, t\}$ such that $bP^{h-1}a$ and aP^hb, there exists a path $\pi' \in \Pi(P, P^h)$ along which a does not overtake any alternative.

The environment G satisfies Property S if, for every $P, P' \in \mathcal{D}$, there exists a path satisfying the antidote property with respect to (P, P').

Environment G^* in Example 2 violates Property S, which establishes that Property S is stronger than Property L. Consider the pair (P^1, P^5). As noted earlier, the clockwise path from P^1 to P^5 has $\{a, b\}$ restoration since aP^1b, bP^4a, and aP^5b. For it to satisfy the antidote property, a should not overtake any alternative in the counterclockwise path from P^1 to P^5. However, a does overtake c on this path. Property L is nevertheless satisfied since there is no restoration with a and any of the alternatives ranked below a in P^1 along this path.

5. Multidimensional voting: The separable domain and the multidimensional single-peaked domain

In this section, we apply our results to a well known voting model. The set of alternatives has a Cartesian product structure, i.e., $A = \times_{j \in M} A_j$, where $M = \{1, 2, \dots, m\}$ is a finite set of *components* with $m \geq 2$. For each $j \in M$, the component set A_j contains a finite number of elements with $|A_j| \geq 2$. For any $j \in M$, $A_{-j} = \times_{i \neq j} A_i$. An alternative $a \in A$ is an m-tuple $a \equiv (a_1, \dots, a_m)$. We sometimes write a in the form (a_j, a_{-j}), where $a_j \in A_j$ and $a_{-j} \in A_{-j}$. A preference P is a linear order over A. A *marginal* preference over component j is a linear order over A_j.

A preference P is *separable* if, for all $a_j, b_j \in A_j$, $c_{-j}, d_{-j} \in A_{-j}$, and $j \in M$, $(a_j, c_{-j})P(b_j, c_{-j})$ implies $(a_j, d_{-j})P(b_j, d_{-j})$. Thus, every separable preference P induces an m-tuple of marginal preferences $(P_1, \dots, P_m)$.[19] Let $\mathcal{D}_S$ denote the set of *all* separable preferences. Note that for every component j and any marginal preference P_j over the component set A_j, there exists $P \in \mathcal{D}_S$ such that P induces the marginal preference P_j over A_j. There is a large literature on committee voting following Barberà et al. (1991), which assumes separable preferences.

Another domain of preferences that we consider is that of multidimensional single-peaked preferences introduced by Barberà et al. (1993). (See also Le Breton and Sen (1999).) This notion generalizes the well known class of single-peaked preferences (see Moulin (1980)). For this purpose, additional structure is introduced on each component set.

Let $\prec_j$ denote a linear order over A_j for each $j \in M$. A *grid* is an m-tuple $(\prec_1, \dots, \prec_m)$.[20] Let P be a preference over A whose first-ranked alternative is x. Then P is *multidimensional single-peaked* with respect to the grid $(\prec_1, \dots, \prec_m)$ if for all distinct $a, b \in A$, we have $[x_j \preceq_j a_j \prec_j b_j$ or $b_j \prec_j a_j \preceq_j x_j$ for all $j \in M$ with $a_j \neq b_j] \Rightarrow [aPb]$.[21]

The domain $\mathcal{D}_{\text{MSP}}$ contains preferences that are not separable (see Section 3 in Le Breton and Sen (1999)). However $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}} \neq \emptyset$. To see this, pick an arbitrary m-tuple of marginal preferences $(P_1, \dots, P_m)$, where each P_j, $j \in M$, is single-peaked with respect to $\prec_j$. Construct P as follows. For all distinct $c, d \in A$ with $c \neq d$, let j be the integer in M such that $c_j \neq d_j$ and $c_r = d_r$ for all $r < j$. Then cPd if and only if $c_j P_j d_j$. It is easy to verify that $P \in \mathcal{D}_S$. We also claim $P \in \mathcal{D}_{\text{MSP}}$. Suppose x is the first-ranked alternative in P. Pick distinct alternatives $a, b \in A$. Clearly, $a_j \neq b_j$ for some $j \in M$. Assume further that $x_j \preceq_j a_j \prec_j b_j$ or $b_j \prec_j a_j \preceq_j x_j$ for all $j \in M$ with $a_j \neq b_j$. Let $k \in M$ be the lowest component such that $a_k \neq b_k$. By virtue of the single-peakedness of P_k, $x_k \preceq_k a_k \prec_k b_k$ or $b_k \prec_k a_k \preceq_k x_k$ implies $a_k P_k b_k$. Then aPb follows directly from the construction of P.

We introduce a new notion of neighbors that applies to any domain that includes separable preferences. Let $P, P' \in \mathcal{D}_S$. We say that P and P' are *separably adjacent* (denoted by $(P, P') \in \mathcal{E}^{\text{SA}}$) if there exist $j \in M$ and $a_j, b_j \in A_j$ such that $[\{x, y\} \in P \triangle P'] \Rightarrow [x_j = a_j, y_j = b_j$ and $x_k = y_k$ for all $k \neq j]$. Thus, P and P' are separably adjacent if all pairs of alternatives that are reversed between P and P' differ in the values of exactly

[19]The converse is not true however. Several preferences can induce the same tuple of marginal preferences. For instance, consider *additively separable* preferences. Preferences over each component j have a utility representation $u_j : A_j \to \Re$. Utility representations over A are obtained by summing utilities over components. By considering different affine transformations of u_j, one can obtain different preferences over A without changing marginal preferences. Details can be found in Le Breton and Sen (1999).

