Axioms for Top Trading Cycles in Multi-Object Reallocation*

Jacob Coreno[a,†] Di Feng[b]

February 4, 2026

Abstract

This paper studies multi-object reallocation without monetary transfers, where agents initially own multiple indivisible objects and have strict preferences over bundles (e.g., shift exchange among workers at a firm). Focusing on marginal rules that elicit only rankings over individual objects, we provide axiomatic characterizations of the generalized Top Trading Cycles rule (TTC) on the lexicographic and responsive domains. On the lexicographic domain, TTC is characterized by balancedness, individual-good efficiency, the worst-endowment lower bound, and either truncation-proofness or drop strategy-proofness. On the responsive domain, TTC is the unique marginal rule satisfying individual-good efficiency, truncation-proofness, and either the worst-endowment lower bound or individual rationality. In the Shapley–Scarf housing market, TTC is characterized by Pareto efficiency, individual rationality, and truncation-proofness. Finally, on the conditionally lexicographic domain, the augmented Top Trading Cycles rule is characterized by balancedness, Pareto efficiency, the worst-endowment lower bound, and drop strategy-proofness. The conditionally lexicographic domain is a maximal domain on which Pareto efficiency coincides with individual-good efficiency.

Keywords: Top Trading Cycles; multi-object exchange; responsive preferences; lexicographic preferences; truncation-proofness; worst-endowment lower bound.

JEL Classification: C78; D47; D71.

*We are especially grateful to Ivan Balbuzanov for his guidance and support throughout this project. We also thank Açelya Altuntaş, Onur Kesten, Bettina Klaus, Alexandru Nichifor, Szilvia Pápai, Tom Wilkening, and Steven Williams, as well as seminar participants at HEC Lausanne, Deakin University, the 2024 Economic Theory Festival at UNSW, ESAM 2024, and MATCH-UP 2024 for their helpful comments and suggestions. Feng gratefully acknowledges financial support from the National Natural Science Foundation of China No. 72503028. Coreno received financial support from an Australian Government RTP Scholarship, a University of Melbourne FBE-GRATS Scholarship, and an M. A. Bartlett Scholarship. Coreno and Feng gratefully acknowledge financial support from the Swiss National Science Foundation (SNSF) through Project 100018-231777. This article is based upon work supported by the National Science Foundation under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778, while authors were in residence at the Simons Laufer Mathematical Sciences Institute (formerly MSRI), Berkeley, CA, Fall 2023.

†Correspondence to: Faculty of Business and Economics, University of Lausanne, Internef 530, 1015 Lausanne, Switzerland.

[a]University of Lausanne, Switzerland. Contact: jacob.coreno@unil.ch.

[b]Dongbei University of Finance and Economics, China. Contact: dfeng@dufe.edu.cn.

1 Introduction

We study *multi-object reallocation* without monetary transfers. A problem instance consists of a group of agents, an initial endowment in which each agent owns a set of heterogeneous and indivisible objects, and each agent's strict preferences over bundles. A *rule* maps reported preferences to a feasible reassignment. When each agent is endowed with a single object, the model reduces to the classic *housing market* of Shapley and Scarf (1974).

Reallocation problems arise in a variety of applications. Firms plan shift schedules months in advance, and employees exchange their assigned shifts (e.g., Manjunath and Westkamp (2021)). Universities exchange students through programs such as the Tuition Exchange Program in the United States and Erasmus in Europe (e.g., Dur and Ünver (2019); Bloch et al. (2020)). Living-donor organ exchange programs provide further examples featuring both single- and multi-object exchange (e.g., Roth et al. (2004, 2005); Ergin et al. (2017, 2020)).[1]

Multi-object reallocation confronts two difficulties that are largely absent in single-object reallocation. The first is the vast number of possible bundles, which makes it infeasible for agents to accurately report their preferences over bundles in practice.[2] As Roth (2015, p. 331) explains, "a practical mechanism must simplify the language in which preferences can be reported, and by doing so it will restrict which preferences can be reported." Accordingly, we focus primarily on *marginal* rules, under which agents report only rankings over individual objects (their *marginal preferences*).[3]

When working with marginal rules, we focus on two preference domains in which marginal preferences are informative about the underlying bundle comparisons. On the lexicographic domain, marginal rules are without loss of generality because an agent's ranking over individual objects fully determines her preferences over bundles. On the more general responsive domain (Roth (1985)), marginal preferences remain informative—especially when comparing bundles that differ by one object—but they need not determine all possible bundle comparisons. Both domains rule out complementarities, which cannot be expressed via marginal preferences.[4]

The second difficulty is the tension among standard objectives relating to efficiency, voluntary participation, and incentive compatibility. While the ideal desiderata—*Pareto efficiency (PE)*, *individual rationality (IR)*, and *strategy-proofness (SP)*—are jointly attainable in housing

[1]Related work studies course allocation as a *multi-object assignment* problem without endowments (e.g., Budish (2011); Budish and Cantillon (2012); Bichler et al. (2021)).

[2]For instance, in shift reallocation with only 20 shifts, there are $\binom{20}{5} = 15{,}504$ distinct bundles of five shifts.

[3]This reporting format is used, for example, in the National Resident Matching Program (NRMP): hospitals submit "rank-order lists" over individual doctors despite having potentially complex preferences over sets of doctors (e.g., Roth and Peranson (1999), Milgrom (2009, 2011)).

[4]To address complementarities, Section 6 considers a richer but still succinct reporting language based on "lexicographic preference trees".

markets,[5] they become incompatible once some agent owns more than one object (Sönmez (1999)). Much of the axiomatic literature sidesteps this incompatibility by insisting on strategy-proofness and sacrificing one of the other desiderata. Strategy-proofness is especially restrictive in multi-object environments. Under mild auxiliary axioms, combining SP with PE leads to *Sequential Dictatorships* (Pápai (2001); Klaus and Miyagawa (2002); Ehlers and Klaus (2003); Hatfield (2009)), while combining SP with IR leads to *Segmented Trading Cycles rules* (Pápai (2003, 2007)). Either way, insisting on strategy-proofness forces major concessions—either by effectively ignoring the initial ownership structure or by accepting severe inefficiency.

