Submitted to *Theoretical Economics*

Property-Rights Withdrawal in Aggregate Allocation

XIN CHENG
Shanghai University of Finance and Economics

YIXIN TAO
Shanghai University of Finance and Economics

WEIQIANG ZHENG
Yale University

We study aggregate allocation when voluntary participation is constrained by property-rights withdrawal: a withdrawing agent redeploys her own endowment, while still being affected by the residual coalition. We formalize participation using weak individual rationality (WIR), which is imposed only on the grand coalition and requires that no agent prefer unilateral withdrawal to the collective allocation. In an ordinal model, WIR alone uniquely selects the peak-proportional outcome: each agent's endowment is assigned to her top-ranked alternative. The characterization uses no efficiency, incentive, transfer, or coalition-stability conditions. The characterization extends to rich restricted preference domains, including graph domains, through a layered-separator richness condition. Cardinal models delineate the boundary of the ordinal result.

KEYWORDS. aggregate allocation, voluntary participation.
JEL CLASSIFICATION. D71, C71.

Xin Cheng: `2024214437@stu.sufe.edu.cn`
Yixin Tao: `taoyixin@mail.shufe.edu.cn`
Weiqiang Zheng: `weiqiang.zheng@yale.edu`

1. Introduction

Many collective allocation problems are not closed systems. Participants enter with property rights, budgets, liabilities, or mandates. A municipality owns its waste obligation, a donor owns her contribution, and a firm controls its investment capital. If a collective mechanism allocates these resources in a way that a participant finds unacceptable, the participant may withdraw her own endowment while the remaining coalition continues to allocate the residual endowment. The resulting outside option is therefore endogenous: it depends both on the withdrawing agent's unilateral action and on the mechanism used by the residual coalition.

Consider a consortium of municipalities assigning a total volume of waste, understood as a public chore, across processing sites. Each municipality generates a specific share of the waste. If the consortium allocates a landfill or waste-processing burden near municipality $\mathcal{A}$, thereby imposing high disutility on it, $\mathcal{A}$ may withdraw and process its own share at the site it finds least burdensome. Exit, however, does not remove $\mathcal{A}$ from the system: it remains exposed to the residual allocation generated by the residual consortium, potentially in a different way because the withdrawing municipality is no longer part of the coalition. Thus, $\mathcal{A}$ compares the collective allocation not with isolation, but with an allocation in which it processes its own share optimally while the remaining municipalities process the residual waste. A parallel logic applies in settings involving divisible resources, such as charitable donors or investment partners who pool endowments to fund public goods. Pooling resources enables coordinated impact, but each agent ultimately owns her contribution.

In such environments, the mechanism designer faces a stringent participation constraint: no agent should prefer unilateral withdrawal to the collective allocation. The following example shows that familiar majoritarian rules can violate this constraint.

Example Suppose there are two sites, x and y, and three equally sized municipalities, each generating one unit of waste. Site x is remote from municipalities 1 and 2 but close to municipality 3, while site y is remote from municipality 3 but less attractive to municipalities 1 and 2. Since waste-processing sites impose negative externalities on nearby municipalities, each municipality prefers waste to be assigned to a more remote site. This induces the following preferences:

$$\text{municipality } 1 \text{ and } 2: \; x \succ y, \qquad \text{municipality } 3: \; y \succ x.$$

Majority rule assigns all three units of waste to x, the majority's preferred site, but this allocation violates participation for municipality 3. Municipality 3 can withdraw its own unit of waste and send it to y. Suppose the residual coalition $\{1,2\}$ assigns the remaining two units to x under majority rule. The resulting outside allocation is $2\mathbf{1}_x + \mathbf{1}_y$, which municipality 3 strictly prefers to $3\mathbf{1}_x$. Thus, municipality 3 withdraws rather than accepting the majority allocation.

The example illustrates that participation with property-rights withdrawal is not a weak afterthought: it can overturn familiar majoritarian allocations.

This failure shows how voluntary participation disciplines collective choice and motivates our central question: Which mechanisms satisfy this participation constraint?

We study this question through *weak individual rationality* (WIR). Individual rationality requires, for any coalition, that no participant prefers its withdrawal allocation to the allocation prescribed by the mechanism. WIR imposes this requirement only at the grand coalition and only against unilateral withdrawals. It asks whether a single participant can profitably withdraw her own endowment from the grand coalition, taking as given the allocation rule used by the remaining participants.

Our results We characterize the grand-coalition outcome of mechanisms satisfying weak individual rationality (WIR). We first consider the unrestricted ordinal domain, in which every strict ranking of the alternatives is admissible. On this domain, we show that WIR alone pins down the grand-coalition allocation as peak-proportional: each agent's endowment weight is assigned to that agent's peak, their most preferred alternative. Remarkably, this characterization requires no efficiency, incentive, transfer, or coalitional-stability requirement. As an implication, voluntary participation alone makes the set of acceptable grand-coalition outcomes collapse to the allocation each agent can guarantee independently.

When the total endowment is normalized to one, the peak-proportional allocation coincides with the weighted random-dictatorship lottery. The random-dictatorship outcome is often characterized by strategy-proofness, where the weights represent the lottery properties. In contrast, the weights in a peak-proportional outcome represent agents' property rights and can be characterized solely by WIR rather than by strategy-proofness on unrestricted ordinal domains.

The unrestricted theorem provides a clean benchmark. If every ranking is admissible, then voluntary participation leaves no room for a mechanism to use one agent's endowment in support of another agent's priority. However, the set of admissible rankings is often restricted by structures in many applications. Sites, projects, facilities, and organizational designs usually have an underlying graphical structure. For example, rankings in reality often treat nearby sites, related projects, or adjacent exposure zones similarly.

Domain restriction is more realistic, but also raises a nontrivial robustness question. Indeed, WIR may not pin down peak-proportional outcome on arbitrary restricted domains (see Example 5). An interesting question is which domain-restriction structure preserves the strong characterization of WIR.

Our main restricted-domain results answer this question for graph-structured environments. Alternatives are vertices of a finite graph, and the graph records spatial proximity, technological relatedness, project similarity, or exposure links. In goods, services, facility-location, and public-investment problems, an agent's attractive alternatives are naturally organized into connected regions of the graph. A donor may view related projects as close substitutes; an investor may value clusters of technologically related opportunities. We capture this by upper k-peaked preferences: above every cutoff, the set of acceptable alternatives has at most k connected components. The parameter k allows several local peaks, but rules out preferences whose acceptable alternatives are scattered arbitrarily across unrelated parts of the network.

For chores, congestion, risk, and exposure problems, the relevant connectedness runs in the opposite direction. What is geographically or technologically coherent is not the set of attractive alternatives, but the set of sufficiently bad ones. A municipality exposed to waste sites, traffic congestion, or environmental risk may view harm as spreading through contiguous exposure zones. We capture this by lower k-dipped preferences: below every cutoff, the set of bad alternatives has at most k connected components. Thus, the same graph can discipline preferences in two economically distinct ways: connected acceptable regions for goods and connected exposure regions for bads.

On both graph domains, WIR again uniquely determines the peak-proportional grand outcome. This is not immediate. In connected domains, some alternatives cannot be tested by simply ranking them last; for example, a cut vertex may have to remain above some alternatives in any admissible order. The proof shows that these missing direct tests can be replaced by indirect contour-set tests. The general property behind this argument is layered separator richness: the domain contains enough admissible contour sets to isolate each coordinate of the allocation, possibly through lower-layer alternatives that have already been identified. Once WIR gives a lower bound for every coordinate, feasibility forces all bounds to bind, yielding peak-proportionality.

Layered separator richness also applies beyond graphs. We illustrate this with a business architecture investment domain in which alternatives are partitions of business functions. The example shows that the relevant object is not graph connectedness per se, but a more general ability of the domain to provide layered tests for every coordinate of the aggregate allocation.

Finally, we examine the power of WIR in cardinal domains, drawing the boundary of the ordinal characterization. When agents have linear utilities, the ordinal logic survives only with an additional incentive restriction. WIR by itself can accommodate residual-coalition improvements and therefore need not select the peak-proportional outcome; strategy-proofness eliminates this residual freedom, and on profiles with unique maximizers the grand allocation again assigns each agent's endowment to her favorite alternative. Nonlinear utilities change the picture more sharply. With concave CES preferences, withdrawal lets an agent use her own endowment to complete the residual allocation, and even the binary equal-weight economy can make the resulting outside options mutually infeasible: no grand allocation satisfies WIR. With convex L_p-norm preferences, withdrawal instead rewards concentration; in the same binary equal-weight setting, this force is incompatible with grand strategy-proofness. Thus, the ordinal theorem is not a generic consequence of participation alone. It relies on ordinal stochastic-dominance comparisons; once cardinal curvature matters, participation can become either too weak to identify peak-proportionality or too strong to be jointly feasible with incentive compatibility.

Literature review This paper is closely related to the literature on random dictatorship and random assignment. In randomized social choice, random dictatorship is typically obtained from strategy-proofness together with ordinal or stochastic-dominance requirements (Gibbard, 1977, Duggan, 1996, Sen, 2011). Related characterizations have also been obtained on cardinal utility domains (Hylland, 1980, Dutta et al., 2007). Moreover, Ozkes and Sanver (2023) characterize uniform random dictatorship using efficiency and an independence condition, without assuming strategy-proofness. In our setting, however, the weights in the weighted random dictatorship formula represent endowment rights rather than lottery probabilities.

In the random assignment literature on ordinal domains, stochastic dominance is used to compare fractional assignments: Bogomolnaia and Moulin (2001) introduce the probabilistic serial mechanism, and Che and Kojima (2010) show that it is asymptotically equivalent to random priority in large markets. By contrast, Hylland and Zeckhauser (1979) study random assignments on a cardinal utility domain and evaluate lotteries by expected utility. In our setting, the stochastic-dominance comparison is instead applied to aggregate allocations over public alternatives. This focus on public alternatives also connects our analysis to Conitzer et al. (2017), who study proportionality and other fairness notions in public decision making, where a single decision may benefit multiple agents simultaneously.

The paper is also related to the participation and no-show literature in social choice. In that literature, participation asks whether an agent is better off joining an election than abstaining. The no-show paradox and its relation to Condorcet consistency are studied by Fishburn and Brams (1983) and Moulin (1988), with later work extending the analysis to set-valued and probabilistic rules (Brandt et al., 2017). Our WIR condition has a similar participation structure, but a withdrawing agent keeps her endowment and faces the residual coalition's allocation.

A closer economic interpretation is provided by the literature on voluntary participation in public-good provision. Saijo and Yamato (2010) show that, once non-excludability is taken seriously, it may be impossible to design a mechanism under which every agent voluntarily participates. Furusawa and Konishi (2011) study which agents choose to contribute rather than free ride and characterize the resulting contributing coalition and public-good provision. More recently, Hirai and Shinohara (2024) examine how subsequent renegotiation with non-participants affects agents' initial participation decisions and the provision of the public good. Related individual rationality constraints appear in allocation problems with outside options, such as house allocation with existing tenants (Abdulkadiroğlu and Sönmez, 1999). Our main ordinal analysis abstracts from transfers and additional feasibility constraints to isolate the force of voluntary participation, while the cardinal results delineate the boundary of this characterization.

Finally, the paper contributes to the literature on restricted preference domains. Single-peakedness (Black, 1948, Moulin, 1980), tree domains (Demange, 1982), generalized median voter schemes (Barberà et al., 1993, Nehring and Puppe, 2007), and network location models (Schummer and Vohra, 2002) show that graph structure can yield positive social choice and mechanism design results. Here graph structure matters because connectedness must preserve enough preference variation for participation to identify the aggregate allocation. Layered separator richness formalizes this idea for the upper k-peaked graph domains and lower k-dipped graph domains, while the business architecture domain shows that the same logic extends beyond graphs.

The rest of the paper is organized as follows. Section 2 introduces the model and WIR. Section 3 defines the peak-proportional rule and proves its sufficiency. Section 4 characterizes WIR on the unrestricted domain. Section 5 develops layered separator richness as the abstract criterion. Section 6 verifies the criterion for upper k-peaked graph domains for goods, lower k-dipped graph domains for chores, and the business-architecture domain. Part of the proof for upper k-peaked graph domains is deferred to Appendix A, while the proof for the business architecture domain is provided in Appendix B. Section 7 proves the layered separator theorem. Section 8 gives linear, CES, and convex L_p-norm cardinal boundary results, with proofs collected in Appendix C.

2. MODEL

In this section, we introduce the basic notations. Throughout, $\mathbb{R}_+ = \mathbb{R}_{\geq 0}$.