[20]A grid can be interpreted as a product of lines. The notion of multidimensional single-peakedness can be generalized on a product of trees where our result still holds. For notational convenience, let $a_j \preceq_j b_j$ denote either $a_j \prec_j b_j$ or $a_j = b_j$.

[21]In the case where $m = 1$, multidimensional single-peakedness reduces to single-peakedness. The definition of multidimensional single-peakedness is silent regarding the comparison of some alternatives. For instance, suppose $m = 2$, $\prec$ is the $<$ ordering on real numbers, and $A_1 = A_2 = \{0, 1\}$. Let $(0, 0)$ be the highest-ranked alternative in the multidimensional single-peaked preference $\bar{P}$. We must have $(0, 0)\bar{P}(1, 0)$, $(0, 0)\bar{P}(0, 1)$, $(0, 0)\bar{P}(1, 1)$, $(1, 0)\bar{P}(1, 1)$, and $(0, 1)\bar{P}(1, 1)$ by definition.

TABLE 3. Domains $\mathcal{D}_S$ and $\mathcal{D}_{\text{MSP}}$.

P^1	P^2	P^3	P^4	P^5	P^6	P^7	P^8
$(0,0)$	$(0,0)$	$(0,1)$	$(0,1)$	$(1,0)$	$(1,0)$	$(1,1)$	$(1,1)$
$(0,1)$	$(1,0)$	$(0,0)$	$(1,1)$	$(0,0)$	$(1,1)$	$(0,1)$	$(1,0)$
$(1,0)$	$(0,1)$	$(1,1)$	$(0,0)$	$(1,1)$	$(0,0)$	$(1,0)$	$(0,1)$
$(1,1)$	$(1,1)$	$(1,0)$	$(1,0)$	$(0,1)$	$(0,1)$	$(0,0)$	$(0,0)$

one component.[22] We emphasize that separable adjacency applies *only* to separable preferences.

Separable adjacency does not cover the standard adjacency case. We, therefore, consider a strengthened version of separable adjacency: P and P' are *adjacent–separably adjacent* (denoted by $(P, P') \in \mathcal{E}^{\text{ASA}}$)[23] if either $(P, P') \in \mathcal{E}^{\text{adj}}$ or $(P, P') \in \mathcal{E}^{\text{SA}}$ holds. Two separable preferences P and P' are neighbors in the ASA sense if one can be obtained from the other by a "minimal" change.

EXAMPLE 3. Let $A = A_1 \times A_2$ with $A_1 = A_2 = \{0, 1\}$. In the special case $|A_j| = 2$ for all $j \in M$, we have $\mathcal{D}_S = \mathcal{D}_{\text{MSP}}$, implying that the environments $\langle \mathcal{D}_S, \mathcal{E}^{\text{ASA}} \rangle$ and $\langle \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{ASA}} \rangle$ are the same. Table 3 lists the preferences in $\mathcal{D}_S$ and $\mathcal{D}_{\text{MSP}}$. Note that the domain satisfies minimal richness.

This environment is shown in Figure 4. The thicker lines in the figure show the environment $\langle \mathcal{D}_S, \mathcal{E}^{\text{adj}} \rangle$, i.e., $\mathcal{E}^{\text{adj}} = \{(P^1, P^2), (P^3, P^4), (P^5, P^6), (P^7, P^8)\}$. The other edges in the figure belong to $\mathcal{E}^{\text{SA}}$. Note that $(P^1, P^2) \notin \mathcal{E}^{\text{SA}}$ since $P^1 \triangle P^2 = \{\{(0,1), (1,0)\}\}$. Also $P^1 \triangle P^3 = \{\{(0,0), (0,1)\}, \{(1,0), (1,1)\}\}$. Observe that the set of alternatives that are reversed between P^1 and P^3 can be obtained by switching the value of component 2 from 0 to 1 at different values of component 1. Clearly $(P^1, P^3) \in \mathcal{E}^{\text{SA}}$. Alternatively, $(P^2, P^4) \notin \mathcal{E}^{\text{SA}}$ since $\{(0,0), (1,1)\} \in P^2 \triangle P^4$.

We show later that the environment $\langle \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{ASA}} \rangle$ satisfies Property L′. Clearly, part (i) of Property L′ is satisfied as indicated by the four thick edges in Figure 4. Now consider the preference P^1 and the alternative $(1, 1)$, which is not first-ranked in P^1. We

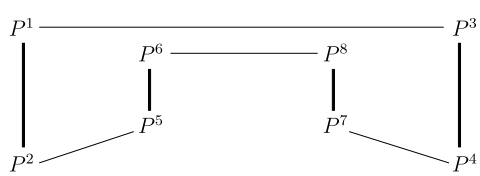

FIGURE 4. $\langle \mathcal{D}_S, \mathcal{E}^{\text{ASA}} \rangle$ and $\langle \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{ASA}} \rangle$.

[22]Separably adjacency is based on a notion of Kemeny distance that applies to separable preferences. Two (separable) preferences are separably adjacent if they disagree on the relative ranking of two alternatives that differ in the values of exactly one component. Further analysis of separable adjacency can be found in Chatterji and Zeng (2019).