In this paper, we circumvent the incompatibility by making modest, practically motivated compromises relative to the ideal notions of efficiency and incentive compatibility. We study the *generalized Top Trading Cycles rule (TTC)*, a marginal rule based on a multi-object extension of the celebrated Top Trading Cycles algorithm (attributed to David Gale by Shapley and Scarf (1974)). We show that TTC comes close to the unattainable ideals: it delivers meaningful guarantees regarding efficiency, participation, and incentive compatibility. More concretely, we characterize TTC on various domains using *balancedness (BAL)*, *individual-good efficiency (IGE)*, the participation guarantees IR and the *worst-endowment lower bound (WELB)*, and the incentive requirements *truncation-proofness (TP)* and *drop strategy-proofness (DSP)*.

On the lexicographic domain, TTC is characterized by BAL, PE, WELB, and either TP or DSP (Theorems 1 and 2). On the responsive domain, where no marginal rule can satisfy IGE, IR, and DSP (Corollary 1), TTC is characterized as the unique marginal rule satisfying IGE, TP, and either WELB or IR (Theorems 3 and 4). In housing markets, our axioms simplify, and TTC is characterized by PE, IR, and TP (Theorem 5). We briefly discuss these properties below, deferring formal definitions and further intuition to Section 2.2.

Balancedness (BAL) requires that each agent receive as many objects as in her endowment. In many applications this is imposed as a feasibility constraint.[6] We treat BAL as an axiom, rather than a model restriction, so that we can study its interaction with other properties (see Lemmata 1 and 3).

Individual rationality (IR) and the worst-endowment lower bound (WELB) are distinct participation guarantees. IR ensures that the outcome weakly Pareto dominates the initial allocation; that is, every agent enjoys a bundle at least as good as her endowment. WELB is an object-wise guarantee that ensures no agent receives any object ranked below the worst in her endowment according to her marginal preferences. IR and WELB coincide in single-object reallocation. In multi-object problems the properties diverge, but IR implies WELB under mild

[5] In housing markets, these properties uniquely characterize the Top Trading Cycles rule (Ma (1994)).

[6] In shift reallocation it may reflect contractual constraints on the number of shifts per worker (Manjunath and Westkamp (2021, 2025)). Similar constraints arise in time banks and exchange programs (Andersson et al. (2021); Biró et al. (2022a); Dur and Ünver (2019)).

additional structure.[7] TTC satisfies both guarantees on each domain we study.

For efficiency, we use Pareto efficiency (PE) when it is attainable and individual-good efficiency (IGE) as a tractable relaxation. IGE relaxes PE by ruling out mutually beneficial exchanges where each participating agent trades a single object. Such trades are comparatively easy to identify and execute, and under responsive preferences they can be detected using only marginal preferences. The two properties coincide in single-object reallocation and multi-object reallocation with lexicographic preferences. On the responsive domain, the relaxation is substantive: full PE is computationally demanding (e.g., De Keijzer et al. (2009); Aziz et al. (2019)) and, among marginal rules, incompatible with IR (Manjunath and Westkamp (2025)).

Drop strategy-proofness (DSP) and truncation-proofness (TP) protect against simple manipulations of agents' marginal preferences. An agent implements a *drop strategy* by demoting an object she does not own to the bottom of her marginal preference list (Altuntaş et al. (2023)); she implements a *truncation strategy* by demoting every such object that she ranks below a chosen cutoff. Intuitively, truncation strategies can be thought of as "shortening the list of acceptable objects" in the standard sense from two-sided matching (e.g., Mongell and Roth (1991); Roth and Rothblum (1999); Ehlers (2008)). DSP and TP require that no agent can improve her assignment using the corresponding class of manipulations.

In addition to marginal preferences, we also consider a richer—yet still succinct—reporting language that can express complementarities. We study *conditionally lexicographic* preferences, which admit a compact representation via "lexicographic preference trees" and preserve the equivalence between IGE and PE found under lexicographic preferences. On this domain, we study the *Augmented Top Trading Cycles rule (ATTC)* introduced by Fujita et al. (2018). We show that ATTC satisfies natural extensions of our key properties and is characterized by BAL, PE, WELB, and DSP (Theorem 6). Finally, we show that conditionally lexicographic preferences form a maximal domain on which PE and IGE are equivalent (Proposition 5); beyond this boundary, ruling out Pareto-improving single-object exchanges no longer guarantees full Pareto efficiency.

The paper proceeds as follows. Section 2 introduces the model, preference domains, properties, and TTC. Sections 3 and 4 present our characterizations on the lexicographic and responsive domains, and Section 5 specializes the results to housing markets. Section 6 studies conditionally lexicographic preferences and characterizes ATTC. Section 7 discusses related literature. The appendices formalize lexicographic preference trees, contain omitted proofs, and provide additional examples.

[7]See Lemma 1 and the alternative preference domains in the surrounding discussion.

2 Model

Let $N = \{1, \ldots, n\}$ be a finite set of $n \geq 2$ *agents*. Let O be a finite set of heterogeneous and indivisible *objects* such that $|O| \geq n$. A *bundle* is a subset of O. Let 2^O denote the set of bundles. We denote generic elements of O by lowercase letters (e.g., x, y, z), and generic elements of 2^O by uppercase letters (e.g., X, Y, Z). To simplify notation, when there is no risk of confusion, we identify a singleton set $\{x\}$ with the element x itself. For example, we write $X \cup x$ to denote $X \cup \{x\}$.

An *allocation* is a function $\mu : N \to 2^O$ such that (i) for all $i \in N$, $\mu(i) \neq \emptyset$, (ii) for all $i, j \in N$, $i \neq j$ implies $\mu(i) \cap \mu(j) = \emptyset$, and (iii) $\bigcup_{i \in N} \mu(i) = O$. Thus, an allocation μ can be represented as a profile $(\mu_i)_{i \in N}$ of nonempty, pairwise disjoint bundles satisfying $\bigcup_{i \in N} \mu_i = O$. For each $i \in N$, μ_i is called agent i's *assignment* at μ. Let $\mathcal{A}$ denote the set of allocations.