Aggregate allocations Let X be a finite set of alternatives. A vector $\mathbf{q} \in \mathbb{R}_+^X$ assigns a nonnegative amount q_x to each alternative $x \in X$. For a subset $A \subseteq X$, write $q(A) = \sum_{x \in A} q_x$. In particular, $q(X)$ is the total mass of $\mathbf{q}$. For $x \in X$, let $\mathbf{1}_x \in \mathbb{R}_+^X$ denote the unit vector at x: $(\mathbf{1}_x)_x = 1$ and $(\mathbf{1}_x)_y = 0$ for every $y \neq x$.

The set of agents is $N = \{1, \ldots, n\}$. Agent $i \in N$ owns an exogenous endowment weight $w_i > 0$. These weights are fixed parameters of the problem, not reports to the mechanism. An allocation for coalition S is a vector whose total mass equals $\sum_{i \in S} w_i$. Its feasible set is

$$\mathcal{Q}_S = \{\mathbf{q} \in \mathbb{R}_+^X : q(X) = \sum_{i \in S} w_i\}.$$

Thus $\mathcal{Q}_\emptyset = \{\mathbf{0}\}$, where $\mathbf{0}$ is the zero vector in $\mathbb{R}_+^X$. The ordinal benchmark imposes no additional feasibility restrictions such as hard capacities, transportation costs, or other technological constraints.

Ordinal preferences and stochastic dominance A strict order $\succ$ on X represents an agent's ordinal ranking of alternatives. Its reflexive closure is denoted by $\succeq$. Since X is finite and $\succ$ is strict, there is a unique maximal element, denoted $p(\succ)$

and called the peak. For $x \in X$, define the upper and lower contour sets

$$U_{\succ}(x) = \{y \in X : y \succeq x\}, \qquad L_{\succ}(x) = \{y \in X : x \succeq y\}.$$

Thus $U_{\succ}(x)$ is the set of alternatives weakly preferred to x, and $L_{\succ}(x)$ is the set of alternatives weakly worse than x.

A domain $\mathcal{D}$ is a nonempty set of strict orders on X. Let $P(\mathcal{D}) = \{p(\succ) : \succ \in \mathcal{D}\}$ denote the set of alternatives that can arise as peaks in $\mathcal{D}$. Each agent i has her own nonempty domain $\mathcal{D}_i$ of strict orders on X. For a coalition S, a preference profile is a tuple $\succ_S = (\succ_i)_{i \in S} \in \prod_{i \in S} \mathcal{D}_i$. If $T \subseteq S$, then $\succ_T$ denotes the restriction of this profile to the agents in T. At a profile $\succ_S$, write $p_i = p(\succ_i)$ for agent i's peak.

DEFINITION 1 (Stochastic dominance). For allocations $\mathbf{q}, \mathbf{r} \in \mathbb{R}_+^X$ with $q(X) = r(X)$, $\mathbf{q}$ stochastically dominates $\mathbf{r}$ for preference $\succ_i$, written $\mathbf{q} \succeq_i^{sd} \mathbf{r}$, if

$$q(U_{\succ_i}(x)) \geq r(U_{\succ_i}(x)) \quad \text{for every } x \in X.$$

Strict dominance, written $\mathbf{q} \succ_i^{sd} \mathbf{r}$, obtains if at least one inequality is strict.

Mechanisms and withdrawal A mechanism is a family of allocation rules

$$M = \{M_S\}_{S \subseteq N}, \qquad M_S : \prod_{i \in S} \mathcal{D}_i \to \mathcal{Q}_S.$$

For each coalition S and profile $\succ_S$, the vector $M_S(\succ_S)$ is the aggregate allocation chosen for S. For the empty coalition, $M_\emptyset$ is the zero allocation in $\mathcal{Q}_\emptyset$.

The full family $\{M_S\}_{S \subseteq N}$ is needed because individual rationality is evaluated against withdrawal. Fix a nonempty coalition S, a profile $\succ_S \in \prod_{i \in S} \mathcal{D}_i$, an agent $i \in S$, and a withdrawal alternative $x \in X$. For any $x \in X$, we write $M_S(\succ_S)_x$ for the amount allocated to alternative x by the mechanism for coalition S at profile $\succ_S$. If agent i withdraws to x, she allocates her own endowment w_i to x, while the residual coalition $S \setminus \{i\}$ receives the allocation prescribed by the same mechanism family. The resulting outside allocation is $w_i \mathbf{1}_x + M_{S \setminus \{i\}}(\succ_{S \setminus \{i\}})$. This vector belongs to $\mathcal{Q}_S$, since its total mass is $w_i + \sum_{j \in S \setminus \{i\}} w_j = \sum_{j \in S} w_j$. The definition also covers the singleton case, where the residual coalition is empty and the residual allocation is zero.

DEFINITION 2 (Individual rationality and WIR). For a nonempty coalition $S \subseteq N$, a mechanism M satisfies *individual rationality at coalition* S, or *IR at* S, if, for every profile $\succ_S \in \prod_{i\in S} \mathcal{D}_i$, every agent $i \in S$, and every withdrawal alternative $x \in X$,

$$M_S(\succ_S) \succeq_i^{sd} w_i \mathbf{1}_x + M_{S\setminus\{i\}}(\succ_{S\setminus\{i\}}).$$

The condition central to the ordinal necessity results is *weak individual rationality, or WIR*: IR is imposed only at the grand coalition N. Equivalently, for every $\succ_N \in \prod_{i\in N} \mathcal{D}_i$, every $i \in N$, and every $x \in X$,

$$M_N(\succ_N) \succeq_i^{sd} w_i \mathbf{1}_x + M_{N\setminus\{i\}}(\succ_{N\setminus\{i\}}).$$

The adjective "weak" refers only to the scope of the axiom: WIR imposes no individual-rationality restriction on proper subcoalitions, except insofar as their rules appear in grand coalition outside options. By contrast, M satisfies *coalitional individual rationality* if it satisfies IR at every nonempty coalition $S \subseteq N$.

REMARK 1. The stochastic-dominance comparison makes the participation constraint robust to unobserved cardinal intensities. For equal-mass allocations on a finite ordered set, $\mathbf{q} \succeq_i^{sd} \mathbf{r}$ is equivalent to

$$\sum_{x\in X} v(x) q_x \geq \sum_{x\in X} v(x) r_x$$

for every cardinal representation $v : X \to \mathbb{R}$ satisfying $x \succ_i y \Rightarrow v(x) > v(y)$. Hence, if the mechanism's allocation fails to stochastically dominate a withdrawal allocation, some utility representation consistent with the reported order makes withdrawal strictly profitable. WIR therefore gives an ordinal, intensity-free participation guarantee.

REMARK 2. In the ordinal model, it is enough to test withdrawal to the agent's peak. The peak p_i belongs to every upper contour set of $\succ_i$, so, for every coalition S, residual allocation $M_{S\setminus\{i\}}(\succ_{S\setminus\{i\}})$, and alternative x,

$$w_i \mathbf{1}_{p_i} + M_{S\setminus\{i\}}(\succ_{S\setminus\{i\}}) \succeq_i^{sd} w_i \mathbf{1}_x + M_{S\setminus\{i\}}(\succ_{S\setminus\{i\}}).$$

Thus peak withdrawal is the strongest pure-withdrawal test. The same observation shows that allowing the withdrawing agent to split her own endowment across alternatives would not change the ordinal IR or WIR constraints: $w_i\mathbf{1}_{p_i}$ stochastically dominates every $\mathbf{d} \in \mathbb{R}^X_+$ with $d(X) = w_i$. We state the ordinal model in terms of pure withdrawal alternatives because the economic interpretation is that a withdrawing agent can place her endowment at a chosen site. The cardinal section allows split withdrawals explicitly, since curvature can make the best cardinal complement depend on the residual allocation.

3. PEAK-PROPORTIONAL RULE

This section introduces the benchmark rule and proves that it always satisfies coalitional individual rationality on any ordinal domain.

The benchmark rule assigns each endowment to the agent's peak. For every coalition S, define

$$M^{\mathrm{PPR}}_S(\succ_S) = \sum_{i\in S} w_i \mathbf{1}_{p_i}.$$

We call M^{PPR} the peak-proportional rule. If $w_N = 1$, this is the weighted random-dictatorship lottery. In the present model the same expression is a deterministic aggregate allocation.

PROPOSITION 1 (Sufficiency of the peak-proportional rule). *On every ordinal domain $\prod_{i\in S}\mathcal{D}_i$ for every coalition $S \subseteq N$, the peak-proportional rule satisfies coalitional individual rationality.*

PROOF. Fix a nonempty coalition S, a profile $\succ_S \in \prod_{i\in S}\mathcal{D}_i$, an agent $i \in S$, and a withdrawal alternative x. We have $M^{\mathrm{PPR}}_S(\succ_S) = w_i\mathbf{1}_{p_i} + M^{\mathrm{PPR}}_{S\setminus\{i\}}(\succ_{S\setminus\{i\}})$. For any upper contour set U of $\succ_i$, the peak p_i belongs to U. Hence

$$\sum_{\ell\in U} M^{\mathrm{PPR}}_S(\succ_S)_\ell = w_i + \sum_{\ell\in U} M^{\mathrm{PPR}}_{S\setminus\{i\}}(\succ_{S\setminus\{i\}})_\ell \geq w_i \mathbf{1}_{\{x\in U\}} + \sum_{\ell\in U} M^{\mathrm{PPR}}_{S\setminus\{i\}}(\succ_{S\setminus\{i\}})_\ell.$$

The inequality holds for every upper contour set, so the coalition allocation stochastically dominates the withdrawal allocation. Since S, $\succ_S$, i, and x were arbitrary, coalitional individual rationality follows. □

The rest of the ordinal analysis proves necessity. The main theorems impose WIR only and therefore characterize only M_N. The residual rules $M_{N\setminus\{i\}}$ are otherwise unrestricted, but they enter the participation constraints as continuation outcomes after unilateral withdrawal. This is the only way in which subcoalition rules are used in the grand coalition characterization.

4. UNRESTRICTED ORDINAL CHARACTERIZATION

We first give the unrestricted proof before introducing the layered separator richness. The proof is a coordinate-identification argument. For each alternative a, we show that the grand allocation must assign at least the total endowment of agents peaking at a. Since these lower bounds sum to the total endowment, feasibility turns all inequalities into equalities.

LEMMA 2 (Direct-test lower bound). *Fix $a \in X$ and a grand profile $\succ_N$. Suppose that every agent j with $p_j = a$ has some admissible order $\tau_j \in \mathcal{D}_j$ that ranks a last. If M satisfies WIR, then*

$$M_N(\succ_N)_a \geq \sum_{j:p_j=a} w_j.$$

PROOF. If $X = \{a\}$, then every feasible grand allocation has a-coordinate w_N, so the result is immediate. Hence assume $|X| \geq 2$.

We prove the result by induction on $k = \big|\{j \in N : p_j = a\}\big|$. If $k = 0$, the claim reduces to $M_N(\succ_N)_a \geq 0$, which follows from nonnegativity.

Assume the result has been proved for all profiles with fewer than k agents peaking at a, and consider a profile with exactly $k \geq 1$ such agents. Choose one of these agents and call her i. Thus $p_i = a$. By the hypothesis of the lemma, there exists $\tau_i \in \mathcal{D}_i$ that ranks a last.

At the modified profile $(\tau_i, \succ_{-i})$, agent i no longer peaks at a. Therefore exactly $k-1$ agents peak at a. The hypothesis of the lemma remains valid at the modified profile: every remaining agent who peaks at a already peaked at a in the original profile, and therefore still has an admissible order in her own domain that ranks

a last. Hence the induction hypothesis gives

$$M_N(\tau_i, \succ_{-i})_a \geq \sum_{j \neq i : p_j = a} w_j. \tag{1}$$

Because a is ranked last under τ_i, the set $X \setminus \{a\}$ is an upper contour set of τ_i. WIR at the modified profile, applied to agent i with withdrawal to $p(\tau_i)$, gives

$$M_N(\tau_i, \succ_{-i})_a \leq M_{N \setminus \{i\}}(\succ_{-i})_a. \tag{2}$$

Combining (1) and (2), we obtain

$$M_{N \setminus \{i\}}(\succ_{-i})_a \geq \sum_{j \neq i : p_j = a} w_j.$$

Return now to the original profile $\succ_N$. Since $p_i = a$, the singleton $\{a\}$ is an upper contour set of $\succ_i$. WIR for agent i, with withdrawal to a, gives

$$M_N(\succ_N)_a \geq w_i + M_{N \setminus \{i\}}(\succ_{-i})_a \geq \sum_{j : p_j = a} w_j.$$

This completes the induction. □

THEOREM 3 (Unrestricted ordinal characterization). *Suppose $\mathcal{D}_i = \mathcal{A}(X)$ for every i, where $\mathcal{A}(X)$ is the set of all strict orders on X. If M satisfies WIR, then, for every grand profile $\succ_N$,*

$$M_N(\succ_N) = \sum_{i \in N} w_i \mathbf{1}_{p_i}.$$

PROOF. The result follows immediately from Lemma 2 and the identity $\sum_a M_N(\succ_N)_a = \sum_{i \in N} w_i$. □

5. LAYERED SEPARATOR RICHNESS

This section introduces the richness condition used in the main ordinal theorem. The condition identifies when a restricted domain contains sufficient preference variation for WIR to determine each coordinate of the grand allocation.