[23]The acronym ASA stands for adjacent–separably adjacent.

have $(1, 1)$ first-ranked in preference P_8 and the path $(P^8, P^7, P^4, P^3, P^1)$ has no restoration for $(1, 1)$ and any other alternative. Consequently, the requirement of part (ii) of Property L′ is satisfied in this case. ◊

Example 3 and Figure 4 also lead to the conclusion that the environments $\langle \mathcal{D}_S, \mathcal{E}^{\mathrm{SA}} \rangle$, $\langle \mathcal{D}_{\mathrm{MSP}}, \mathcal{E}^{\mathrm{SA}} \rangle$, $\langle \mathcal{D}_S, \mathcal{E}^{\mathrm{adj}} \rangle$, and $\langle \mathcal{D}_{\mathrm{MSP}}, \mathcal{E}^{\mathrm{adj}} \rangle$ fail LGE. The graphs in these environments are not connected, which can be verified by inspection and by our earlier remarks.

According to the main result in this section, combining the adjacency and separable adjacency notions of neighbors with the separable and multidimensional single-peaked domains leads to LGE.

Proposition 2. *The environments $\langle \mathcal{D}_S, \mathcal{E}^{\mathrm{ASA}} \rangle$ and $\langle \mathcal{D}_{\mathrm{MSP}}, \mathcal{E}^{\mathrm{ASA}} \rangle$ satisfy LGE.*

The proof of Proposition 2 can be found in the Appendix.

6. LGE and random social choice functions

In this section, we examine LGE in the context of random social choice functions. Our result is that an environment that satisfies LGE for deterministic social choice functions may not satisfy LGE for random social choice functions.

Let $\Delta(A)$ denote the set of probability distributions over A. An element $\lambda \in \Delta(A)$ is referred to as a *lottery*. We let λ_a denote the probability with which $a \in A$ is selected by λ. Thus, $0 \leq \lambda_a \leq 1$ and $\sum_{a \in A} \lambda_a = 1$.

A *random social choice function* (RSCF) is a map $\varphi : \mathcal{D} \to \Delta(A)$ that associates a lottery $\varphi(P)$ with each $P \in \mathcal{D}$.

For every $P \in \mathcal{D}$, and $k = 1, 2, \ldots, |A|$, let $r_k(P) \in A$ denote the kth ranked alternative in P, i.e., $r_k(P) = a$ implies $|\{b \in A : bPa\}| = k - 1$. The lottery λ *stochastically dominates* (sd) lottery λ' at $P \in \mathcal{D}$ (denoted by $\lambda P_{\mathrm{sd}} \lambda'$) if $\sum_{k=1}^{t} \lambda_{r_k(P)} \geq \sum_{k=1}^{t} \lambda'_{r_k(P)}$ for all $t = 1, \ldots, |A|$.

Let $G = \langle \mathcal{D}, \mathcal{E} \rangle$ be an environment. A RSCF $\varphi : \mathcal{D} \to \Delta(A)$ is locally sd-strategy-proof if $\varphi(P) P_{\mathrm{sd}} \varphi(P')$ for all $(P, P') \in \mathcal{E}$. A RCSF $\varphi : \mathcal{D} \to \Delta(A)$ is sd-strategy-proof if $\varphi(P) P_{\mathrm{sd}} \varphi(P')$ for all $P, P' \in \mathcal{D}$.

The environment $G = \langle \mathcal{D}, \mathcal{E} \rangle$ satisfies *random local-global equivalence* (RLGE) if every locally sd-strategy-proof RSCF $\varphi : \mathcal{D} \to \Delta(A)$ is also sd-strategy-proof.

In the case where a RSCF is deterministic, local sd-strategy-proofness and sd-strategy-proofness reduce to local strategy-proofness and strategy-proofness, respectively. An immediate consequence of this observation is that an environment that satisfies RLGE also satisfies LGE. The results of Carroll (2012) and Cho (2016) show that the converse is true for several special domains. The example below shows that LGE does not imply RLGE.

Example 4. Let $A = \{a, b, c, v, w, x, y, z\}$. The domain $\tilde{\mathcal{D}}$ is described in Table 4. The environment $\tilde{G} = \langle \tilde{\mathcal{D}}, \mathcal{E}^{\mathrm{adj}} \rangle$ is shown in Figure 5.

By using arguments similar to those in Example 2, we can show that $\tilde{G}$ satisfies Property L. Therefore, Theorem 1 implies that $\tilde{G}$ satisfies LGE. We construct a RSCF that satisfies local sd-strategyproofness but not sd-strategy-proofness.

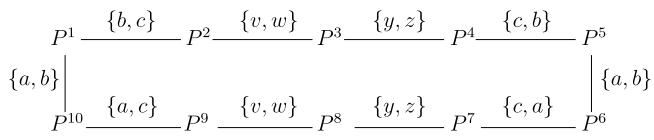

FIGURE 5. The environment $\tilde{G} = \langle \tilde{\mathcal{D}}, \mathcal{E}^{\text{adj}} \rangle$.

For any $d \in A$, we let e_d denote the degenerate lottery that picks d with probability 1. Consider the RSCF $\varphi : \tilde{\mathcal{D}} \to \Delta(A)$:

$$\varphi(P^k) = \begin{cases} \frac{1}{2}e_a + \frac{1}{2}e_b & \text{if } k \in \{1, 10\}, \\ \frac{1}{2}e_a + \frac{1}{4}e_b + \frac{1}{4}e_c & \text{if } k \in \{2, 3, 4, 5\}, \\ \frac{1}{4}e_a + \frac{1}{2}e_b + \frac{1}{4}e_c & \text{if } k \in \{6, 7, 8, 9\}. \end{cases}$$

So as to verify the local sd-strategy-proofness of φ, it suffices to show that the voter cannot gain by manipulation in each of the following cases: (i) from P^1 to P^2 and vice versa, (ii) from P^5 to P^6 and vice versa, and (iii) from P^9 to P^{10} and vice versa. This can be verified easily in each of the cases. Consider (i), for instance. Observe that c locally overtakes b from P^1 to P^2. Correspondingly, probability $\frac{1}{4}$ is transferred from b to c (keeping other probabilities fixed) as we move from $\varphi(P^1)$ to $\varphi(P^2)$. Therefore, $\varphi(P^2)P^2_{\text{sd}}\varphi(P^1)$ and, symmetrically, $\varphi(P^1)P^1_{\text{sd}}\varphi(P^2)$. The same argument can be made in cases (ii) and (iii).