The *initial allocation*, also referred to as the *endowment allocation*, is denoted by $\omega = (\omega_i)_{i \in N}$. For each $i \in N$, ω_i is called agent i's *endowment*. The initial *owner* of an object o is the unique agent i with $o \in \omega_i$.

Each agent i has a *(strict) preference* relation P_i on the set of bundles. We assume that P_i belongs to some specified subset $\mathcal{P}_i$ of all strict preference relations on 2^O. In subsequent sections, we impose further structure on the sets $\mathcal{P}_i$. If agent i prefers bundle X to Y, then we write $X \; P_i \; Y$. Let R_i denote the "at least as good as" relation associated with P_i, defined by $X \; R_i \; Y$ if and only if ($X \; P_i \; Y$ or $X = Y$).[8] Given a nonempty bundle $X \in 2^O$, $\max_{P_i}(X)$ denotes the most-preferred object in X at P_i, i.e., $\max_{P_i}(X) = x$ if $x \in X$ and $x \; R_i \; y$ for all $y \in X$. Similarly, $\min_{P_i}(X)$ denotes the least-preferred object in X at P_i, i.e., $\min_{P_i}(X) = x$ if $x \in X$ and $y \; R_i \; x$ for all $y \in X$. A *preference profile* is an indexed family $P = (P_i)_{i \in N}$ of preference relations. The *domain* is the set $\mathcal{P} \coloneqq \prod_{i \in N} \mathcal{P}_i$, representing all possible preference profiles under consideration.

An object reallocation problem (or simply a *problem*) is a triple (N, ω, P). Since (N, ω) remains fixed throughout, we identify a problem with its preference profile P. Thus, the domain $\mathcal{P} = \prod_{i \in N} \mathcal{P}_i$ of preference profiles represents the set of all problems.

A *rule (on $\mathcal{P}$)* is a function $\varphi : \mathcal{P} \to \mathcal{A}$ that associates with each preference profile P an allocation $\varphi(P)$. For each $i \in N$, $\varphi_i(P)$ denotes agent i's assignment at $\varphi(P)$.

Marginal rules

A rule is called *marginal* if it can be implemented with a simple reporting language consisting of linear orders over individual objects. Marginal rules are attractive in applications because they ask agents only for a ranking over individual objects, rather than a ranking over all feasible

[8] Formally, R_i is a linear order (i.e., a *complete*, *transitive*, and *antisymmetric* binary relation) on 2^O , and P_i is the strict (i.e., *irreflexive* and *asymmetric*) part of R_i.

bundles, and thus substantially reduce the informational and cognitive burden on participants. A prominent example is the National Resident Matching Program (NRMP), in which hospitals submit rank-order lists over individual doctors, even though they may have complex preferences over sets of doctors (see, e.g., Roth and Peranson (1999), Roth (2002), Milgrom (2009, 2011)).

Formally, given a preference relation P_i on 2^O, the *marginal preference* over a subset $X \subseteq O$ of individual objects, denoted $P_i|_X$, is the restriction of P_i to singleton subsets of X.[9] Similarly, $R_i|_X$ denotes the restriction of R_i to singleton subsets of X. We often represent a marginal preference $P_i|_X$ as an ordered list of objects; for example, $P_i|_X : x_1, x_2, \ldots, x_{|X|}$ means that $x_1 \; P_i \; x_2 \; P_i \cdots P_i \; x_{|X|}$, and $P_i|_X : x_1, x_2, \ldots, x_k, \ldots$ means that $x_1 \; P_i \; x_2 \; P_i \cdots P_i \; x_k \; P_i \; o$ for all $o \in X \setminus \{x_1, x_2, \ldots, x_k\}$. Given a preference profile $P = (P_i)_{i \in N}$, the *marginal preference profile* over X is the profile $P|_X = (P_i|_X)_{i \in N}$.

A rule is marginal if it depends solely on agents' marginal preferences over O.

Definition 1. A rule φ is *marginal (MAR)* if, for all $P, P' \in \mathcal{P}$, $P|_O = P'|_O$ implies $\varphi(P) = \varphi(P')$.

Marginal rules are most natural in settings where objects behave like substitutes, so that agents' preferences over bundles are captured reasonably well by their rankings over individual objects. Most of our analysis focuses on these settings (Sections 3–5). The main limitation of marginal rules is that they do not allow agents to express complementarities across objects. To address complementarities, Section 6 (and Appendix A) considers a richer but still succinct reporting language based on "lexicographic preference trees".

2.1 Preference domains

We now describe the lexicographic and responsive domains, two structured preference domains where objects behave like substitutes.

An agent has lexicographic preferences if, when evaluating distinct bundles X and Y, she prefers the bundle containing the most-preferred object in $X \cup Y$; if the most-preferred object in $X \cup Y$ is common to X and Y, then she prefers the bundle containing the second-most-preferred object in $X \cup Y$, and so on.[10] Formally, a preference relation P_i on 2^O is *lexicographic* if, for any two distinct bundles X and Y,

$$X \; P_i \; Y \iff \max_{P_i} (X \triangle Y) \in X,$$

where $X \triangle Y = (X \setminus Y) \cup (Y \setminus X)$ is the symmetric difference between X and Y. For each $i \in N$, let $\mathcal{L}_i$ be the set of lexicographic preferences on 2^O, and let $\mathcal{L} \coloneqq \prod_{i \in N} \mathcal{L}_i$ be the *lexicographic domain.*

[9] Equivalently, $P_i|_X$ is the strict linear order on X such that, for all $x, y \in X$, $x \; P_i|_X \; y$ if and only if $x \; P_i \; y$.

[10] Similar "take the best" decision heuristics have been documented in psychology (e.g., Gigerenzer and Goldstein (1996), Gigerenzer and Todd (1999)).