The benchmark case is direct testability. If every peakable alternative can also be ranked in last place, then WIR identifies the corresponding coordinate directly, and the grand allocation is uniquely pinned down by the peak-proportional rule.

Direct testability alone, however, is too restrictive for more general domains. As a simple example, suppose $X = \{a, b, c, d\}$, and suppose all alternatives can be ranked first, but only b, c, and d can be ranked last. In this case, the domain is not directly testable.

Layered separator richness covers such situations by testing a indirectly. It uses admissible upper contour sets that exclude a but contain lower-ranked peaks that have already been tested. The separator identity then combines the resulting inequalities so that, through feasibility, all non-target coordinates cancel, isolating the coordinate of a.

DEFINITION 3 (Layered separator richness). A domain $\mathcal{D}$ is layered separator-rich if there exist an integer $R \geq 0$ and a rank function

$$\rho : P(\mathcal{D}) \to \{0, \ldots, R\}$$

such that every $a \in P(\mathcal{D})$ satisfies the following rank-dependent requirements.

(L0) If $\rho(a) = 0$, then there is an order $\tau \in \mathcal{D}$ that ranks a last.

(L1) If $\rho(a) > 0$, then there are finitely many orders $\tau_h \in \mathcal{D}$, upper contour sets C_h of τ_h, coefficients $\alpha_h > 0$, and a scalar $\beta \geq 0$ such that, with $H_h = X \setminus C_h$,

$$p(\tau_h) \in C_h, \qquad a \notin C_h, \qquad \rho(p(\tau_h)) < \rho(a) \quad \text{for every } h,$$

and

$$\sum_h \alpha_h \mathbf{1}_{H_h} = \beta \mathbf{1}_X + \mathbf{1}_{\{a\}} \quad \text{as functions on } X.$$

The definition assigns a rank $\rho(a)$ to each peakable alternative a. Rank-zero alternatives are those that the domain can place last. A positive-rank alternative need not be rankable last, but it must admit a finite separator certificate. Such a certificate consists of admissible orders τ_h and upper contour sets C_h that exclude a, while the peaks $p(\tau_h)$ have already appeared at lower ranks. The complements

$H_h = X \setminus C_h$ must satisfy the displayed identity: their weighted indicators count every alternative uniformly β times and count the target a one additional time.

The theorem below turns this combinatorial condition into the characterization. The proof is a lower-bound argument. WIR gives coordinate lower bounds for alternatives that can be tested directly. For a higher-rank target a, the separator certificate combines WIR inequalities from lower-rank test peaks and isolates the coordinate a. Once every coordinate is bounded below by the total endowment of agents who peak there, feasibility makes all lower bounds bind.

THEOREM 4 (Layered separator theorem). *Suppose, for every* $i \in N$, $\mathcal{D}_i$ *is layered separator-rich. If the mechanism* M *satisfies WIR, then, for every profile* $\succ_N \in \prod_{i\in N} \mathcal{D}_i$,

$$M_N(\succ_N) = \sum_{i\in N} w_i \mathbf{1}_{p_i}.$$

Conversely, the peak-proportional rule satisfies coalitional individual rationality on $\prod_{i\in N} \mathcal{D}_i$.

The proof of Theorem 4 is postponed to Section 7.

We remark that layered separator richness is not monotone under domain inclusion. Adding a preference may create a new peakable alternative without adding the contour-set tests needed to identify its coordinate; hence, an enlarged domain may fail the condition even when a subdomain satisfies it. The next example illustrates this point and shows why WIR does not imply peak-proportionality on arbitrary restricted domains: on the enlarged domain, a WIR mechanism can have a non-peak-proportional grand coalition outcome.

EXAMPLE 5. Let $X = \{a, b, c, d\}$. Consider the following three strict orders:

$$\succ^1:\quad a \succ c \succ d \succ b, \qquad \succ^2:\quad b \succ c \succ d \succ a, \qquad \succ^3:\quad c \succ a \succ b \succ d.$$

Define two domains $\mathcal{D} = \{\succ^1, \succ^2\}$ and $\mathcal{D}' = \{\succ^1, \succ^2, \succ^3\}$.

First, $\mathcal{D}$ is layered separator-rich. Indeed, $P(\mathcal{D}) = \{a, b\}$. The alternative a is ranked last under $\succ^2$, and the alternative b is ranked last under $\succ^1$. Hence, $\mathcal{D}$ satisfies the rank-zero condition with $\rho(a) = \rho(b) = 0$. By Theorem 4, WIR uniquely determines the peak-proportional outcome in this case.

We now show that WIR does not uniquely determine the peak-proportional outcome on $\mathcal{D}'$. Let there be two agents, $N = \{1, 2\}$, with equal weights $w_1 = w_2 = 1$, and suppose both agents have domain $\mathcal{D}'$. We now construct a different outcome that also satisfies the WIR condition.

Define the singleton rules as follows. For every $\{i\}$,

$$M_{\{i\}}(\succ^1) = 1_a, \qquad M_{\{i\}}(\succ^2) = 1_b, \qquad M_{\{i\}}(\succ^3) = 1_d.$$

Define the grand-coalition rule M_N symmetrically ($M_N(\succ, \succ') = M_N(\succ', \succ)$) by

$$M_N(\succ^1, \succ^1) = 2\mathbf{1}_a, \qquad M_N(\succ^2, \succ^2) = 2\mathbf{1}_b, \qquad M_N(\succ^1, \succ^2) = \mathbf{1}_a + \mathbf{1}_b,$$
$$M_N(\succ^1, \succ^3) = \mathbf{1}_a + \mathbf{1}_c, \qquad M_N(\succ^2, \succ^3) = \mathbf{1}_b + \mathbf{1}_c, \qquad M_N(\succ^3, \succ^3) = \mathbf{1}_c + \mathbf{1}_d.$$

This rule is not peak-proportional. At the profile $(\succ^3, \succ^3)$, both agents have peak c, so the peak-proportional allocation is $2\mathbf{1}_c$, whereas the rule selects $\mathbf{1}_c + \mathbf{1}_d$.

It remains to verify WIR. All comparisons are immediate. At profiles $(\succ^1, \succ^1)$, $(\succ^2, \succ^2)$, and $(\succ^1, \succ^2)$, the peak-withdrawal allocation coincides with the grand allocation. At $(\succ^1, \succ^3)$, the only nontrivial comparison is for the $\succ^1$-agent:

$$\mathbf{1}_a + \mathbf{1}_c \succeq^{sd}_{\succ^1} \mathbf{1}_a + \mathbf{1}_d,$$

because $c \succ^1 d$. The profile $(\succ^2, \succ^3)$ is analogous, since $c \succ^2 d$. Finally, at $(\succ^3, \succ^3)$, peak withdrawal gives $\mathbf{1}_c + M_{\{i\}}(\succ^3) = \mathbf{1}_c + \mathbf{1}_d$, which is exactly the grand allocation. Therefore M satisfies WIR on $\mathcal{D}'$, but it is not peak-proportional at $(\succ^3, \succ^3)$. ♢

6. SEPARATOR-RICH EXAMPLES

This section verifies layered separator richness for the domains used in the main applications. We first consider graph-based domains, where alternatives are vertices of a finite graph. We then give a non-graph example based on business-architecture partitions.

6.1 *Graph Domains*

DEFINITION 4 (Connected sets in a graph). Let $G=(X,E)$ be an undirected graph. For $C\subseteq X$, let $G[C]$ denote the subgraph induced by C, and let $\kappa_G(C)$ denote the number of connected components of $G[C]$, with the convention $\kappa_G(\varnothing)=0$. A nonempty set $C\subseteq X$ is connected in G if $G[C]$ is connected; equivalently, any two vertices in C can be joined by a path whose vertices all lie in C.

Throughout this subsection, $G=(X,E)$ is finite, and we allow it to have at most k connected components, where k will be specified later.

Connectedness can be imposed on preferences in two natural orientations. The upper orientation restricts upper contour sets: alternatives that remain acceptable above a cutoff must form at most k connected regions. This is the natural formulation for goods, services, facilities, or investments, where acceptable alternatives expand from favored sites or projects. The lower orientation restricts lower contour sets: alternatives that are sufficiently bad below a cutoff must form at most k connected regions. This is the natural formulation for chores, congestion, risk, or other exposure-type burdens.

6.1.1 *Upper k-peaked Graph Domains*

The upper k-peaked domain consists of all strict preferences whose upper contour sets have at most k connected components.

DEFINITION 5 (Upper k-peaked graph domain). Fix an integer $k\geq 1$. The upper k-peaked graph domain on G is

$$\mathcal{D}^{\uparrow,k}(G)=\{\succ\in\mathcal{A}(X):\kappa_G(U_\succ(x))\leq k\text{ for every }x\in X\},$$

where $\mathcal{A}(X)$ is the set of all strict orders on X.

This domain permits an agent to have several locally attractive regions, but rules out rankings in which acceptable alternatives are scattered across more than k unrelated parts of the network. The key difficulty, relative to direct testability, is that some vertices need not be rankable last. For example, let G be the path $a - b - c$ and take $k = 1$. If the middle vertex b were ranked last, then the upper contour set of the alternative ranked immediately above b would be $\{a, c\}$, which has two connected components. This violates upper 1-peakedness. Thus the cut vertex b cannot be tested directly by placing it at the bottom of an admissible order. More generally, direct testing fails for a vertex a whenever $\kappa_G(X \setminus \{a\}) > k$. The main claim is that this obstruction can nevertheless be handled by separator tests: whenever $\kappa_G(X) \leq k$, the upper k-peaked graph domain is layered separator-rich. In fact, separator depth one is enough.

LEMMA 6 (Upper k-peaked graph domains). *Let $G = (X, E)$ be a finite graph with $\kappa_G(X) \leq k$, and let $k \geq 1$. Then the upper k-peaked graph domain $\mathcal{D}^{\uparrow,k}(G)$ is layered separator-rich.*

The proof of Lemma 6 relies on the following three graph lemmas. The first lemma is a realization result: any connected set can be made an upper contour set of an admissible order, with any prescribed vertex in that set as the peak.

LEMMA 7. *Let $G = (X, E)$ be a finite graph with $\kappa_G(X) \leq k$. Let $C \subseteq X$ be nonempty and connected in G, and let $z \in C$. There exists an order τ on X such that:*

- *the peak of τ is z;*
- *the first $|C|$ alternatives in τ are exactly the vertices of C;*
- *every prefix of τ has at most k connected components in G.*

Consequently, $\tau \in D^{\uparrow,k}(G)$, and C is an upper contour set of τ.

The second lemma gives the direct rank-zero tests: if deleting a vertex leaves at most k connected components, then that vertex can be ranked last by an order in $D^{\uparrow,k}(G)$.

LEMMA 8. *Let $G = (X, E)$ be a finite graph with $\kappa_G(X) \le k$. Fix $a \in X$. If $\kappa_G(X \setminus \{a\}) \le k$, then there is an order $\tau \in D^{\uparrow,k}(G)$ that ranks a last.*

The third lemma handles the remaining case. If deleting a vertex creates more than k connected components, then each component of the deletion graph contains a vertex that is directly testable. These directly testable vertices will serve as the lower-rank peaks in the separator certificate for the original vertex.

LEMMA 9. *Let $G = (X, E)$ be a finite graph with $\kappa_G(X) \le k$. Fix $a \in X$ and suppose that $\kappa_G(X \setminus \{a\}) > k$. If C is a connected component of $G[X \setminus \{a\}]$, then there exists $\lambda \in C$ such that $\kappa_G(X \setminus \{\lambda\}) \le k$.*

The proofs of these three lemmas are deferred to Appendix A.

PROOF OF LEMMA 6. Let $D = D^{\uparrow,k}(G)$. We prove that D is layered separator-rich with separator depth at most one.

First, every alternative in X is peakable. Fix $a \in X$. Since the singleton $\{a\}$ is connected, Lemma 7, applied to $C = \{a\}$ and $z = a$, gives an order in $D^{\uparrow,k}(G)$ whose peak is a. Hence $P(D^{\uparrow,k}(G)) = X$.

For each $a \in X$, define

$$\rho(a) = \begin{cases} 0, & \kappa_G(X \setminus \{a\}) \le k, \\ 1, & \kappa_G(X \setminus \{a\}) > k. \end{cases}$$

This is a rank function from $P(D) = X$ into $\{0, 1\}$. We verify the two conditions in the definition of layered separator richness.

First consider (L0). Let $a \in X$ satisfy $\rho(a) = 0$. Then $\kappa_G(X \setminus \{a\}) \le k$. By Lemma 8, there exists an order $\tau \in D^{\uparrow,k}(G)$ that ranks a last. This proves (L0).