However, it is not the case that $\varphi(P^5)P^5_{\text{sd}}\varphi(P^1)$ (in fact, $\varphi(P^1)P^5_{\text{sd}}\varphi(P^5)$). Consequently, φ is not sd-strategy-proof. $\Diamond$

We make two observations about Example 4.

OBSERVATION 1. As mentioned earlier, Carroll (2012) and Cho (2016) have established the equivalence of local sd-strategy-proofness and sd-strategy-proofness in specific adjacency environments. These environments all satisfy Property U. The environment $\tilde{G}$ in Example 4 violates Property U since both the clockwise and counterclockwise paths between P^1 and P^5 have restorations.

TABLE 4. Domain $\tilde{\mathcal{D}}$.

P^1	P^2	P^3	P^4	P^5	P^6	P^7	P^8	P^9	P^{10}
a	a	a	a	a	b	b	b	b	b
b	c	c	c	b	a	c	c	c	a
c	b	b	b	c	c	a	a	a	c
v	v	w	w	w	w	w	w	v	v
w	w	v	v	v	v	v	v	w	w
x	x	x	x	x	x	x	x	x	x
y	y	y	z	z	z	z	y	y	y
z	z	z	y	y	y	y	z	z	z

Observation 2. The key feature of the example in Example 4 that makes the LGE and RLGE results differ is that some lotteries under φ have support $\{a, b, c\}$, e.g., $\varphi(P^k)$, $k = 2, \dots, 9$. However, no locally strategy-proof SCF can have a range that includes all three alternatives a, b, and c. To see this, let $f : \tilde{\mathcal{D}} \to A$ be a locally strategy-proof SCF. Theorem 1 implies that f is strategy-proof. Suppose $\{a, b, c\} \subseteq \text{Range}(f) = \{d \in A : f(P) = d \text{ for some } P \in \tilde{\mathcal{D}}\}$. Thus, there exists a preference where f takes value a and another preference where f takes value b. Strategy-proofness immediately implies $f(P^k) = a$ for all $1 \le k \le 5$ and $f(P^l) = b$ for all $6 \le l \le 10$. Hence, we have a contradiction.

A characterization for RLGE appears to be significantly more difficult than that for LGE. In our companion paper Kumar et al. (2021), we derive a weak sufficient condition for RLGE in multi-voter models where RSCFs satisfy the additional property of unanimity.

Appendix: Proof of Proposition 2

We begin by observing that both the separable domain $\mathcal{D}_S$ and the multidimensional single-peaked (MSP) domain $\mathcal{D}_{\text{MSP}}$ satisfy the minimal richness property. Applying Theorem 1 and Proposition 1, it suffices to show that both domains satisfy Property L′. Furthermore both domains satisfy part (i) of Property L′ as is shown in Appendices E.2 and E.5 of Chatterji and Zeng (2019). Therefore, we only verify part (ii) of Property L′.[24]

We first investigate the separable domain $\mathcal{D}_S$. Next, we show part (ii) of Property L′ on the intersection of the separable domain and the multidimensional single-peaked domain $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$, and then extend the result to the multidimensional single-peaked domain $\mathcal{D}_{\text{MSP}}$.

In the proofs, we occasionally employ a special type of separable preferences called lexicographic separable preferences. Let $(P_1, \dots, P_m)$ be an m-tuple of marginal preferences and let P_0 be strict order over the set M. The preference P is *lexicographically separable* with respect to the $(m+1)$-tuple $(P_0, P_1, \dots P_m)$ if, for all $a, b \in A$, $[a_j P_j b_j$ and $a_r = b_r$ for all r such that $rP_0 j] \Rightarrow [aPb]$. In other words, a is ranked strictly better than b according to P if a_j is ranked higher than b_j according to the marginal preference P_j and $a_r = b_r$ for all components r that are ranked strictly higher than j according to the component preference P_0. We write a lexicographically separable preference P as $P \equiv (P_0, P_1, \dots, P_m)$.

We first prove two preliminary lemmas.

Lemma 2. *Let distinct $P, P' \in \mathcal{D}_S$ induce the same marginal preferences. Then there exists a path from P to P' in $\langle \mathcal{D}_S, \mathcal{E}^{\text{adj}} \rangle$ such that there is no restoration for any pair of alternatives.*

This lemma follows from Fact 5 of Chatterji and Zeng (2019).

[24]Part (i) of Property L′ is the same as the interior$^+$ property of Chatterji and Zeng (2019). Hence, we can directly apply their result for this part. However, part (ii) of Property L′ is stronger than their exterior$^+$ property, so we have to show this independently.