Although it is somewhat restrictive, the lexicographic domain is a natural starting point in our analysis. Each lexicographic preference $P_i \in \mathcal{L}_i$ is uniquely determined by its marginal preference $P_i|_O$ over individual objects; that is, for all $P_i, P_i' \in \mathcal{L}_i$, $P_i|_O = P_i'|_O$ if and only if $P_i = P_i'$. Consequently, any rule defined on $\mathcal{L}$ is automatically marginal. We therefore identify each $P_i \in \mathcal{L}_i$ with its associated marginal preference $P_i|_O$ and write $P_i : x_1, x_2, \ldots, x_{|O|}$ if $x_1 \; P_i \; x_2 \; P_i \cdots P_i \; x_{|O|}$.

Responsiveness is a more general condition first studied by Roth (1985) for many-to-one matching models. An agent has responsive preferences if, for any two bundles that differ in one object, she prefers the bundle containing the more-preferred object. Formally, a preference relation P_i on 2^O is *responsive* if, for any bundle X and any objects $y, z \in O \setminus X$,

$$(X \cup y) \; P_i \; (X \cup z) \iff y \; P_i \; z.$$

Responsive preferences rule out complementarities because the relative ranking between any two objects is independent of the other objects they are obtained with. For each $i \in N$, let $\mathcal{R}_i$ denote the set of responsive preferences on 2^O, and let $\mathcal{R} \coloneqq \prod_{i \in N} \mathcal{R}_i$ be the *responsive domain*.

It is helpful to compare the lexicographic and responsive domains with the domain of monotonic preferences. A preference relation P_i on 2^O is *monotonic* if, for any bundles X and Y, $X \; R_i \; Y$ whenever $X \supseteq Y$. For each $i \in N$, let $\mathcal{M}_i$ be the set of monotonic preferences on 2^O, and let $\mathcal{M} \coloneqq \prod_{i \in N} \mathcal{M}_i$ be the *monotonic domain*.

Every lexicographic preference is both responsive and monotonic, while responsiveness and monotonicity are logically independent. Formally, for each $i \in N$, we have $\mathcal{L}_i \subseteq \mathcal{R}_i \cap \mathcal{M}_i$ (with strict inclusion if $|O| \geq 3$), $\mathcal{M}_i \nsubseteq \mathcal{R}_i$ (whenever $|O| \geq 3$), and $\mathcal{R}_i \nsubseteq \mathcal{M}_i$.[11]

2.2 Properties of allocations and rules

This section introduces several desirable properties of allocations and rules. Unless stated otherwise, these properties are defined for an arbitrary domain $\mathcal{P}$ of strict preferences over bundles. Our main focus will be the lexicographic and responsive domains introduced above.

Our first requirement is that each agent ends up with the same number of objects as initially endowed. An allocation μ is *balanced* if, for each agent $i \in N$, $|\mu_i| = |\omega_i|$.

Definition 2. A rule φ satisfies *balancedness (BAL)* if, for each $P \in \mathcal{P}$, $\varphi(P)$ is balanced.

Balancedness is often treated as an inviolable constraint in practical reallocation problems. In shift reallocation, for instance, it prevents overwork or underemployment and may reflect contractual constraints on the number of shifts per worker (Manjunath and Westkamp (2021,

[11]For example, if $O = \{a, b\}$, then $P_i : \{a\}, \emptyset, \{a, b\}, \{b\}$ and $P_i' : \emptyset, \{a\}, \{b\}, \{a, b\}$ are responsive but not monotonic. If $O = \{a, b, c\}$, then $P_i' : \{a, b, c\}, \{a, c\}, \{a, b\}, \{b, c\}, \{a\}, \{b\}, \{c\}, \emptyset$ is monotonic but not responsive.

2025)). In student and tuition exchange programs, BAL helps to maintain reciprocity and control education costs at popular schools (Andersson et al. (2021), Biró et al. (2022a), Dur and Ünver (2019)).[12] We model BAL as a property of rules, rather than as an explicit feasibility constraint, so that we can study its interaction with other properties (see Lemmata 1 and 3).

2.2.1 Efficiency

An allocation $\overline{\mu}$ *Pareto dominates* another allocation μ at a preference profile P if (i) for all $i \in N$, $\overline{\mu}_i \; R_i \; \mu_i$, and (ii) for some $i \in N$, $\overline{\mu}_i \; P_i \; \mu_i$. An allocation μ is *Pareto efficient* at P if it is not Pareto dominated at P by any allocation.

Definition 3. A rule φ satisfies *Pareto efficiency (PE)* if, for each $P \in \mathcal{P}$, $\varphi(P)$ is Pareto efficient at P.

In principle, any allocation that violates Pareto efficiency could be destabilized by a coalition of agents carrying out a mutually beneficial exchange. In practice, however, such exchanges may be hard to realize. A coalition would typically need detailed information about the agents' preferences over bundles and a way to coordinate intricate trades of multiple objects among many agents. Even with the required preference information, finding a Pareto improvement—or certifying that none exists—is computationally demanding. In particular, under additive preferences (a subclass of responsive preferences), deciding whether a Pareto improvement exists is NP-complete (e.g., De Keijzer et al. (2009); Aziz et al. (2019)).

Individual-good efficiency relaxes PE by ruling out a restricted class of mutually beneficial exchanges—multilateral trades where each participating agent offers a single object in return for another. Formally, a *trading cycle* at an allocation μ is a cyclic sequence

$$C = (o_1, i_1, o_2, i_2, \dots, i_{k-1}, o_k, i_k, o_{k+1} = o_1)$$

consisting of $k \geq 1$ distinct objects and k distinct agents such that, for each $\ell \in \{1, \dots, k\}$, $o_\ell \in \mu_{i_\ell}$. For each $\ell \in \{1, \dots, k\}$, we interpret o_ℓ as the object relinquished by agent i_ℓ and $o_{\ell+1}$ as the object received. We say that C is *Pareto improving* at a preference profile P if

$$\text{for each } \ell \in \{1, \dots, k\}, \quad (\mu_{i_\ell} \cup o_{\ell+1}) \setminus o_\ell \; P_{i_\ell} \; \mu_{i_\ell}.$$

An allocation μ is *individual-good efficient* at P if no trading cycle at μ is Pareto improving at P.[13]

[12] Dur and Ünver (2019) document several cases in which such long-run imbalances led to the failure of an exchange program.