Now consider (L1). Let $a \in X$ satisfy $\rho(a) = 1$. Then $\kappa_G(X \setminus \{a\}) > k$. Let $C_1, \dots, C_d$ be the connected components of $G[X \setminus \{a\}]$. Thus $d = \kappa_G(X \setminus \{a\}) > k$.

For each $h \in \{1, \dots, d\}$, Lemma 9 gives a vertex $\lambda_h \in C_h$ such that $\kappa_G(X \setminus \{\lambda_h\}) \leq k$. Therefore $\rho(\lambda_h) = 0 < 1 = \rho(a)$.

Since C_h is connected, Lemma 7, applied to $C = C_h$ and $z = \lambda_h$, gives an order $\tau_h \in D^{\uparrow,k}(G)$ whose peak is λ_h and whose first $|C_h|$ alternatives are exactly the vertices of C_h. Hence C_h is an upper contour set of τ_h. Moreover,

$$p(\tau_h) = \lambda_h \in C_h, \qquad a \notin C_h, \qquad \rho(p(\tau_h)) = \rho(\lambda_h) < \rho(a).$$

Set $H_h = X \setminus C_h$, $\alpha_h = 1$ for every h and $\beta = d - 1$. We claim that

$$\sum_{h=1}^{d} \mathbf{1}_{H_h} = (d-1)\mathbf{1}_X + \mathbf{1}_{\{a\}}.$$

Indeed, if $x = a$, then $a \notin C_h$ for every h, so $a \in H_h$ for every h. Hence the left-hand side equals d, while the right-hand side also equals $(d-1) + 1 = d$. If $x \neq a$, then x belongs to exactly one component C_h of $G[X \setminus \{a\}]$. Therefore x is excluded from exactly one complement H_h and belongs to the remaining $d - 1$ complements. Hence the left-hand side equals $d - 1$, which is also the value of the right-hand side at x.

Thus the orders τ_h, the upper contour sets C_h, the coefficients $\alpha_h = 1$, and the scalar $\beta = d - 1$ satisfy all requirements of (L1).

We have verified both (L0) and (L1) with ranks in $\{0, 1\}$. Therefore $D^{\uparrow,k}(G)$ is layered separator-rich, with separator depth at most one. □

6.1.2 *Lower k-dipped Graph Domains* This orientation is appropriate when the relevant network object is the set of bad, costly, or exposed alternatives. In a waste, congestion, or risk allocation interpretation, alternatives below a cutoff form a bad exposure region that may have up to k connected components. As the cutoff moves upward, these exposure components may expand through adjacent vertices and may also merge with one another. Thus, the parameter k allows several separated bad exposure regions while still ruling out arbitrarily fragmented lower contour sets.

DEFINITION 6 (Lower k-dipped graph domain). Fix an integer $k \geq 1$. The lower k-dipped graph domain on G is

$$\mathcal{D}^{\downarrow,k}(G) = \{\succ \in \mathcal{A}(X) : \kappa_G(L_{\succ}(x)) \leq k \text{ for every } x \in X\},$$

where $\mathcal{A}(X)$ is the set of all strict orders on X.

The key mathematical difference from the upper k-peaked case is that the lower-contour restriction does not obstruct the direct-test argument. When $\kappa_G(X) \leq k$, for any vertex a, one can construct an admissible order that ranks a last and whose every lower contour set has at most k connected components. Hence every peakable alternative is directly testable, and no positive separator layer is needed.

LEMMA 10 (Lower k-dipped graph domain). *Let $G = (X, E)$ be a finite graph with $\kappa_G(X) \leq k$, and let $k \geq 1$. Then the lower k-dipped graph domain $\mathcal{D}^{\downarrow,k}(G)$ is layered separator-rich. In fact, separator depth zero is sufficient.*

PROOF OF LEMMA 10. Set $\rho(a) = 0$ for every $a \in P(\mathcal{D}^{\downarrow,k}(G))$. It remains only to verify condition (L0).

Fix $a \in P(\mathcal{D}^{\downarrow,k}(G))$. Let $X_1, \ldots, X_m$ be the connected components of G and $m = \kappa_G(X) \leq k$. Relabel the components so that $a \in X_1$.

We first list the vertices of X_1 starting from a. Since $G[X_1]$ is connected, we can write $a = x_1, x_2, \ldots, x_{|X_1|}$ so that every prefix $\{x_1, \ldots, x_t\}$, $t \leq |X_1|$, is connected. Indeed, whenever the current prefix is not all of X_1, connectedness of $G[X_1]$ gives a vertex outside the prefix adjacent to it, and we append such a vertex.

For each remaining component X_ℓ, $\ell = 2, \ldots, m$, choose an ordering of its vertices $x_{\ell,1}, x_{\ell,2}, \ldots, x_{\ell,|X_\ell|}$ such that every prefix inside X_ℓ is connected. Concatenate these lists after the list of X_1. Thus we obtain a listing of all vertices of X whose first vertex is a, and whose every prefix has at most $m \leq k$ connected components in G.

Write this full list as $a = y_1, y_2, \ldots, y_{|X|}$. Now define the strict order $y_{|X|} \succ y_{|X|-1} \succ \cdots \succ y_1 = a$. The alternative a is ranked last. Moreover, every lower contour set of this order is one of the prefixes $\{y_1, \ldots, y_t\}$ for $t = 1, \ldots, |X|$. By construction, each such prefix has at most k connected components. Hence the order belongs to $\mathcal{D}^{\downarrow,k}(G)$ and ranks a last.

Therefore every peakable alternative satisfies (L0). Since all alternatives in $P(\mathcal{D}^{\downarrow,k}(G))$ have rank zero, condition (L1) is vacuous. Hence $\mathcal{D}^{\downarrow,k}(G)$ is layered separator-rich with separator depth zero. $\square$

6.2 *A non-graph example: business architectures*

This example gives a non-graph domain of exact separator depth three. Consider an investment partnership that pools divisible capital and invests across firms with different business architectures.

Let $F = \{a, b, c, d, e\}$ be five business functions, and let $X = \Pi(F)$ be the set of partitions of F. A partition represents a business architecture: functions in the same block are integrated within a common business unit, operating platform, or revenue-responsibility structure; while functions in different blocks are linked through commercial interfaces, partnerships, service agreements, or organizational handoffs. For instance, $ab \mid c \mid d \mid e$ integrates a and b, whereas $abcde$ is fully integrated.

Partitions are ordered by refinement: $\pi \sqsubseteq \sigma$ if σ is obtained from π by merging blocks. Thus larger elements are more integrated. Let $b(\pi)$ be the number of blocks of π, and define

$$r(\pi) = |F| - b(\pi).$$

Hence the fully separated architecture $\hat{0} = a \mid b \mid c \mid d \mid e$ has rank zero, while the fully integrated architecture $\hat{1} = abcde$ has rank four.

An allocation $q \in \mathbb{R}_+^X$ assigns capital across architectures, with q_π denoting investment in firms or projects with architecture π.

Each investor has an existing architecture $P \in \Pi(F)$. Matching this architecture yields an exact-integration surplus, so P is her peak. Away from this exact match, the investor prefers greater integration, since it reduces interfaces, coordination failures, and execution risk. Formally, let $\mathcal{D}^{BA}$ be the set of strict orders $\succ$ on X

for which there exists $P \in \Pi(F)$ such that $p(\succ) = P$ and, for all $\pi, \sigma \in X$,

$$\pi \sqsubset \sigma, \quad \pi \neq P \quad \Longrightarrow \quad \sigma \succ \pi.$$

LEMMA 11 (Business architectures have exact separator depth three). *The domain $\mathcal{D}^{BA}$ is layered separator-rich, and the least integer R in Definition 3 is three.*

PROOF SKETCH. Every architecture is peakable. Assign rank $\rho(\pi) = \max\{r(\pi) - 1, 0\}$. Architectures of lattice rank zero or one can be ranked last, giving the direct tests. For an architecture σ with $r(\sigma) \geq 2$, split $X \setminus \{\sigma\}$ into its strict refinements $D_\sigma = \{\pi : \pi \sqsubset \sigma\}$ and the alternatives $O_\sigma = \{\pi : \pi \not\sqsubseteq \sigma\}$ that are not below σ. Singleton tests identify the strict refinements. The set O_σ is upward closed under refinement, so one rank-one architecture outside σ generates an admissible upper contour set O_σ. These cells partition $X \setminus \{\sigma\}$, so their complements satisfy the separator identity. Exact depth three follows from the chain $ab \mid c \mid d \mid e \sqsubset abc \mid d \mid e \sqsubset abcd \mid e \sqsubset abcde$. The full proof is in Appendix B. □

7. PROOF OF THE LAYERED SEPARATOR THEOREM

This section proves our main theorem used above.

The proof has two steps. First, if all agents who peak at a can rank a last in their domains, WIR gives the coordinate lower bound $M_N(\succ_N)_a \geq \sum_{i:p_i=a} w_i$. This is the direct-test lemma. Second, if a cannot be ranked last, the separator identity combines WIR inequalities associated with lower-rank peaks. The common $\beta\mathbf{1}_X$ term cancels by feasibility, leaving the same lower bound for coordinate a. Induction on the total separator rank gives the lower bound for every coordinate. Since total mass is fixed, all lower bounds bind.

PROOF OF THEOREM 4. For a grand profile $\succ_N$, define its heterogeneous total rank by $R(\succ_N) = \sum_{i\in N} \rho_i(p_i)$. This is a nonnegative integer. We prove necessity by induction on $R(\succ_N)$ that $M_N(\succ_N) = \sum_{i\in N} w_i \mathbf{1}_{p_i}$. It is enough to prove the following coordinate lower bound for every $a \in X$:

$$M_N(\succ_N)_a \geq \sum_{j:p_j=a} w_j. \tag{3}$$

Indeed, summing (3) over all $a \in X$ gives

$$\sum_{a \in X} M_N(\succ_N)_a \geq \sum_{a \in X} \sum_{j : p_j = a} w_j = \sum_{j \in N} w_j = w_N.$$

Feasibility gives $M_N(\succ_N)(X) = w_N$. Therefore, every coordinate lower bound in (3) must bind, and the allocation must be peak-proportional.

Base case Suppose $R(\succ_N) = 0$. Then every agent j satisfies $\rho_j(p_j) = 0$. Fix $a \in X$. If no agent peaks at a, then (3) is trivial. Otherwise, for every agent j with $p_j = a$, we have $\rho_j(a) = 0$. By layered separator richness of $\mathcal{D}_j$, each such agent has an admissible order in $\mathcal{D}_j$ that ranks a last. Lemma 2 applies and yields $M_N(\succ_N)_a \geq \sum_{j : p_j = a} w_j$. Thus (3) holds for every a, and feasibility implies the peak-proportional formula.

Induction step Assume the result is true for every profile whose total rank is strictly smaller than $R(\succ_N)$, and consider a profile $\succ_N$ with positive total rank. Fix $a \in X$. We prove (3).

If no agent peaks at a, then the bound is trivial. Hence, suppose that at least one agent peaks at a.

First consider the case in which every agent peaking at a has rank zero; that is $\rho_j(a) = 0$ for every j with $p_j = a$. Then every such agent has an admissible order in her own domain that ranks a last. Lemma 2 gives (3).