Lemma 3. *Fix marginal preferences $P_1, \dots, P_m$. Let a be an alternative such that a_j is not the first-ranked element in P_j for some $j \in M$. For each component k, let $X_k = \{x_k \in A_k : x_k P_k a_k\} \cup \{a_k\}$. Let $X = X_1 \times \cdots \times X_m$. Pick component j and let $b_j, c_j \in X_j$ or $b_j, c_j \in A_j \backslash X_j$ be consecutively ranked elements in P_j. Then there exists a separable ordering $\bar{P}(j)$ that satisfies the properties*

(i) $\bar{P}(j)$ induces the marginal preferences $P_1, \dots, P_m$

(ii) $[x\bar{P}(j)a] \Rightarrow$ [for each $k \in M$, either $x_k P_k a_k$ or $x_k = a_k$, i.e., $x \in X$]

(iii) (b_j, z_{-j}) and (c_j, z_{-j}) are consecutively ranked in $\bar{P}(j)$ for all $z_{-j} \in A_{-j}$.

Proof. We construct a partition of the set A. To do so, define the sets $A_{-j} = \times_{k \neq j} A_k$, $X_{-j} = \times_{k \neq j} X_k$, $Y_j = A_j \setminus X_j$, and $Y_{-j} = A_{-j} \setminus X_{-j}$. The sets X, $B = X_j \times Y_{-j}$, $C = Y_j \times X_{-j}$, and $D = Y_j \times Y_{-j}$ constitute a partition of the set A. The ordering $\bar{P}(j)$ is defined by two conditions:

(a) We have that $X\bar{P}(j)B\bar{P}(j)C\bar{P}(j)D$, i.e., all alternatives in X are ranked above those in B, which in turn are ranked above those in C, while all alternatives in D are ranked below those in C.

(b) We have that $\bar{P}(j)$ over X is lexicographically separable according to $(P_0(j), P_1, \dots, P_m)$, where j is ranked last in the component preference $P_0(j)$, i.e., given $x, y \in X$, $[x_k P_k y_k$ and $x_r = y_r$ for all $r P_0(j) k] \Rightarrow [x\bar{P}(j)y]$. Similarly, $\bar{P}(j)$ is lexicographically separable over alternatives in B, C, and D relative to $(P_0(j), P_1, \dots, P_m)$, respectively.

Observe that a_k is the lowest-ranked element in X_k according to P_k for all $k \in M$. Therefore, by construction, a is the worst alternative in X according to $\bar{P}(j)$. As X is the highest-ranked block according to $\bar{P}(j)$, it follows that all alternatives x that are ranked higher than a according to $\bar{P}(j)$ must satisfy $x \in X$. This establishes part (ii) of Lemma 3.

To show that $\bar{P}(j)$ is a separable preference and satisfies part (i) of Lemma 3, it suffices to show that for an arbitrary pair of alternatives that disagree in exactly one component, say $x = (x_k, z_{-k})$ and $y = (y_k, z_{-k})$, we have $[(x_k, z_{-k})\bar{P}(j)(y_k, z_{-k})] \Rightarrow [x_k P_k y_k]$. If x and y both belong to one of the sets X, B, C, or D, the result follows immediately. Henceforth, assume that x and y belong to two different sets of X, B, C, and D.

Suppose $k = j$. We know that either $z_{-j} \in X_{-j}$ or $z_{-j} \in Y_{-j}$. If $z_{-j} \in X_{-j}$, $(x_k, z_{-k})\bar{P}(j)(y_k, z_{-k})$ implies $x \in X$ and $y \in C$. Similarly, if $z_{-j} \in Y_{-j}$, $(x_k, z_{-k})\bar{P}(j)(y_k, z_{-k})$ implies $x \in B$ and $y \in D$. Consequently, in both cases, $x_j \in X_j$ and $y_j \in Y_j$, and, hence, $x_j P_j y_j$.

Suppose $k \neq j$. Let z_{-jk} denote the vector z_{-k} with its element of component j deleted. Since $x\bar{P}(j)y$, and x and y agree on component j, we know that either $x \in X$ and $y \in B$ or $x \in C$ and $y \in D$, both of which imply $(x_k, z_{-jk}) \in X_{-j}$ and $(y_k, z_{-jk}) \in Y_{-j}$. Since X_{-j} is a Cartesian product set, $(x_k, z_{-jk}) \in X_{-j}$ implies $x_k \in X_k$ and $z_{-jk} \in \times_{r \neq j,k} X_r$. Last, since $z_{-jk} \in \times_{r \neq j,k} X_r$, $(y_k, z_{-jk}) \notin X_{-j}$ implies $y_k \notin X_k$. Therefore, $x_k P_k y_k$.

Hence, $\bar{P}(j)$ is a separable preference and induces marginal preferences $P_1, \dots, P_m$.

Part (iii) of Lemma 3 is an immediate consequence of the fact that $\bar{P}(j)$ over alternatives of X and B, respectively, is lexicographically separable with respect to the component preference $P_0(j)$, where component j is ranked last. □

We now show that the separable domain $\mathcal{D}_S$ satisfies part (ii) of Property L′.

Proof of Proposition 2 in the environment $\langle \mathcal{D}_S, \mathcal{E}^{\text{ASA}} \rangle$. Consider $P' \in \mathcal{D}_S$ and $a \in A$ such that a is not the first-ranked alternative in P'. Let $P'_1, \dots, P'_m$ be the induced marginal preferences of P'. Without loss of generality, assume that $a_1, a_2, \dots, a_r$, $r \leq m$, are not first-ranked in $P'_1, P'_2, \dots, P'_r$, respectively, while $a_v = r_1(P'_v)$ for all $v = r + 1, \dots, m$. We construct a sequence of preferences that are edges in $\langle \mathcal{D}_S, \mathcal{E}^{\text{ASA}} \rangle$ with the property that a keeps "rising" along the sequence. The sequence terminates in a preference $P \in \mathcal{D}_S$, where a is first-ranked. Then the reverse path from P to P' has no $\{a, b\}$ restoration for all $b \in A \backslash \{a\}$, as required by part (ii) of Property L′.