[13] This terminology is borrowed from Biró et al. (2022a). Similar properties are studied in Aziz et al. (2019), Caspari (2020), and Coreno and Balbuzanov (2022).

Definition 4. A rule φ satisfies *individual-good efficiency (IGE)* if, for each $P \in \mathcal{P}$, $\varphi(P)$ satisfies individual-good efficiency at P.

Although PE and IGE coincide on the lexicographic domain (Aziz et al. (2019)), on more general preference domains an allocation may satisfy IGE even when further Pareto improvements are possible through more complex trades. This loss of efficiency is the price we pay for a notion that is tractable and achievable in practice. IGE focuses on single-object exchanges, which are relatively easy to organize,[14] and under responsive preferences such exchanges can be detected using only the information elicited by a marginal rule. In addition, there are polynomial-time algorithms for deciding whether a given allocation satisfies IGE on each of the domains we consider (e.g., Cechlárová et al. (2014); Aziz et al. (2019); Fujita et al. (2018)).

2.2.2 Participation guarantees

An allocation μ is *individually rational* at a preference profile P if, for each $i \in N$, $\mu_i \; R_i \; \omega_i$.

Definition 5. A rule φ satisfies *individual rationality (IR)* if, for each $P \in \mathcal{P}$, $\varphi(P)$ is individually rational at P.

Individual rationality is a participation guarantee that ensures no agent is made worse off by taking part in the reallocation. Our second participation guarantee restricts which individual objects may appear in an agent's assignment. An allocation μ satisfies the *worst-endowment lower bound* at a preference profile P if, for each $i \in N$ and each $o \in \mu_i$, $o \; R_i \; \min_{P_i}(\omega_i)$; that is, no agent is assigned an object that she ranks below her least-preferred endowed object.

Definition 6. A rule φ satisfies the *worst-endowment lower bound (WELB)* if, for each $P \in \mathcal{P}$, $\varphi(P)$ satisfies the worst-endowment lower bound at P.

Note that WELB is formulated purely in terms of agents' marginal preferences over individual objects, whereas IR is defined in terms of preferences over bundles.

IR and WELB coincide in the single-object environment, where each agent is endowed with a single object. In general multi-object problems, however, the two properties are logically independent. Intuitively, IR permits an agent to receive an individually unattractive object as part of a sufficiently desirable bundle, whereas WELB rules out such objects altogether. One can interpret WELB as giving agents veto power over any individual object owned by others: by ranking an object below every object in her endowment, an agent can ensure that she is never assigned that object.

Under additional structure, IR becomes strictly stronger than WELB. For example, on the responsive domain, any marginal and IR rule automatically satisfies BAL and WELB.

[14] A similar relaxation is "pair efficiency" (Ekici (2024)), which precludes only mutually beneficial *bilateral* trades.

Lemma 1. *On the responsive domain, if a marginal rule satisfies IR, then it satisfies BAL and WELB.*

A similar implication holds on any domain $\mathcal{P}^*$ that treats bundles that would violate an agent's WELB condition as "unacceptable" (see, e.g., Andersson et al. (2021), Biró et al. (2022a)).[15] We do not impose this additional restriction; instead we treat IR and WELB as independent properties in order to separate their respective roles in the results.

2.2.3 Incentive properties

Given a preference profile $P = (P_i)_{i \in N}$ and a preference relation P'_i for agent i, we denote by (P'_i, P_{-i}) the preference profile in which agent i reports P'_i and every other agent $j \in N \setminus \{i\}$ reports P_j. Given a rule φ, we say that agent i can *manipulate* φ at P by reporting P'_i if $\varphi_i(P'_i, P_{-i})\ P_i\ \varphi_i(P)$. A rule is *strategy-proof* if no agent can manipulate it by reporting any preference relation.

Definition 7. A rule φ satisfies *strategy-proofness (SP)* if, for each $P \in \mathcal{P}$, each $i \in N$, and each $P'_i \in \mathcal{P}_i$, $\varphi_i(P)\ R_i\ \varphi_i(P'_i, P_{-i})$.

Because PE, IR, and SP are incompatible once some agent owns more than one object (Sönmez (1999), Todo et al. (2014)), we will not insist on full SP. Instead, we relax SP by restricting the set of manipulation strategies under consideration. In particular, we focus on simple manipulations—subset-drop, drop, and truncation strategies—that act directly on agents' marginal preferences over individual objects and are especially salient when studying marginal rules (see, e.g., Biró et al. (2022a,b), Altuntaş et al. (2023)).

Subset-drop, drop, and truncation strategies. Fix an agent i and a preference relation $P_i \in \mathcal{P}_i$. A subset-drop strategy for P_i is a report whose marginal preference is obtained from $P_i|_O$ by "dropping" some subset $X \subseteq O \setminus \omega_i$ to the bottom, preserving the relative order within X and within $O \setminus X$. Any preference $P'_i \in \mathcal{P}_i$ with this marginal preference is called a subset-drop strategy for P_i. Formally, $P'_i \in \mathcal{P}_i$ is a *subset-drop strategy* for P_i if there exists $X \subseteq O \setminus \omega_i$ such that

(i) $P_i|_X = P'_i|_X$,

(ii) $P_i|_{O \setminus X} = P'_i|_{O \setminus X}$, and

(iii) for all $x \in X$ and all $y \in O \setminus X$, $y\ P'_i\ x$.

In this case, we say that P'_i (respectively $P'_i|_O$) is obtained from P_i (respectively $P_i|_O$) by *dropping* X. Let $\mathcal{S}_i(P_i)$ denote the set of all subset-drop strategies for P_i.