It remains to consider the case in which some agent peaking at a has positive rank. Choose such an agent i, so that $p_i = a$ and $\rho_i(a) > 0$. By layered separator richness of $\mathcal{D}_i$, there is a separator certificate in agent i's own domain. Thus there exists orders $\tau_h \in \mathcal{D}_i$, upper contour sets C_h of τ_h, coefficients $\alpha_h > 0$, and a scalar $\beta \geq 0$, such that, with $H_h = X \setminus C_h$,

$$p(\tau_h) \in C_h, \qquad a \notin C_h, \qquad \rho_i(p(\tau_h)) < \rho_i(a) \quad \text{for every } h, \tag{4}$$

and

$$\sum_h \alpha_h \mathbf{1}_{H_h} = \beta \mathbf{1}_X + \mathbf{1}_{\{a\}}. \tag{5}$$

For each h, define the modified grand profile in which only agent i's report is changed to τ_h. By (4), $\rho_i(p(\tau_h)) < \rho_i(a)$, while all other agents' reports are unchanged. Hence, the modified grand profile has a strictly smaller total rank than the original grand profile, i.e., $R(\tau_h, \succ_{-i}) < R(\succ_N)$. By the induction hypothesis, the modified grand allocation Q^h is peak-proportional:

$$M_N(\tau_h, \succ_{-i}) = w_i \mathbf{1}_{p(\tau_h)} + \sum_{j \neq i} w_j \mathbf{1}_{p_j}. \tag{6}$$

Since $p(\tau_h) \in C_h$, we have $p(\tau_h) \notin H_h$. Evaluating (6) on H_h gives

$$\sum_{\ell \in H_h} M_N(\tau_h, \succ_{-i})_\ell = \sum_{j \neq i : p_j \in H_h} w_j. \tag{7}$$

Now apply WIR at the modified profile $(\tau_h, \succ_{-i})$ to agent i, with withdrawal to $p(\tau_h)$. Since C_h is an upper contour set of τ_h and contains $p(\tau_h)$, WIR gives

$$\sum_{\ell \in C_h} M_N(\tau_h, \succ_{-i})_\ell \geq w_i + \sum_{\ell \in C_h} M_{N \setminus \{i\}}(\succ_{-i})_\ell.$$

Using feasibility and passing to complements, we obtain

$$\sum_{\ell \in H_h} M_{N \setminus \{i\}}(\succ_{-i})_\ell \geq \sum_{\ell \in H_h} M_N(\tau_h, \succ_{-i})_\ell = \sum_{j \neq i : p_j \in H_h} w_j \qquad \text{for every } h. \tag{8}$$

Multiply (8) by α_h and sum over h. The left-hand side becomes

$$\begin{aligned}
\sum_h \alpha_h \sum_{\ell \in H_h} M_{N \setminus \{i\}}(\succ_{-i})_\ell &= \sum_{x \in X} \left(\sum_h \alpha_h \mathbf{1}_{H_h}(x) \right) M_{N \setminus \{i\}}(\succ_{-i})_x \\
&\overset{(*)}{=} \sum_{x \in X} \big(\beta \mathbf{1}_X(x) + \mathbf{1}_{\{a\}}(x) \big) M_{N \setminus \{i\}}(\succ_{-i})_x \\
&= \beta \sum_{\ell \in X} M_{N \setminus \{i\}}(\succ_{-i})_\ell + M_{N \setminus \{i\}}(\succ_{-i})_a \\
&\overset{(\dagger)}{=} \beta \sum_{j \neq i} w_j + M_{N \setminus \{i\}}(\succ_{-i})_a,
\end{aligned} \tag{9}$$

where $(*)$ follows from the separator identity (5) and $(\dagger)$ uses feasibility for the residual economy $N \setminus \{i\}$.

Similarly, applying (5) at each residual peak p_j, the right-hand side becomes

$$\begin{aligned}\sum_h \alpha_h \sum_{j\neq i:p_j\in H_h} w_j &= \sum_{j\neq i} w_j \sum_h \alpha_h \mathbf{1}_{H_h}(p_j)\\ &= \sum_{j\neq i}\big(\beta+\mathbf{1}_{\{a\}}(p_j)\big)w_j\\ &= \beta\sum_{j\neq i} w_j + \sum_{j\neq i:p_j=a} w_j. \qquad (10)\end{aligned}$$

Combining (8), (9), and (10), and then canceling the common term $\beta\sum_{j\neq i} w_j$, gives

$$M_{N\setminus\{i\}}(\succ_{-i})_a \geq \sum_{j\neq i:p_j=a} w_j. \qquad (11)$$

Finally return to the original profile $\succ_N$. Since $p_i = a$, the singleton $\{a\}$ is an upper contour set of $\succ_i$. WIR for agent i, with withdrawal to a, gives $M_N(\succ_N)_a \geq w_i + M_{N\setminus\{i\}}(\succ_{-i})_a$. Using (11), we have

$$M_N(\succ_N)_a \geq w_i + \sum_{j\neq i:p_j=a} w_j = \sum_{j:p_j=a} w_j.$$

This proves the coordinate lower bound (3) in the positive-rank case as well.

Since the coordinate lower bound holds for every $a\in X$, feasibility implies that all these bounds bind. Hence $M_N(\succ_N)_a = \sum_{j:p_j=a} w_j$ for every $a\in X$ or equivalently, $M_N(\succ_N) = \sum_{j\in N} w_j\mathbf{1}_{p_j}$. This completes the induction and proves the necessity. □

8. CARDINAL BOUNDARY RESULTS

The preceding analysis is deliberately ordinal. Participation is evaluated by stochastic dominance, so the comparison is robust to all cardinal representations of the reported ranking. This is what makes WIR so sharp: in the ordinal model, a withdrawing agent's strongest pure withdrawal is to place her endowment at her peak, and no cardinal tradeoff among alternatives has to be specified.

Cardinal utilities change the participation problem. A withdrawing agent no longer merely chooses a favorite alternative. She chooses the best allocation of her own endowment given the residual allocation selected by the mechanism. The outside option is therefore shaped by the curvature of utility over aggregate allocations. This section shows that the ordinal characterization is not a generic participation theorem. It is a boundary result: once cardinal curvature matters, WIR may be too weak to identify the peak-proportional allocation, or too strong to be jointly feasible with incentive compatibility.

We consider three benchmark cardinal environments. With linear utilities, withdrawal remains separable across alternatives. WIR is then close to its ordinal counterpart, but it is no longer sufficient: residual coalitions may exploit Pareto improvements that move the grand allocation away from peak-proportionality. Adding strategy-proofness eliminates this residual freedom and restores the peak-proportional conclusion on profiles with unique maximizers. With concave CES utilities, the difficulty is different. Withdrawal creates a completion value: an agent can use her own endowment to complement the residual allocation. Even in a binary equal-weight economy, the resulting outside options can be mutually infeasible. With convex L_p-norm utilities, withdrawal rewards concentration. In the binary equal-weight economy, this force is incompatible with strategy-proofness.

Cardinal participation A cardinal domain is a set $\mathcal{U}$ of utility functions on $\mathbb{R}_+^X$. Each agent $i \in N$ has a private utility function $u_i : \mathbb{R}_+^X \to \mathbb{R}$. A cardinal mechanism is a family $M = \{M_S\}_{S \subseteq N}$, where M_S maps the reported utility profile $u_S \triangleq \{u_i(\cdot)\}_{i \in S}$ to an allocation $\mathbf{q}$ over X. If agent $i \in S$ withdraws from coalition S, she chooses an allocation $\mathbf{d} \in \mathbb{R}_+^X$ of her own endowment and the residual coalition receives $M_{S\setminus\{i\}}(u_{S\setminus\{i\}})$. Cardinal individual rationality at S requires

$$u_i(M_S(u_S)) \geq \max_{\mathbf{d} \in \mathbb{R}_+^X : \sum_{a \in X} d_a = w_i} u_i\big(\mathbf{d} + M_{S\setminus\{i\}}(u_{S\setminus\{i\}})\big)$$

for every $i \in S$ and every profile u_S. In the cardinal section, *WIR* again means that this individual-rationality requirement is imposed at $S = N$.

We say the grand-coalition rule M_N is strategy-proof if

$$u_i(M_N(u_i, u_{-i})) \geq u_i(M_N(u_i', u_{-i}))$$

for every agent i, true utility u_i, report u_i', and profile u_{-i}, where $u_{-i} = u_{N\setminus\{i\}}$.

8.1 *Linear utilities*

Let $u_i(\mathbf{q}) = \sum_{x\in X} v_{ix}q_x$, where $\mathbf{v}_i \in \mathbb{R}^X_{\geq 0}$. Under linear utilities, an agent who withdraws optimally allocates her endowment among her utility-maximizing alternatives. This keeps the participation constraint close to its ordinal counterpart. Nevertheless, WIR alone leaves room for non-peak-proportional outcomes.

PROPOSITION 12. *On the linear utility domain with three agents and three alternatives, there exists a mechanism satisfying WIR whose grand coalition outcome is not peak-proportional and, at some profile, gives an agent strictly lower utility than the peak-proportional benchmark.*

The detailed construction is provided in Appendix C.1. The next result shows that strategy-proofness removes precisely this freedom. The proof draws on techniques from Hylland (1980).

THEOREM 13. *If* $|X| \geq 3$*, utilities are linear and the mechanism* M_N *satisfies WIR and strategy-proofness. There exist vectors* $\boldsymbol{\gamma}_i$ *for each agent* i*, such that* $\boldsymbol{\gamma}_i \geq 0$*,* $\sum_{x\in X} \boldsymbol{\gamma}_{ix} = 1$*, and* $\operatorname{supp}(\gamma_i) \subseteq \arg\max_{x\in X} v_{ix}$*, satisfying*

$$M_N(u_N) = \sum_{i\in N} w_i \boldsymbol{\gamma}_i.$$

Additionally, with two alternatives, $X = \{x, y\}$*, WIR alone gives the peak-proportional outcome,*

$$M_N(u_N)_x \geq \sum_{i:v_{ix}>v_{iy}} w_i, \qquad M_N(u_N)_y \geq \sum_{i:v_{iy}>v_{ix}} w_i.$$

The full proof is provided in Appendix C.1.

REMARK 3. When $|X| \geq 3$, we note that the representation $M_N(u_N) = \sum_{i\in N} w_i \boldsymbol{\gamma}_i$ does not by itself guarantee strategy-proofness, because the vector $\boldsymbol{\gamma}_i$ may depend implicitly on the reports of other agents, $u_{N\setminus\{i\}}$. The mechanism is strategy-proof if each $\boldsymbol{\gamma}_i$ depends only on u_i.

As a special case of Theorem 13, when $|X| \geq 3$ and M satisfies WIR and strategy-proofness, we have, for every utility profile in which each agent i has a unique optimal alternative p_i, $M_N(u_N) = \sum_{i\in N} w_i \mathbf{1}_{p_i}$.

8.2 *Concave CES utilities*

We next consider concave CES utilities. Fix $\rho < 1$, $\rho \neq 0$. The CES utility function is

$$u_i(\mathbf{q}) = \left(\sum_{a\in X} v_{ia} q_a^{\rho}\right)^{1/\rho} \qquad \mathbf{v}_i \in \mathbb{R}^X_{++}.$$

For $\rho < 0$, the utility is interpreted by the standard continuous convention: if a coordinate with a positive coefficient is zero, the utility is zero.

CES curvature changes the participation constraint in a more fundamental way. A withdrawing agent can use her endowment to complete the residual allocation. As a result, the outside option is not simply the value of receiving one unit at a favorite alternative. It depends on what the residual coalition leaves behind. This can make the participation requirements for different agents mutually incompatible.

THEOREM 14. *Let $N = \{1, \dots, n\}$, $n \geq 2$, let $X = \{x, y\}$, and let $w_i = 1$ for every $i \in N$. Fix $\rho < 1$, $\rho \neq 0$. Suppose each agent has a positive binary CES utility function*

$$u_i(\mathbf{q}) = \left(v_{ix} q_x^{\rho} + v_{iy} q_y^{\rho}\right)^{1/\rho}, \qquad v_{ix}, v_{iy} > 0.$$

Then no mechanism satisfies WIR.

The full proof is deferred to Appendix C.2.

8.3 L_p*-norm utilities*

Finally, consider the convex L_p-norm family, with $p > 1$:

$$u_i(\mathbf{q}) = \left(\sum_{a\in X} (v_{ia} q_a)^p\right)^{1/p}, \qquad p > 1,$$

with coefficient vectors $\mathbf{v}_i \in \mathbb{R}^X_{\geq 0} \setminus \{\mathbf{0}\}$.

This domain has the opposite geometry from the concave CES case. Convexity makes concentration valuable. A withdrawing agent benefits from placing her own endowment where it reinforces an already favorable residual allocation. The next result shows that, even in the binary equal-weight economy, this concentration force is incompatible with combining WIR and strategy-proofness.

THEOREM 15. *Let $N = \{1, \ldots, n\}$ with $n \geq 2$, let $X = \{x, y\}$, and let $w_i = 1$ for every i. Fix $p > 1$. Suppose each agent has a binary L_p-norm utility function*

$$u_i(\mathbf{q}) = \left((v_{ix}q_x)^p + (v_{iy}q_y)^p\right)^{1/p},$$

no mechanism satisfies both WIR and strategy-proofness.

A detailed proof is provided in Appendix C.3.

The three cardinal cases, therefore, locate the ordinal theorem precisely. Linear utilities show that WIR alone need not identify peak-proportionality once residual coalitions can exploit cardinal Pareto improvements. Concave CES utilities show that participation itself may be infeasible when withdrawal values depend on completing the residual allocation. Convex L_p-norm utilities show that participation and strategy-proofness can conflict when withdrawal rewards are concentrated. The ordinal characterization is therefore not a consequence of participation in the abstract. It relies on the stochastic-dominance formulation, which removes cardinal curvature from the outside option and turns voluntary participation into a coordinate-identification condition.

9. CONCLUSION

This paper studies aggregate allocation when exit preserves property rights: a withdrawing agent redeploys her own endowment, while the remaining agents continue under the same mechanism family. In the ordinal model, this outside option gives weak individual rationality (WIR) strong identifying power. On the unrestricted domain, and more generally on layered separator-rich domains, WIR alone implies

$$M_N(\succ_N) = \sum_{i \in N} w_i \mathbf{1}_{p_i}.$$

This characterization requires neither efficiency nor strategy-proofness and imposes no participation condition on proper subcoalitions. Conversely, the peak-proportional family satisfies individual rationality at every coalition.