We start from P'_1. Let $\mathcal{P}_1$ denote the set of all marginal preferences over A_1. Pick a marginal ordering P_1 such that a_1 is first-ranked. By Proposition 4.1 of Sato (2013), we have a path $\pi^1 = (P_1^1, \dots, P_1^t)$ from P'_1 to P_1 in $\langle \mathcal{P}_1, \mathcal{E}^{\text{adj}} \rangle$ which has no restoration for any pair of elements of A_1.[25] Since $L(a_1, P'_1) \subset L(a_1, P_1)$, a_1 must keep rising along the path π^1, i.e., $L(a_1, P_1^k) \subseteq L(a_1, P_1^{k+1})$ for all $1 \leq k < t$. Therefore, for all $1 \leq k < t$, if a_1 is involved in the local switching elements across P_1^k and P_1^{k+1}, it is true that $x_1 P_1^k a_1$ and $a_1 P_1^{k+1} x_1$ for some $x_1 \in A_1$.

For each $k = 1, \dots, t$, let $X_1^k = \{x_1 \in A_1 : x_1 P_i^k a_1\} \cup \{a_1\}$. For each $k = 1, \dots, t-1$, consider (P_1^k, P_1^{k+1}) and let $P_1^k \triangle P_1^{k+1} = \{\{b_1^k, c_1^k\}\}$. Since $L(a_1, P_1^k) \subseteq L(a_1, P_1^{k+1})$, it must be the case that either $b_1^k, c_1^k \in X_1^k$ or $b_1^k, c_1^k \in A_1 \backslash X_1^k$. Next, for each $k = 1, \dots, t$, by Lemma 3, let $\bar{P}^k(1) \in \mathcal{D}_S$ be such that (i) it induces the marginal preferences $P_1^k, P'_2, \dots, P'_m$, (ii) if $x \bar{P}^k(1) a$, then for all $j \in M$, either $x_j = a_j$, or x_j is strictly better than a_j according to the jth marginal ordering of $\bar{P}^k(1)$, and (iii) (b_1^k, z_{-1}) and (c_1^k, z_{-1}) are consecutively ranked in $\bar{P}^k(1)$ for all $z_{-1} \in A_{-1}$. Let $\hat{P}^k(1)$ be the ordering obtained by switching all alternatives of the type (b_1^k, z_{-1}) and (c_1^k, z_{-1}) for some $z_{-1} \in A_{-1}$. It is clear that $\hat{P}^k(1)$ is a separable preference with the same marginal preferences as $\bar{P}^k(1)$ for all components other than 1. For component 1, c_1^k is now ranked immediately above b_1^k, while the rankings of other elements are unchanged. Therefore, there are three properties of $\hat{P}^k(1)$ that are important: (a) $(\bar{P}^k(1), \hat{P}^k(1)) \in \mathcal{E}^{\text{SA}}$ and $\bar{P}^k(1) \triangle \hat{P}^k(1) = \{\{(b_1^k, z_{-1}), (c_1^k, z_{-1})\} : z_{-1} \in A_{-1}\}$; (b) $L(a, \bar{P}^k(1)) \subseteq L(a, \hat{P}^k(1))$, where the strict inclusion holds if and only if $a_1 = c_1^k$; (c) $\hat{P}^k(1)$ and $\bar{P}^{k+1}(1)$ have the same marginal preferences, and $L(a, \hat{P}^k(1)) \subseteq L(a, \bar{P}^{k+1}(1))$ by part (ii) of Lemma 3 in the construction of $\bar{P}^{k+1}(1)$.

Now, we have a sequence

$$P' \to \bar{P}^1(1) \to \hat{P}^1(1) \to \bar{P}^2(1) \to \cdots \to \bar{P}^{t-1}(1) \to \hat{P}^{t-1}(1) \to \bar{P}^t(1).$$

[25]For instance, we generate P_1 by moving a_1 directly to the top of P'_1 while keeping the rankings of other elements unchanged, and then construct the path from P'_1 to P_1 in $\langle \mathcal{P}_1, \mathcal{E}^{\text{adj}} \rangle$ by progressively moving a_1 to the top of P'_1.