[15] Formally, consider a domain $\mathcal{P}^*$ such that, for each $i \in N$ and each $P_i \in \mathcal{P}^*_i$, any bundle that intersects $\{o \in O \mid \min_{P_i}(\omega_i)\ P_i\ o\}$ is worse than any bundle that does not. Clearly, any IR rule on $\mathcal{P}^*$ also satisfies WELB. Furthermore, if preferences P_i are assumed to be lexicographic (or responsive) when restricted to "acceptable" bundles (subsets of $\{o \in O \mid o\ R_i\ \min_{P_i}(\omega_i)\}$), then our characterization results go through unchanged on $\mathcal{P}^*$.

Within the class of subset-drop strategies, we distinguish two important subclasses. A subset-drop strategy P_i' is called a *drop strategy* for P_i if it is obtained by dropping a singleton subset $\{x\} \subseteq O \setminus \omega_i$, i.e., by moving a single object $x \in O \setminus \omega_i$ to the bottom of the marginal preferences. Let $\mathcal{D}_i(P_i)$ denote the set of all drop strategies for P_i.

A subset-drop strategy P_i' is called a *truncation strategy* for P_i if it is obtained by dropping a *tail subset* of the form $\{o \in O \setminus \omega_i \mid x \; R_i \; o\}$ for some $x \in O$, i.e., by moving all objects that are weakly worse than some object x to the bottom of the marginal preferences. Let $\mathcal{T}_i(P_i)$ denote the set of all truncation strategies for P_i.

It is convenient to single out those truncation strategies with a specified cutoff object. A truncation strategy P_i' is called a *truncation of* P_i *at* y if it is obtained by dropping the strict tail subset $\{o \in O \setminus \omega_i \mid y \; P_i \; o\}$.[16] In this case, there is no $z \in O \setminus \omega_i$ such that $y \; P_i' \; z \; P_i' \; \min_{P_i}(\omega_i)$.

In the single-object environment (Section 5), where ω_i is a singleton and thus $\omega_i = \min_{P_i}(\omega_i)$, a subset-drop strategy can be interpreted as "declaring some of the other agents' objects unacceptable", and a truncation strategy simply "shortens the list of acceptable objects" in the standard sense from two-sided matching (see, e.g., Mongell and Roth (1991); Roth and Rothblum (1999); Ehlers (2008)). This intuition extends to the multi-object environment, though objects ranked below $\min_{P_i}(\omega_i)$ need not be unacceptable in the usual sense.[17] Nevertheless, under WELB, a subset-drop strategy effectively "vetoes" some of the other agents' objects: no rule satisfying WELB ever assigns to agent i any object ranked below $\min_{P_i}(\omega_i)$.

Definition 8. A rule φ satisfies

- *truncation-proofness (TP)* if, for each $P \in \mathcal{P}$, each $i \in N$, and each $P_i' \in \mathcal{T}_i(P_i)$, we have $\varphi_i(P) \; R_i \; \varphi_i(P_i', P_{-i})$.
- *drop strategy-proofness (DSP)* if, for each $P \in \mathcal{P}$, each $i \in N$, and each $P_i' \in \mathcal{D}_i(P_i)$, we have $\varphi_i(P) \; R_i \; \varphi_i(P_i', P_{-i})$.
- *subset-drop strategy-proofness (SDSP)* if, for each $P \in \mathcal{P}$, each $i \in N$, and each $P_i' \in \mathcal{S}_i(P_i)$, we have $\varphi_i(P) \; R_i \; \varphi_i(P_i', P_{-i})$.

SDSP implies both TP and DSP, as it defends against the largest set of manipulation strategies. TP and DSP are logically independent in general.

Example 1. Let $O = \{a, b, c, d, e, x, y\}$, $\omega_i = \{x, y\}$, and $P_i \in \mathcal{P}_i$ be such that $P_i|_O$: $a, b, \underline{x}, c, d, \underline{y}, e$ (agent i's endowment is underlined for emphasis). Then:

- any $P_i^1 \in \mathcal{P}_i$ with $P_i^1|_O : b, \underline{x}, c, d, \underline{y}, e, a$ is a drop strategy for P_i, obtained by dropping object a. Note that P_i^1 is not a truncation strategy for P_i (e.g., the relations disagree

[16] Equivalently, P_i' is a truncation of P_i at y if (i) $P_i|_{\omega_i} = P_i'|_{\omega_i}$, (ii) $P_i|_{O \setminus \omega_i} = P_i'|_{O \setminus \omega_i}$, and (iii) for each $x \in O \setminus \omega_i$ with $y \; P_i \; x$, $\min_{P_i'}(\omega_i) \; P_i' \; x$.

[17] Our results continue to hold on related domains in which bundles containing objects ranked below $\min_{P_i}(\omega_i)$ are explicitly unacceptable; see footnote 15 and the surrounding discussion.

about the best object in $O \setminus \omega_i$);

- any $P_i^2, P_i^3 \in \mathcal{P}_i$ with $P_i^2|_O : a, b, \underline{x}, c, \underline{y}, d, e$ and $P_i^3|_O : a, \underline{x}, \underline{y}, b, c, d, e$, are truncation strategies for P_i, obtained by dropping the tail subsets $\{d, e\}$ and $\{b, c, d, e\}$, respectively. Neither P_i^2 nor P_i^3 is a drop strategy for P_i (however, each can be obtained from P_i via a sequence of drop strategies). $\diamond$

Remark 1. Fix $i \in N$ and $P_i \in \mathcal{P}_i$. For any object $x \in O \setminus \omega_i$, there exists a subset-drop strategy for P_i that ranks x at the top of agent i's marginal preferences over $O \setminus \omega_i$. The same reported preferences can be obtained by dropping, one at a time, the objects in $O \setminus \omega_i$ that agent i ranks above x. This observation simplifies proofs of characterizations of TTC based on DSP and SDSP (and on SP for the single-object environment). By contrast, truncation strategies preserve the ranking of objects in $O \setminus \omega_i$ (see footnote 16), so they do not permit this kind of "push-up" manipulation.