The proof reveals the domain structure behind the result. Direct preference tests identify some coordinates, while separator certificates combine contour-set inequalities to identify the rest. This method covers upper and lower graph-restricted domains as well as the non-graph business-architecture domain. The essential requirement is therefore not graph structure itself, but a finite hierarchy of admissible tests capable of separating every peakable alternative.

The cardinal results mark the boundary of the ordinal characterization. With linear utilities, WIR alone need not imply peak-proportionality, although strategy-proofness restores it at profiles with unique maximizers. With positive binary CES utilities and equal endowments, WIR is infeasible; with binary L_p-norm utilities, $p > 1$, WIR and strategy-proofness are incompatible. Thus, the ordinal theorem depends on the intensity-robustness of stochastic dominance: once utility curvature affects the value of withdrawal, participation may become either too weak to identify the benchmark or too demanding to coexist with incentives.

APPENDIX A: PROOFS FOR GRAPH DOMAINS

This appendix provides the full proof of Lemma 7, Lemma 8 and Lemma 9.

PROOF OF LEMMA 7. Let $X_1, \ldots, X_m$ be the connected components of G. By assumption, $m = \kappa_G(X) \le k$. Since C is connected, it is contained in one connected component of G. Relabel the components so that $C \subseteq X_1$.

We first construct an ordering of X_1. Choose a spanning tree of the connected graph $G[C]$ and root it at z. List the vertices of C as $z = x_1, x_2, \ldots, x_{|C|}$ so that every vertex other than z appears after its parent in the rooted tree. Then each prefix $\{x_1, \ldots, x_t\}$, $t \le |C|$, is connected.

If $C \neq X_1$, extend this list to all vertices of X_1 as follows. Suppose a connected set $S \subsetneq X_1$ has already been listed. Since $G[X_1]$ is connected, there is an edge between S and $X_1 \setminus S$. Append a vertex of $X_1 \setminus S$ adjacent to S. The enlarged listed set is again connected. Repeating this step gives a list of all vertices of X_1 whose first $|C|$ vertices are exactly C and whose every prefix is connected.

For each remaining component X_ℓ, $\ell = 2, \dots, m$, choose an arbitrary root $z_\ell \in X_\ell$. By the same rooted-tree argument, list the vertices of X_ℓ so that every prefix within X_ℓ is connected. Now concatenate the lists in the order $X_1, X_2, \dots, X_m$, and let τ be the strict order induced by this concatenated list, from first to last.

The first vertex is z, so z is the peak of τ. The first $|C|$ vertices are exactly the vertices of C. Moreover, every prefix of the concatenated list consists of a connected prefix of X_1, together with some complete components $X_2, \dots, X_{\ell-1}$, and possibly a connected prefix of X_ℓ. Hence every prefix has at most $m \le k$ connected components in G.

Since upper contour sets of τ are exactly prefixes of this list, every upper contour set has at most k connected components. Therefore $\tau \in D^{\uparrow,k}(G)$. Finally, if $x_{|C|}$ is the last vertex among the first $|C|$ listed vertices, then $U_\tau(x_{|C|}) = C$. Thus C is an upper contour set of τ. □

PROOF OF LEMMA 8. If $X = \{a\}$, the unique order ranks a last. Since $\kappa_G(X) = 1 \le k$, this order belongs to $D^{\uparrow,k}(G)$.

Assume $X \setminus \{a\} \neq \emptyset$. Let $C_1, \dots, C_d$ be the connected components of $G[X \setminus \{a\}]$. By assumption, $d = \kappa_G(X \setminus \{a\}) \le k$.

For each component C_ℓ, choose any vertex $z_\ell \in C_\ell$. Choose a spanning tree of $G[C_\ell]$ rooted at z_ℓ, and list the vertices of C_ℓ as $x_{\ell,1}, x_{\ell,2}, \dots, x_{\ell,|C_\ell|}$ so that every vertex other than z_ℓ appears after its parent in the rooted tree. Then every prefix $\{x_{\ell,1}, \dots, x_{\ell,t}\}$, $t = 1, \dots, |C_\ell|$, is connected inside C_ℓ.

Define the order τ by concatenating these component lists and putting a last:

$$x_{1,1} \succ \cdots \succ x_{1,|C_1|} \succ x_{2,1} \succ \cdots \succ x_{2,|C_2|} \succ \cdots \succ x_{d,1} \succ \cdots \succ x_{d,|C_d|} \succ a.$$

Consider any upper contour set of τ. If it does not contain a, then it is a prefix of the concatenated list before a is added. Such a prefix consists of some complete components among $C_1, \dots, C_d$, possibly together with a connected prefix of the next component. Hence it has at most $d \le k$ connected components.

The only upper contour set that contains a is X itself, since a is ranked last. By assumption, $\kappa_G(X) \le k$. Therefore every upper contour set of τ has at most k connected components. Hence $\tau \in D^{\uparrow,k}(G)$, and by construction a is ranked last. □

PROOF OF LEMMA 9. Let $X_1, \dots, X_m$ be the connected components of G. By assumption, $m = \kappa_G(X) \le k$. Relabel the components so that $a \in X_1$.

Every connected component of $G[X \setminus \{a\}]$ is either a connected component of $G[X_1 \setminus \{a\}]$, or one of the original components $X_2, \dots, X_m$. We consider these two cases separately.

First suppose that C is a connected component of $G[X_1 \setminus \{a\}]$. Since $\kappa_G(X \setminus \{a\}) > k \ge m = \kappa_G(X)$, removing a increases the number of connected components. Hence a is a cut vertex of the connected graph $G[X_1]$. Choose $\lambda \in C$ maximizing the graph distance from a inside $G[X_1]$.

We claim that $G[X_1 \setminus \{\lambda\}]$ is connected. Suppose not. Then there is a connected component C' of $G[X_1 \setminus \{\lambda\}]$ that does not contain a. Choose $y \in C'$. Every path from a to y in $G[X_1]$ must pass through λ, so $d_{G[X_1]}(a, y) > d_{G[X_1]}(a, \lambda)$. Moreover, $y \in C$: since C' is a component of $G[X_1 \setminus \{\lambda\}]$ not containing a, and $G[X_1]$ is connected, some vertex of C' is adjacent to λ. Thus y is connected to λ in $G[X_1 \setminus \{a\}]$, and since $\lambda \in C$, we have $y \in C$. This contradicts the choice of λ as a vertex in C with maximal distance from a. Therefore $G[X_1 \setminus \{\lambda\}]$ is connected.

It follows that removing λ does not increase the number of connected components of G. The remaining components are $X_1 \setminus \{\lambda\},\ X_2, \dots, X_m$, so $\kappa_G(X \setminus \{\lambda\}) = m \le k$.

Now suppose that $C = X_\ell$ for some $\ell \ge 2$. If C is a singleton, let λ be its unique vertex. Then removing λ deletes this connected component, so $\kappa_G(X \setminus \{\lambda\}) = m - 1 \le k$. If C has at least two vertices, choose a spanning tree of $G[C]$ and let λ be a leaf of this tree. Then the tree with λ removed is still connected and spans $C \setminus \{\lambda\}$. Hence $G[C \setminus \{\lambda\}]$ is connected. Therefore removing λ does not increase the number of connected components of G, and again $\kappa_G(X \setminus \{\lambda\}) = m \le k$.

In both cases, we have found $\lambda \in C$ such that $\kappa_G(X \setminus \{\lambda\}) \leq k$. $\square$

APPENDIX B: PROOF FOR THE BUSINESS ARCHITECTURE DOMAIN

PROOF OF LEMMA 11. Every architecture is peakable, so $P(\mathcal{D}^{BA}) = X$. Define $\rho(\pi) = \max\{r(\pi) - 1, 0\}$. Thus ranks zero and one form layer zero, rank two layer one, rank three layer two, and rank four layer three.

We verify condition (L0). The fully separated architecture $\hat{0} = a \mid b \mid c \mid d \mid e$ can be ranked last by choosing an admissible order with peak $\hat{1} = abcde$. Now let λ be any rank-one architecture. Choose an admissible order with peak $\hat{0}$. The only strict predecessor of λ in the refinement order is $\hat{0}$, and the comparison from $\hat{0}$ to λ is exempt because $\hat{0}$ is the peak. All strict successors of λ, if any, are required to be above λ, which is consistent with placing λ last. Hence λ can be ranked last in some admissible order. This proves (L0) for every architecture with $\rho = 0$.

We next verify condition (L1). Fix $\sigma \in X$ with $r(\sigma) \geq 2$. Define $D_\sigma = \{\pi \in X : \pi \sqsubset \sigma\}$ and $O_\sigma = \{\pi \in X : \pi \not\sqsubseteq \sigma\}$. Then, $X \setminus \{\sigma\} = D_\sigma \sqcup O_\sigma$.

For each $\pi \in D_\sigma$, the singleton $\{\pi\}$ is an admissible upper contour set: choose an admissible order with peak π, and take the upper contour set at the peak. Moreover, $\rho(\pi) < \rho(\sigma)$, because $\pi \sqsubset \sigma$ implies $r(\pi) < r(\sigma)$.

It remains to handle O_σ. If $O_\sigma = \varnothing$, there is nothing to add. Suppose instead that $O_\sigma \neq \varnothing$. Since $r(\sigma) \geq 2$, and $\sigma \neq \hat{1}$ in this case, σ has at least two blocks. Choose two functions that lie in different blocks of σ, and let λ_σ be the rank-one architecture that merges exactly those two functions and leaves all other functions as singleton blocks. Then $\lambda_\sigma \in O_\sigma$ and $\rho(\lambda_\sigma) = 0 < \rho(\sigma)$. The first inclusion holds because λ_σ merges two functions that are separate in σ, so σ cannot be obtained from λ_σ by further merging blocks. Therefore there is an admissible order with peak λ_σ for which O_σ is an upper contour set.

Define $\Gamma_\sigma = \big\{\{\pi\} : \pi \sqsubset \sigma\big\} \cup \big\{O_\sigma : O_\sigma \neq \varnothing\big\}$. The sets in Γ_σ partition $X \setminus \{\sigma\}$. For each $C \in \Gamma_\sigma$, let $H_C = X \setminus C$. Therefore $\sum_{C \in \Gamma_\sigma} 1_{H_C} = \big(|\Gamma_\sigma| - 1\big) 1_X + 1_{\{\sigma\}}$. This is the separator identity required in (L1). Thus $\mathcal{D}^{BA}$ is layered separator-rich with separator depth at most three.

It remains to show that depth three is necessary. Let $\tilde{\rho}$ be any rank function satisfying Definition 3 for $\mathcal{D}^{BA}$. We first observe that every architecture σ with $r(\sigma) \geq 2$ must have positive $\tilde{\rho}$-rank. Such a σ has at least two distinct strict predecessors in the refinement order. In any admissible order with peak P, at least one of these predecessors, call it π, is different from P. Hence σ cannot be ranked last. By (L0), $\tilde{\rho}(\sigma) \neq 0$.

Now consider the chain

$$\pi_1 = ab \mid c \mid d \mid e \sqsubset \pi_2 = abc \mid d \mid e \sqsubset \pi_3 = abcd \mid e \sqsubset \pi_4 = abcde.$$

For $t = 2,3,4$, the architecture π_t has rank at least two, and therefore has positive $\tilde{\rho}$-rank. Take a separator certificate for π_t. Its separator identity has the form

$$\sum_h \alpha_h 1_{H_h} = \beta 1_X + 1_{\{\pi_t\}}, \qquad H_h = X \setminus C_h,$$

where each C_h is an admissible upper contour set, $\pi_t \notin C_h$, and $\tilde{\rho}(p(\tau_h)) < \tilde{\rho}(\pi_t)$. Evaluating the identity at π_t gives $\sum_h \alpha_h = \beta + 1$. Since $1_{H_h} = 1_X - 1_{C_h}$, the separator identity is equivalently $\sum_h \alpha_h 1_{C_h} = 1_{X \setminus \{\pi_t\}}$. Hence every alternative different from π_t, and in particular π_{t-1}, belongs to at least one certificate cell C_h.

Choose such a cell C_h with $\pi_{t-1} \in C_h$. Let $P_h = p(\tau_h)$ be the peak of the admissible order generating this upper contour set. We claim that $P_h = \pi_{t-1}$. Suppose not. Since $\pi_{t-1} \sqsubset \pi_t$ and $\pi_{t-1} \neq P_h$, admissibility of τ_h implies $\pi_t \succ_{\tau_h} \pi_{t-1}$. Because C_h is an upper contour set of τ_h and contains π_{t-1}, it must also contain π_t, contradicting $\pi_t \notin C_h$. Therefore $P_h = \pi_{t-1}$. Since the certificate uses only lower-rank peaks, $\tilde{\rho}(\pi_{t-1}) < \tilde{\rho}(\pi_t)$ for $t = 2,3,4$. Since ranks are nonnegative integers, $\tilde{\rho}(\pi_4) \geq 3$. Therefore every valid layered-separator representation has depth at least three. □

APPENDIX C: PROOFS FOR CARDINAL BOUNDARY RESULTS

This appendix collects the proofs of the cardinal results stated in Section 8.