Note that $\bar{P}^t(1)$ has marginal preference P_1, where a_1 is the first-ranked element. Since P' and $\bar{P}^1(1)$ have the same marginal preferences $P'_1, P'_2, \ldots, P'_m$, we know that either $P = \bar{P}^1(1)$ or there exists a path $\bar{\pi}^0$ from P to $\bar{P}^1(1)$ in $\langle \mathcal{D}_S, \mathcal{E}^{\text{adj}} \rangle$ that has no restoration for any pair of alternatives (by Lemma 2). Similarly, for all $1 \leq k < t$, we know that either $\hat{P}^k(1) = \bar{P}^{k+1}(1)$ or there exists a path $\bar{\pi}^k$ from $\hat{P}^k(1)$ to $\bar{P}^{k+1}(1)$ in $\langle \mathcal{D}_S, \mathcal{E}^{\text{adj}} \rangle$ that has no restoration for any pair of alternatives. Since $(\bar{P}^k(1), \hat{P}^k(1)) \in \mathcal{E}^{\text{SA}}$ for all $k = 1, \ldots, t-1$, we construct a concatenated path $\bar{\pi} = (\bar{\pi}^0, \bar{\pi}^1, \ldots, \bar{\pi}^{t-1})$ from P' to $\bar{P}^t(1)$ in $\langle \mathcal{D}_S, \mathcal{E}^{\text{ASA}} \rangle$.[26] Recall that $L(a, P') \subseteq L(a, \bar{P}^1(1))$, $L(a, \bar{P}^k(1)) \subseteq L(a, \hat{P}^k(1))$, and $L(a, \hat{P}^k(1)) \subseteq L(a, \bar{P}^{k+1}(1))$ for all $k = 1, \ldots, t-1$. Then no restoration on subpaths $\bar{\pi}^0, \bar{\pi}^1, \ldots, \bar{\pi}^{t-1}$ implies that a keeps rising along the path $\bar{\pi}$.

We can clearly repeat this procedure, progressively moving a_1 to the top in the marginal preference P_1 and then doing the same for a_2 until a_r. The procedure generates a path in $\langle \mathcal{D}_S, \mathcal{E}^{\text{ASA}} \rangle$ that culminates in a preference $P \in \mathcal{D}_S$, where a is first-ranked. Moreover if a overtakes some x at some preference on the path, it beats x at all preferences further along the path. It follows immediately that the reverse path from P to P' satisfies no $\{a, b\}$ restoration for all $b \in A \backslash \{a\}$. This establishes part (ii) of Property L′ and, hence, proves Proposition 2 for the separable domain $\mathcal{D}_S$. □

To show part (ii) of Property L′ in the multidimensional single-peaked domain $\mathcal{D}_{\text{MSP}}$, we first consider the domain $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$. We make several observations. First, $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$ satisfies part (i) of Property L′ by Appendix E.4 of Chatterji and Zeng (2019). Second, Lemma 2 remains valid in $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$ according to Fact 11 of Chatterji and Zeng (2019). Third, Lemma 3 holds when we set the marginal preferences $P_1, \ldots, P_m$ to be single-peaked with respect to $\prec_1, \ldots, \prec_m$, respectively, and change preference $\bar{P}(j)$ to be both separable and multidimensional single-peaked. Finally, in the verification of part (ii) of Property L′ in the separable domain, if we replace $\mathcal{D}_S$ with $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$, replace $\mathcal{P}_1$ with $\mathcal{S}_1$, which is the set of all single-peaked marginal preferences with respect to $\prec_1$, and replace the reference to Proposition 4.1 of Sato (2013) with a reference to Proposition 4.2 of Sato (2013), our earlier proof works for verifying part (ii) of Property L′ in $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$. Therefore, $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$ satisfies Property L′.

To extend the result to the multidimensional single-peaked domain, we use the following lemma, which follows from Lemma 8 of Chatterji and Zeng (2018).

Lemma 4. *Given distinct $P, P' \in \mathcal{D}_{\text{MSP}}$, let $r_1(P) = r_1(P')$. Then there exists a path from P to P' in $\langle \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{adj}} \rangle$ such that there is no restoration for any pair of alternatives.*

We now show part (ii) of Property L′ in the multidimensional single-peaked domain $\mathcal{D}_{\text{MSP}}$.

[26]The concatenated path $\bar{\pi}$ has no repeated preference. Given two preferences $\hat{P}$ and $\tilde{P}$ in $\bar{\pi}$, we know that $\hat{P} \in \bar{\pi}^k$ and $\tilde{P} \in \bar{\pi}^{k'}$ for some $0 \leq k, k' \leq t-1$. If $k = k'$, it is evident that $\hat{P} \neq \tilde{P}$ by the definition of the path $\bar{\pi}^k$. Next assume $k < k'$. Note that $\hat{P}^{k'}(1)$ and $\bar{P}^{k'+1}(1)$ induce the same marginal preference $P_1^{k'+1}$, and the path $\bar{\pi}^{k'}$ connecting $\hat{P}^{k'}(1)$ and $\bar{P}^{k'+1}(1)$ has no restoration for any pair of alternatives. Then $\tilde{P} \in \bar{\pi}^{k'}$ implies that $\tilde{P}$ induces the marginal preference $P_1^{k'+1}$. Symmetrically, $\hat{P}$ induces the marginal preference P_1^{k+1}, which is distinct from $P_1^{k'+1}$. Therefore, $\hat{P}$ and $\tilde{P}$ must be distinct.

Proof of Proposition 2 in the environment $\langle \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{ASA}} \rangle$. Consider $P' \in \mathcal{D}_{\text{MSP}}$ and $a \in A$ such that a is not the first-ranked alternative in P'. Let $r_1(P') = \bar{a}$. Fix $k \in M$. If $a_k = \bar{a}_k$, we pick an arbitrary single-peaked marginal preference P'_k that has a_k as the first-ranked element. If $a_k \neq \bar{a}_k$, we identify a particular single-peaked marginal preference P'_k that satisfies the condition $[x_k P'_k a_k] \Rightarrow [\bar{a}_k \preceq_k x_k \prec_k a_k$ or $a_k \prec_k x_k \preceq_k \bar{a}_k]$. The marginal preferences $P'_1, \dots, P'_m$ are single-peaked by construction. Applying the counterpart of Lemma 3, we have $\bar{P}' \in \mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$ such that $\bar{P}'$ induces $P'_1, \dots, P'_m$, and $[x\bar{P}'a] \Rightarrow [$for all $k \in M$, either $x_k = a_k$ or $x_k P'_k a_k]$. Note that $L(a, \bar{P}') \supseteq L(a, P')$. By Lemma 4, since $r_1(P') = r_1(\bar{P}')$, we have a path $\hat{\pi}$ from $\bar{P}'$ to P' in $\langle \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{adj}} \rangle$ that has no restoration for any pair of alternatives. Moreover, since $\mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$ satisfies Property L′, we have $P \in \mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}$ that has a first-ranked, and a path $\bar{\pi}$ from P to $\bar{P}'$ in $\langle \mathcal{D}_S \cap \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{ASA}} \rangle$ that has no $\{a, b\}$ restoration for all $b \neq a$.