2.3 Top Trading Cycles

A trading cycle at ω (or simply a *trading cycle*[18]) is a cyclic sequence

$$C = (o_1, i_1, o_2, i_2, \ldots, i_{k-1}, o_k, i_k, o_{k+1} = o_1)$$

consisting of $k \geq 1$ distinct objects and k distinct agents such that, for each $\ell \in \{1, \ldots, k\}$, $o_\ell \in \omega_{i_\ell}$. Each object o_ℓ on C precedes (or "points to") its owner i_ℓ, and each agent i_ℓ on C precedes (or "points to") the object $o_{\ell+1}$. The sets of agents and objects on C are denoted by $N(C) = \{i_1, i_2, \ldots, i_k\}$ and $O(C) = \{o_1, o_2, \ldots, o_k\}$, respectively. An allocation μ is said to *execute* the trading cycle C if it assigns to each agent in $N(C)$ the object she points to on C; that is, for each $i_\ell \in N(C)$, $o_{\ell+1} \in \mu_{i_\ell}$.

We study a multi-object extension of Gale's Top Trading Cycles algorithm (Shapley and Scarf (1974)), which determines an allocation by executing a sequence of carefully chosen trading cycles at ω. At each step of this procedure, every agent points to her most-preferred unassigned object, and every unassigned object points to its owner. The resulting directed graph contains at least one trading cycle, and each agent involved in a trading cycle is assigned the object to which she points. All objects involved in a trading cycle are then removed. If unassigned objects remain, then the procedure continues to the next step; otherwise, it terminates with the resulting allocation.

We now formalize the *(generalized) Top Trading Cycles rule (TTC)*, which we denote by φ^{TTC}. Given a preference profile $P \in \mathcal{P}$, the allocation $\varphi^{\text{TTC}}(P)$ is determined by running the following *(generalized) TTC algorithm* at P. We denote this specific instance as $\text{TTC}(P)$.

[18]That is, a *trading cycle* refers to a trading cycle at the particular allocation $\mu = \omega$ (as defined in Section 2.2.1).

Algorithm: TTC(P)

Initialization: Set $\mu^0 := (\emptyset)_{i\in N}$ and $O^1 := O$.

Step $t \geq 1$: Each agent i points to $\max_{P_i}(O^t)$, and each object in O^t points to its owner. Let $\mathcal{C}_t(P)$ be the resulting set of trading cycles. For each $C \in \mathcal{C}_t(P)$, assign to each $i \in N(C)$ the object she points to in C; this yields the partial allocation μ^t with

$$\mu_i^t = \begin{cases} \mu_i^{t-1} \cup \{\max_{P_i}(O^t)\}, & \text{if } i \in \bigcup_{C\in\mathcal{C}_t(P)} N(C), \\ \mu_i^{t-1}, & \text{otherwise.} \end{cases}$$

Remove all objects involved in trading cycles to form the set $O^{t+1} = O^t \setminus \bigcup_{C\in\mathcal{C}_t(P)} O(C)$. Proceed to step $t+1$ if $O^{t+1} \neq \emptyset$; proceed to Termination otherwise.

Termination: Because at least one object is removed at each step, the algorithm terminates at some step T. Return the allocation $\varphi^{\mathrm{TTC}}(P) := \mu^T$.

Example 2 illustrates the TTC algorithm with a simple three-agent problem.

Example 2. Suppose $N = \{1, 2, 3\}$, $O = \{a, b, c, d\}$, and $\omega = (\{a, b\}, \{c\}, \{d\})$. Consider a preference profile P with marginal preferences $P_1|_O : c, a, d, b$, $P_2|_O : a, b, c, d$, and $P_3|_O : a, c, b, d$. The algorithm TTC(P) operates as follows.

Step 1: Each agent points to her most-preferred object in O: agent 1 points to c, while agents 2 and 3 point to a. Each object in O points to its owner. There is a trading cycle $C_1(P) = (c, 2, a, 1, c)$, which is executed to yield partial allocation $\mu^1 = (\{c\}, \{a\}, \emptyset)$ and the remaining objects $O^2 = \{b, d\}$.

Step 2: Each agent points to her most-preferred object in $O^2 = \{b, d\}$: agent 1 points to d, while agents 2 and 3 point to b. Each object in O^2 points to its owner. There is a trading cycle $C_2(P) = (d, 3, b, 1, d)$, which is executed to yield the final allocation $\mu^2 = (\{c, d\}, \{a\}, \{b\})$ and the remaining objects $O^3 = \emptyset$.

The algorithm therefore terminates with the allocation $\varphi^{\mathrm{TTC}}(P) = (\{c, d\}, \{a\}, \{b\})$. ◇

The following fact is immediate from the definition of φ^{TTC}.

Fact 1. *On any domain of strict preferences, φ^{TTC} is marginal and satisfies BAL and WELB.*

Beyond these basic properties, the behavior of φ^{TTC} depends on the preference domain.

3 Lexicographic preferences

The lexicographic domain $\mathcal{L}$ is a natural starting point for our analysis because of its tractability. As discussed in Sections 2.1–2.2, this domain has two convenient features. First, each $P_i \in \mathcal{L}_i$ is uniquely determined by its marginal preference $P_i|_O$, so any rule defined on $\mathcal{L}$ is automatically marginal. Second, on $\mathcal{L}$, PE coincides with IGE.

On the lexicographic domain, φ^{TTC} behaves especially well. Fujita et al. (2018) show that it always selects an allocation in the core of the associated exchange economy, and hence it satisfies both PE and IR. Although φ^{TTC} necessarily fails SP,[19] Altuntaş et al. (2023) prove that it satisfies SDSP; thus it also satisfies TP and DSP.[20] These observations yield the following fact.

Fact 2. *On the lexicographic domain, φ^{TTC} satisfies PE, IR, and SDSP.*

Facts 1 and 2 together imply that φ^{TTC} satisfies BAL, IGE, WELB, and TP. Theorem 1 shows that these four properties characterize φ^{TTC} on $\mathcal{L}$.