C.1 *Proofs for linear utilities*

Proof of Proposition 12. Fix a deterministic tie-breaking order on X, and let p_i denote agent i's selected utility-maximizing alternative. Define a mechanism $\widehat{M}$ as follows. Singleton coalitions use the peak-proportional rule, as do the coalitions $\{1,2\}$ and $\{1,3\}$. Coalition $\{2,3\}$ uses the peak-proportional allocation $w_2\mathbf{1}_{p_2}+w_3\mathbf{1}_{p_3}$, unless assigning the entire residual endowment w_2+w_3 to some pure alternative gives both agents 2 and 3 strictly higher utility. If such alternatives exist, $\widehat{M}_{\{2,3\}}$ selects one according to the fixed tie-breaking order. Finally, set $\widehat{M}_N(u_N)=w_1\mathbf{1}_{p_1}+\widehat{M}_{\{2,3\}}(u_2,u_3)$.

Agent 1's WIR constraint holds with equality: if she withdraws, she takes w_1 to p_1, and the residual coalition $\{2,3\}$ receives exactly $\widehat{M}_{\{2,3\}}$. For agent 2, the outside allocation is $w_2\mathbf{1}_{p_2}+w_1\mathbf{1}_{p_1}+w_3\mathbf{1}_{p_3}$. The grand allocation keeps the term $w_1\mathbf{1}_{p_1}$ fixed and weakly improves the $\{2,3\}$ component for agent 2. Thus agent 2's WIR constraint holds. The same argument applies to agent 3.

Now take equal weights and alternatives a,b,c. Let the value vectors, in the order (a,b,c), be

$$\mathbf{v}_1=(1,0.9,0),\qquad \mathbf{v}_2=(0,1,0.9),\qquad \mathbf{v}_3=(1,0,0.9).$$

The peak-proportional grand allocation is $2\mathbf{1}_a+\mathbf{1}_b$, giving agent 1 utility 2.9. For coalition $\{2,3\}$, assigning both endowments to c gives each of agents 2 and 3 utility 1.8, strictly above the utility 1 they receive from the subcoalition peak-proportional allocation $\mathbf{1}_b+\mathbf{1}_a$. Therefore $\widehat{M}_N=\mathbf{1}_a+2\mathbf{1}_c$, and agent 1's utility is 1. □

Proof of Theorem 13. We first consider $|X|\geq 3$; we prove the result for profiles in which every agent has a unique optimal alternative $|\arg\max_x v_{ix}|=1$ for all i. The extension to profiles with non-unique optimal alternatives follows by combining the argument below with Lemmas A8–A9 of Hylland (1980).

We first prove a dichotomous lower bound. A utility profile u_N is said to be dichotomous if for every agent $i\in N$, $v_{ix}\in\{0,1\}$ for all $x\in X$ and there exists at least one alternative x where $v_{ix}=1$. Let $Z_i\triangleq\{x\in X|v_{ix}=1\}$ denote the set of preferred alternatives of agent i. For any dichotomous utility profile u_N, if a subset of agents $S \subseteq N$ share the same set Z (i.e., $Z_j = Z$ for all $j \in S$), then

$$\sum_{x\in Z} M_N(u_N)_x \geq \sum_{j\in S} w_j. \tag{12}$$

The proof is the direct-test induction. Induct on $|S|$. If $|S| = 0$, there is nothing to prove. Fix $i \in S$. Construct a modified profile and set $Z_i' = X \setminus Z$. At the modified profile, the induction hypothesis gives the lower bound in (12) for the agents in $S \setminus \{i\}$. WIR for agent i at the modified profile implies that the modified grand allocation puts at most $\sum_{x\in Z} M_{N\setminus\{i\}}(u_{-i})_x$ on Z. Hence $\sum_{x\in Z} M_{N\setminus\{i\}}(u_{-i})_x \geq \sum_{j\in S\setminus\{i\}} w_j$. Returning to the original profile, WIR for agent i gives

$$\sum_{x\in Z} M_N(u_N)_x \geq w_i + \sum_{x\in Z} M_{N\setminus\{i\}}(u_{-i})_x = \sum_{j\in S} w_j,$$

which proves (12).

If all agents have a common unique maximizer a, $|Z_j| = 1$, then (12) gives $M_N = w_N \mathbf{1}_a$. Strategy-proofness extends this unanimous conclusion to all linear profiles with a common unique maximizer a. Starting from the dichotomous unanimous profile, replace agents' reports one by one with their true utilities. If the outcome first ceased to be $w_N\mathbf{1}_a$ when agent i's report was replaced, then agent i, under her true utility, would profit by reporting the dichotomous utility and restoring $w_N\mathbf{1}_a$, contradicting strategy-proofness.

Since M_N satisfies both strategy-proofness and unanimity,we can directly invoke the fixed-weight random-dictatorship theorem of Dutta et al. (2007). Therefore there exist constants $\beta_i \geq 0$, independent of reports and satisfying $\sum_i \beta_i = w_N$, such that at every profile with unique maximizers,

$$M_N(u_N) = \sum_{i\in N} \beta_i \mathbf{1}_{p_i}. \tag{13}$$

WIR identifies the coefficients β_i. Fix i and a. Consider a profile with unique maximizers at which agent i is a dichotomous with $Z_i = \{a\}$ and no other agent has maximizer a. Equation (12) gives $M_N(u_N)_a \geq w_i$. In (13), this coordinate equals β_i. Hence $\beta_i \geq w_i$ for every i. Since $\sum_i \beta_i = \sum_i w_i$, all inequalities bind.

Finally, we consider the case when $X = \{x, y\}$. Let $T_x = \{i : v_{ix} > v_{iy}\}$. Induct on $|T_x|$. The case $|T_x| = 0$ is immediate. Now choose $i \in T_x$ and modify only agent i's report so that y is strictly better than x. By the induction hypothesis, at the modified profile the grand outcome allocates at least $\sum_{j\in T_x\setminus\{i\}} w_j$ to x. WIR for the modified agent implies that the modified grand outcome has x-mass at most $M_{N\setminus\{i\}}(u_{-i})_x$. Thus $M_{N\setminus\{i\}}(u_{-i})_x \geq \sum_{j\in T_x\setminus\{i\}} w_j$. Returning to the original profile, WIR for agent i gives $M_N(u_N)_x \geq w_i + M_{N\setminus\{i\}}(u_{-i})_x$. This proves the bound of x. The bound of y is symmetric. □

C.2 *Proofs for Concave CES utilities*

We provide the formal proof of Theorem 14. Before proceeding to the formal proof, we introduce some definitions.

For $|S| > 0$, write $\mathcal{Q}_S = \{\mathbf{q} \in \mathbb{R}^2_+ : q_x + q_y = |S|\}$. For notational convenience, whenever a grand profile u_N and an agent $i \in N$ are fixed and clear from context, we slightly abuse notation by writing $\mathbf{q} := M_N(u_N) \in \mathcal{Q}_N$ and $\mathbf{r} := M_{N\setminus\{i\}}(u_{-i}) \in \mathcal{Q}_{N\setminus\{i\}}$. Recall that we assume $w_i = 1$ in Theorem 14, the WIR constraint for agent i at u_N can be written as

$$u_i(\mathbf{q}) \geq \max_{\mathbf{d}\in\mathbb{R}^2_+ :\, d_x+d_y=1} u_i(\mathbf{r} + \mathbf{d}).$$

We use three CES types. The balanced type is $u^b(\mathbf{q}) = \left(q_x^\rho + q_y^\rho\right)^{1/\rho}$. For $\eta > 0$, define the almost-x-single-minded type $u^x(\mathbf{q}) = \left(q_x^\rho + \eta q_y^\rho\right)^{1/\rho}$, and the almost-$y$-single-minded type $u^y(\mathbf{q}) = \left(\eta q_x^\rho + q_y^\rho\right)^{1/\rho}$. All three types belong to the positive CES domain.

We first record three elementary consequences as follows.

LEMMA 16. *Suppose* $n \geq 2$. *Fix any grand profile containing at least one balanced agent. If* M *satisfies WIR at this profile, then the grand outcome satisfies* $1 \leq q_x \leq n-1$ *and* $1 \leq q_y \leq n-1$.

PROOF. Let i be a balanced agent. Since $r_x + r_y = n - 1$, agent i can choose $\mathbf{d} \in \mathbb{R}^2_+$ with $d_x + d_y = 1$ such that $r_x + d_x \geq 1$ and $r_y + d_y \geq 1$. Because the balanced CES utility is symmetric and strictly concave, agent i can guarantee a minimum utility of $u^b(1, n-1)$ or $u^b(n-1, 1)$. By WIR, $u^b(\mathbf{q}) \geq u^b(1, n-1)$. Again by symmetry and strict concavity, any allocation in $\mathcal{Q}_N$ with one coordinate strictly below 1 has utility strictly below $u^b(1, n-1)$. It follows that $q_x \geq 1$ and $q_y \geq 1$. Since $q_x + q_y = n$, we also have $q_x, q_y \leq n-1$. □

LEMMA 17. *Fix $\gamma \in (0,1)$. There exists $\bar{\eta} > 0$ such that, for every $\eta \in (0, \bar{\eta})$, the following holds:*

For an almost-x-single-minded agent i, if $u^x(\mathbf{q}) \geq \max_{\mathbf{d} \in \mathbb{R}^2_+ : d_x + d_y = 1} u^x(\mathbf{r} + \mathbf{d})$, then $q_x \geq r_x + 1 - \gamma$; Symmetrically, for an almost-y-single-minded agent i, if $u^y(\mathbf{q}) \geq \max_{\mathbf{d} \in \mathbb{R}^2_+ : d_x + d_y = 1} u^y(\mathbf{r} + \mathbf{d})$, then $q_y \geq r_y + 1 - \gamma$.

PROOF. We prove the statement for u^x. The argument for u^y is symmetric.

First suppose $0 < \rho < 1$. On the compact set $\mathcal{Q}_N = \{\mathbf{q} \in \mathbb{R}^2_+ : q_x + q_y \leq n\}$, we have uniform convergence $u^x(\mathbf{q}) \to q_x$ as $\eta \downarrow 0$. Hence, for sufficiently small $\eta > 0$, $|u^x(\mathbf{q}) - q_x| < \frac{\gamma}{3}$ for all $\mathbf{q} \in \mathcal{Q}_N$.

If $q_x < r_x + 1 - \gamma$, then using the feasible withdrawal $\mathbf{d} = \mathbf{1}_x$, we get

$$u^x(\mathbf{r} + \mathbf{1}_x) > r_x + 1 - \frac{\gamma}{3} > q_x + \frac{2\gamma}{3} > u^x(\mathbf{q}),$$

contradicting the assumed WIR inequality. Therefore $q_x \geq r_x + 1 - \gamma$.

Now suppose $\rho < 0$. For every q, the CES utility satisfies $u^x(\mathbf{q}) \leq q_x$, because adding the positive term ηq_y^ρ inside the CES expression weakly lowers the value when $1/\rho < 0$. Choose the feasible withdrawal $\mathbf{d} = \left(1 - \frac{\gamma}{2}\right) \mathbf{1}_x + \frac{\gamma}{2} \mathbf{1}_y$. Then we have $u^x(\mathbf{r} + \mathbf{d}) \to r_x + 1 - \frac{\gamma}{2}$ uniformly as $\eta \downarrow 0$. Thus, for sufficiently small $\eta > 0$, $u^x(\mathbf{r} + \mathbf{d}) > r_x + 1 - \frac{\gamma}{2} - \frac{\gamma}{2} = r_x + 1 - \gamma$. If $q_x < r_x + 1 - \gamma$, then

$$u^x(\mathbf{q}) \leq q_x < r_x + 1 - \gamma < u^x(\mathbf{r} + \mathbf{d}),$$

again contradicting the assumed WIR inequality. Hence $q_x \geq r_x + 1 - \gamma$. □

LEMMA 18. *Suppose $n \geq 2$. Fix an integer $k \geq 0$ such that $2k + 1 \leq n$, and let $\gamma \in (0,1)$. For sufficiently small $\eta > 0$, at every grand profile containing at least k agents of type u^x, at least k agents of type u^y, and at least one balanced agent u^b, WIR* *implies*

$$q_x \geq k+1-2k\gamma, \qquad q_y \geq k+1-2k\gamma.$$

PROOF. We prove the claim by induction on k.

The base case $k=0$ follows directly from Lemma 16, which implies $q_x \geq 1$ and $q_y \geq 1$.