Now we have a concatenated path $\pi = (\bar{\pi}, \hat{\pi})$ from P to P' in $\langle \mathcal{D}_{\text{MSP}}, \mathcal{E}^{\text{ASA}} \rangle$.[27] We show that π has no $\{a, b\}$ restoration for all $b \neq a$. Fix an arbitrary $b \neq a$. If b overtakes a on path $\bar{\pi}$, then no $\{a, b\}$ restoration on $\bar{\pi}$ implies that b overtakes a on $\bar{\pi}$ exactly once, and $b\bar{P}'a$. Then $L(a, \bar{P}') \supseteq L(a, P')$ implies $bP'a$, and no restoration on $\hat{\pi}$ from $\bar{P}'$ to P' implies $b\hat{P}a$ for all $\hat{P} \in \hat{\pi}$. Hence, the concatenated path π has no $\{a, b\}$ restoration. If b does not overtake a on path $\bar{\pi}$, then no $\{a, b\}$ restoration on $\bar{\pi}$ implies $a\bar{P}b$ for all $\bar{P} \in \bar{\pi}$ and, hence, $a\bar{P}'b$. Furthermore, no restoration on $\hat{\pi}$ implies that b can overtake a on $\hat{\pi}$ for at most once. Hence, the concatenated path π has no $\{a, b\}$ restoration. This establishes part (ii) of Property L′ and, hence, proves Proposition 2 for the multidimensional single-peaked domain $\mathcal{D}_{\text{MSP}}$. □

References

Archer, A. and R. Kleinberg (2014), "Truthful germs are contagious: A local-to-global characterization of truthfulness." *Games and Economic Behavior*, 86, 340–366. [1198]

Barberà, S., F. Gul, and E. Stacchetti (1993), "Generalized median voter schemes and committees." *Journal of Economic Theory*, 61, 262–289. [1211]

Barberà, S., H. Sonnenschein, and L. Zhou (1991), "Voting by committees." *Econometrica*, 59, 595–609. [1211]

Carroll, G. (2012), "When are local incentive constraints sufficient?" *Econometrica*, 80, 661–686. [1196, 1198, 1199, 1209, 1213, 1214]

Chatterjee, S. (2020), "Computational complexity of local-global equivalence." Available at https://drive.google.com/file/d/1XdHYBqxzc3CEXt2ahJP1giMSUoKQFOPp/view. [1209]

Chatterji, S. and H. Zeng (2018), "On random social choice functions with the tops-only property." *Games and Economic Behavior*, 109, 413–435. [1218]

Chatterji, S. and H. Zeng (2019), "Random mechanism design on multidimensional domains." *Journal of Economic Theory*, 182, 25–105. [1212, 1215, 1218]

[27] By an argument similar to the earlier one, the concatenated path π has no repeated preference.

Cho, W. J. (2016), "Incentive properties for ordinal mechanisms." *Games and Economic Behavior*, 95, 168–177. [1198, 1213, 1214]

Cormen, T., C. E. Leiserson, R. L. Rivest, and C. Stein (2001), *Introduction to Algorithms*, second edition. MIT Press, Boston. [1209]

Fischbacher, U. and F. Föllmi-Heusi (2013), "Lies in disguise—An experimental study on cheating." *Journal of the European Economic Association*, 11, 525–547. [1196]

Gibbard, A. (1973), "Manipulation of voting schemes: A general result." *Econometrica*, 41, 587–601. [1196]

Gibbard, A. (1977), "Manipulation of schemes that mix voting with chance." *Econometrica*, 45, 665–681. [1197]

Kumar, U., S. Roy, A. Sen, S. Yadav, and H. Zeng (2021), "Local global equivalence for unanimous social choice functions." *Games and Economic Behavior*, 130, 299–308. [1198, 1215]

Le Breton, M. and A. Sen (1999), "Separable preferences, strategyproofness and decomposability." *Econometrica*, 67, 605–628. [1211]

Mishra, D., A. Pramanik, and S. Roy (2016), "Local incentive compatibility with transfers." *Games and Economic Behavior*, 100, 149–165. [1198]

Moulin, H. (1980), "On strategy-proofness and single peakedness." *Public Choice*, 35, 437–455. [1211]

Sato, S. (2013), "A sufficient condition for the equivalence of strategy-proofness and nonmanipulability by preferences adjacent to the sincere one." *Journal of Economic Theory*, 148, 259–278. [1196, 1198, 1199, 1209, 1210, 1217, 1218]

Satterthwaite, M. A. (1975), "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions." *Journal of Economic Theory*, 10, 187–217. [1196]

Co-editor Federico Echenique handled this manuscript.

Manuscript received 11 March, 2020; final version accepted 17 November, 2020; available online 6 December, 2020.