Now suppose the claim holds for $k-1$, where $k \geq 1$. Consider a profile u_N containing at least k agents of type u^x, at least k agents of type u^y, and at least one balanced agent u^b.

We first prove the lower bound for q_x. Choose an agent h of type u^x, and construct a modified grand profile u'_N by changing only agent h's report from u^x to u^y. Let $\mathrm{q}' = M_N(u'_N)$. The modified profile contains at least $k-1$ agents of type u^x, at least $k+1$ agents of type u^y, and at least one balanced agent u^b. Hence, by the induction hypothesis, $q'_x \geq k-2(k-1)\gamma$. On the other hand, applying Lemma 17 to agent h at the modified profile gives $q'_y \geq r_y+1-\gamma$. Since $q'_x+q'_y=n$ and $r_x+r_y=n-1$, this implies $q'_x = n-q'_y \leq n-(r_y+1-\gamma) = r_x+\gamma$. Therefore $r_x \geq q'_x-\gamma \geq k-2(k-1)\gamma-\gamma = k-(2k-1)\gamma$. Returning to the original profile, agent h is of type u^x. Applying Lemma 17 again yields $q_x \geq r_x+1-\gamma \geq k+1-2k\gamma$.

The proof of the lower bound for q_y is symmetric: choose an agent of type u^y, change her report to u^x, apply the induction hypothesis to the modified profile, and then apply Lemma 17 twice. Hence $q_y \geq k+1-2k\gamma$. This completes the induction. □

PROOF OF THEOREM 14. Suppose, toward a contradiction, that a mechanism M satisfies WIR. We now derive the contradiction.

First suppose n is odd. Write $n=2k+1$ with $k \geq 1$. Choose a grand profile with exactly k agents of type u^x, exactly k agents of type u^y, and one agent of the balanced type u^b. By Lemma 18, we have $q_x \geq k+1-2k\gamma$ and $q_y \geq k+1-2k\gamma$. Thus $n = q_x+q_y \geq 2k+2-4k\gamma$. But $n=2k+1$. Choosing $0<\gamma<\frac{1}{4k}$ gives a contradiction.

Now suppose n is even. Write $n=2k+2$ with $k \geq 0$. Choose a grand profile u_N with $k+1$ agents of type u^x, k agents of type u^y, and one agent of the balanced type u^b. Applying the same argument as in Lemma 18, we have $q_x \geq k+2-(2k+2)\gamma$ and $q_y \geq k+1-2k\gamma$. Combining these bounds with $q_x + q_y = n$ yields

$$n = q_x + q_y \geq \big(k+2-(2k+2)\gamma\big) + \big(k+1-2k\gamma\big) = 2k+3-(4k+2)\gamma.$$

Since $n = 2k+2$, choosing $0 < \gamma < \frac{1}{4k+2}$ gives a contradiction.

In both parity cases, we obtain a contradiction for sufficiently small $\gamma > 0$ and then sufficiently small $\eta > 0$. Hence no mechanism satisfies WIR on the positive binary CES domain. □

C.3 *Proofs for L_p-norm utilities*

PROOF OF THEOREM 15. Write $\mathbf{e}^x = (1,0)$ and $\mathbf{e}^y = (0,1)$ for the single-minded coefficient vectors. The proof has two steps.

First, WIR pins down all single-minded binary profiles. In this step, we restrict attention to profiles in which every agent reports one of these two single-minded coefficient vectors. For any $i \in N$, let $T_x = \{j \in N \setminus \{i\} : \mathbf{v}_j = \mathbf{e}^x\}$. We claim that $M_{N\setminus\{i\}}(u_{-i})_x = |T_x|$.

To prove the lower bound, use induction on $|T_x|$. The case $|T_x| = 0$ is trivial. Suppose $|T_x| = k \geq 1$, choose $j \in T_x$, and consider the grand profile in which exactly the agents in T_x report e_x. Applying agent j's WIR constraint and the induction hypothesis to the residual coalition $N \setminus \{j\}$, we obtain $M_N(u_N)_x \geq 1 + M_{N\setminus\{j\}}(u_{-j})_x \geq k$. Choose an agent i who is single-minded for y. WIR for agent i gives $n - M_N(u_N)_x \geq 1 + n - 1 - M_{N\setminus\{i\}}(u_{-i})_x$, so $M_{N\setminus\{i\}}(u_{-i})_x \geq M_N(u_N)_x \geq k$. Thus $M_{N\setminus\{i\}}(u_{-i})_x \geq |T_x|$ for all i and T_x. Applying the same argument with x and y interchanged gives $n-1-M_{N\setminus\{i\}}(u_{-i})_x \geq n-1-|T_x|$, and hence $M_{N\setminus\{i\}}(u_{-i})_x \leq |T_x|$. This proves the claim.

It follows that, if exactly $K \subseteq N$ report e_x at a grand profile, then the grand outcome is $|K|\mathbf{1}_x + (n-|K|)\mathbf{1}_y$. Indeed, if $i \in K$, then agent i reports $\mathbf{e}^x$, and WIR gives $M_N(u_N)_x \geq 1 + M_{N\setminus\{i\}}(u_{-i})_x$. By the claim above, $M_{N\setminus\{i\}}(u_{-i})_x = |K| - 1$, so $M_N(u_N)_x \geq |K|$. Similarly, if $i \notin K$, then agent i reports $\mathbf{e}^y$, and WIR gives $M_N(u_N)_y \geq n - |K|$. Feasibility then makes both inequalities bind.

Second, choose $\varepsilon > 0$ small. The only restriction on ε needed below is that, for every $t \in (0, n-1]$,

$$t^p + (1-\varepsilon)^p (n-t)^p < (1-\varepsilon)^p n^p. \tag{14}$$

Such an ε exists because the continuous function $\big(n^p - (n-t)^p\big)/t^p$ has a minimum strictly larger than one on $(0, n-1]$.

Let $\mathbf{v}^y = (1-\varepsilon, 1)$ and $\mathbf{v}^x = (1, 1-\varepsilon)$. We show first that the profile $(\mathbf{v}^y, \mathbf{v}^x, \ldots, \mathbf{v}^x)$ must have grand outcome $n\mathbf{1}_y$. Start from the unanimous profile $(\mathbf{e}^y, \ldots, \mathbf{e}^y)$, whose outcome is $n\mathbf{1}_y$ by the previous argument. Change agents $2, \ldots, n$, one at a time, from $\mathbf{e}^y$ to $\mathbf{v}^x$. At each intermediate profile before the last change, at least one agent still reports $\mathbf{e}^y$. Consider the agent whose report has just been changed to $\mathbf{v}^x$. If she instead reported $\mathbf{e}^y$, the preceding profile would obtain outcome $n\mathbf{1}_y$. Strategy-proofness therefore requires her truthful utility to be at least $n(1-\varepsilon)$. Applying WIR to a profile in which there exists an agent who reports $\mathbf{e}^y$ gives $M_N(u_N)_y \geq 1$. If the truthful outcome had $M_N(u_N)_x = t > 0$, then $t \in (0, n-1]$, and (14) would imply $\left(t^p + (1-\varepsilon)^p (n-t)^p\right)^{1/p} < n(1-\varepsilon)$, contradicting strategy-proofness. Hence the truthful outcome remains $n\mathbf{1}_y$ after each such change. Finally change agent 1 from $\mathbf{e}^y$ to $\mathbf{v}^y$. If she reported $\mathbf{e}^y$, the outcome would be $n\mathbf{1}_y$, which gives her utility n. Since n is the maximal feasible utility for type $\mathbf{v}^y$, strategy-proofness forces the outcome to remain $n\mathbf{1}_y$.

A symmetric argument starts from the profile $(\mathbf{e}^x, \ldots, \mathbf{e}^x)$, whose outcome is $n\mathbf{1}_x$. First change agent 1 from $\mathbf{e}^x$ to $\mathbf{v}^y$. If she reported $\mathbf{e}^x$, she would obtain utility $n(1-\varepsilon)$ at $n\mathbf{1}_x$. WIR for any remaining agents who report $\mathbf{e}^x$ gives $M_N(u_N)_x \geq 1$. If the truthful outcome had positive mass on y, the analogue of (14) would make agent 1's truthful utility strictly less than $n(1-\varepsilon)$, contradicting strategy-proofness. Hence the outcome remains $n\mathbf{1}_x$. Then change agents $2, \ldots, n$, one at a time, from $\mathbf{e}^x$ to $\mathbf{v}^x$. At each step, if the just-changed agent instead reported $\mathbf{e}^x$, the preceding profile would obtain outcome $n\mathbf{1}_x$, giving her truthful utility n. Since n is the maximal feasible utility for type $\mathbf{v}^x$, and is attained only at $n\mathbf{1}_x$, strategy-proofness forces the outcome to remain $n\mathbf{1}_x$ after each such change. This forces the same profile $(\mathbf{v}^y, \mathbf{v}^x, \ldots, \mathbf{v}^x)$ to have grand outcome $n\mathbf{1}_x$, contradicting the previous paragraph. Therefore no mechanism can satisfy WIR and strategy-proofness. □

REFERENCES

Abdulkadiroğlu, Atila and Tayfun Sönmez (1999), "House allocation with existing tenants." *Journal of Economic Theory*, 88(2), 233–260. [7]

Barberà, Salvador, Faruk Gul, and Ennio Stacchetti (1993), "Generalized median voter schemes and committees." *Journal of Economic Theory*, 61(2), 262–289. [7]

Black, Duncan (1948), "On the rationale of group decision-making." *Journal of Political Economy*, 56(1), 23–34. [7]

Bogomolnaia, Anna and Hervé Moulin (2001), "A new solution to the random assignment problem." *Journal of Economic Theory*, 100(2), 295–328. [6]

Brandt, Felix, Christian Geist, and Dominik Peters (2017), "Optimal bounds for the no-show paradox via sat solving." *Mathematical Social Sciences*, 90, 18–27. [7]

Che, Yeon-Koo and Fuhito Kojima (2010), "Asymptotic equivalence of probabilistic serial and random priority mechanisms." *Econometrica*, 78(5), 1625–1672. [6]

Conitzer, Vincent, Rupert Freeman, and Nisarg Shah (2017), "Fair public decision making." In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 629–646. [6]

Demange, Gabrielle (1982), "Single-peaked orders on a tree." *Mathematical Social Sciences*, 3(4), 389–396. [7]

Duggan, John (1996), "A geometric proof of Gibbard's random dictatorship theorem." *Economic Theory*, 7(2), 365–369. [6]

Dutta, Bhaskar, Hans Peters, and Arunava Sen (2007), "Strategy-proof cardinal decision schemes." *Social Choice and Welfare*, 28(1), 163–179. [6, 37]

Fishburn, Peter C and Steven J Brams (1983), "Paradoxes of preferential voting." *Mathematics Magazine*, 56(4), 207–214. [7]

Furusawa, Taiji and Hideo Konishi (2011), "Contributing or free-riding? voluntary participation in a public good economy." *Theoretical Economics*, 6(2), 219–256. [7]

Gibbard, Allan (1977), "Manipulation of schemes that mix voting with chance." *Econometrica: Journal of the Econometric Society*, 665–681. [6]

Hirai, Toshiyuki and Ryusuke Shinohara (2024), "Voluntary participation in a negotiation on providing public goods and renegotiation opportunities." *Journal of Economic Behavior & Organization*, 224, 1–19. [7]

Hylland, Aanund (1980), "Strategy proofness of voting procedures with lotteries as outcomes and infinite sets of strategies." Technical report, mimeo. [6, 28, 36]

Hylland, Aanund and Richard Zeckhauser (1979), "The efficient allocation of individuals to positions." *Journal of Political Economy*, 87(2), 293–314. [6]

Moulin, Hervé (1980), "On strategy-proofness and single peakedness." *Public Choice*, 35(4), 437–455. [7]

Moulin, Hervé (1988), "Condorcet's principle implies the no show paradox." *Journal of Economic Theory*, 45(1), 53–64. [7]

Nehring, Klaus and Clemens Puppe (2007), "The structure of strategy-proof social choice—part i: General characterization and possibility results on median spaces." *Journal of Economic Theory*, 135(1), 269–305. [7]

Ozkes, Ali I and M Remzi Sanver (2023), "Uniform random dictatorship: A characterization without strategy-proofness." *Economics Letters*, 227, 111127. [6]

Saijo, Tatsuyoshi and Takehiko Yamato (2010), "Fundamental impossibility theorems on voluntary participation in the provision of non-excludable public goods." *Review of Economic Design*, 14(1), 51–73. [7]

Schummer, James and Rakesh V Vohra (2002), "Strategy-proof location on a network." *Journal of Economic Theory*, 104(2), 405–428. [7]

Sen, Arunava (2011), "The Gibbard random dictatorship theorem: a generalization and a new proof." *SERIEs*, 2(4), 515–527. [6]

Co-editor [Name Surname; will be inserted later] handled this manuscript.