Theorem 1. *On the lexicographic domain, a rule satisfies BAL, IGE, WELB, and TP if and only if it equals φ^{TTC}.*

The proof proceeds by minimal counterexample, using a novel minimality criterion that combines a "size" function (Sethuraman (2016)) and a "similarity" function (Ekici (2024)). Suppose φ satisfies the properties but differs from φ^{TTC}. Call $P \in \mathcal{L}$ a *conflict profile* if $\varphi(P) \neq \varphi^{\mathrm{TTC}}(P)$. For each conflict profile P, define its similarity $\rho(P)$ as the earliest step t of $\mathrm{TTC}(P)$ such that $\varphi(P)$ does not execute all trading cycles in $\mathcal{C}_t(P)$. Among all conflict profiles that minimize ρ, choose a profile that further minimizes the size function $s(P) \coloneqq \sum_{i \in N} |\{o \in O \mid o \ R_i \ \min_{P_i}(\omega_i)\}|$. At the chosen conflict profile P, the deviation from $\varphi^{\mathrm{TTC}}(P)$ forces a Pareto-improving trading cycle at $\varphi(P)$, contradicting IGE (see Appendix B).

Theorem 1 characterizes φ^{TTC} using TP as its incentive requirement. Our next theorem shows that TP can be replaced with the alternative requirement DSP. The key observation is that DSP becomes strictly stronger once WELB is imposed: on the lexicographic domain, DSP together with WELB implies SDSP, and hence TP.[21]

Lemma 2. *On the lexicographic domain, if a rule satisfies DSP and WELB, then it satisfies SDSP.*

Theorem 2 refines two characterizations provided by Altuntaş et al. (2023, Theorems 1 and 3).

Theorem 2. *On the lexicographic domain, a rule satisfies BAL, IGE, WELB, and DSP if and only if it equals φ^{TTC}.*

Proof. By Facts 1 and 2, φ^{TTC} satisfies the stated properties. Conversely, if φ is a rule satisfying

[19] PE, IR, and SP are incompatible even for lexicographic preferences (Todo et al. (2014)).

[20] Altuntaş et al. (2023) refer to SDSP as "subset total drop strategy-proofness". Their notion of drop strategy-proofness is stronger than ours, as it also precludes manipulations obtained by dropping objects an agent already owns. Throughout, we use the weaker DSP as it suffices for our characterization results.

[21] TP is not implied by DSP alone. Conversely, DSP is not implied by any proper subset of the properties in Theorem 1.

BAL, IGE, WELB, and DSP, then Lemma 2 implies that φ satisfies SDSP and hence TP. Therefore, φ satisfies the properties in Theorem 1, which yields $\varphi = \varphi^{\mathrm{TTC}}$. $\square$

Note that IR cannot replace WELB in the statement of Theorem 1 or 2.[22] Example 9 (Appendix C) exhibits a rule $\varphi^{\neg\mathrm{WELB}}$ on $\mathcal{L}$ that differs from φ^{TTC} yet also satisfies BAL, PE, IR, TP, and DSP. The construction relies on a key distinction between the two participation guarantees: IR allows an agent to receive an individually unattractive object as part of a sufficiently desirable bundle, whereas WELB prevents such objects from appearing in her assignment. Interestingly, the analogous construction is not possible with marginal rules on the responsive domain; in that setting, IR suffices for the characterization (see Lemma 1 and Theorem 4).[23]

More generally, Appendix C provides several alternative rules that demonstrate the independence of our properties. Table 1 summarizes the properties satisfied by these rules on the lexicographic domain; it shows that each of the properties in Theorems 1 and 2 is indispensable.

Table 1: Properties of selected rules on $\mathcal{L}$

Rule	BAL	PE	WELB	TP	DSP	IR	SP
φ^{TTC}	✓	✓	✓	✓	✓	✓	
No-trade rule (Ex. 5)	✓		✓	✓	✓	✓	✓
Balanced Serial Dictatorship (Ex. 6)	✓	✓		✓	✓		✓
$\varphi^{\neg\mathrm{TP}}$ (Ex. 7)	✓	✓	✓			✓	
$\varphi^{\neg\mathrm{BAL}}$ (Ex. 8)		✓	✓	✓	✓		
$\varphi^{\neg\mathrm{WELB}}$ (Ex. 9)	✓	✓		✓	✓	✓	

4 Responsive preferences

We now turn from lexicographic to responsive preferences. Objects remain substitutes, but the tension between efficiency, individual rationality, and incentive compatibility becomes even starker on this more general domain.

On the responsive domain, PE becomes a demanding benchmark: Pareto improvements may require coordinating intricate multi-object exchanges among several agents, and even deciding whether a Pareto improvement exists is NP-complete (De Keijzer et al. (2009), Aziz et al. (2019)). Furthermore, Manjunath and Westkamp (2025) show that no marginal rule can simultaneously satisfy PE and IR. Their argument is based on the following simple two-agent example.

[22] Of course, IR can be *added* to the list of properties but it is *implied* by the four properties in each theorem.

[23] Similarly, this construction is not possible on the domain $\mathcal{P}^*$ described in footnote 15, as IR implies WELB on that domain.

Example 3. Let $N = \{1, 2\}$, $O = \{a, b, c, d\}$, and $\omega = (\{a, d\}, \{b, c\})$. Let φ be a rule on $\mathcal{R}$ satisfying PE and IR.

Consider $P \in \mathcal{L}$ with $P_1|_O = P_2|_O : a, b, c, d$. Since φ satisfies PE and IR, it must assign $\varphi(P) = \omega$. Now consider $P' \in \mathcal{R}$ with $P'|_O = P|_O$, and assume that each agent prefers the other's endowment: $\{b, c\}\ P'_1\ \{a, d\}$ and $\{a, d\}\ P'_2\ \{b, c\}$. Then ω is Pareto dominated at P' by the bundle swap $\mu := (\{b, c\}, \{a, d\})$. Since φ satisfies PE, we must have $\varphi(P') \neq \omega = \varphi(P)$. Thus $P'|_O = P|_O$ and $\varphi(P') \neq \varphi(P)$, so φ is not marginal. $\diamond$

At the lexicographic profile P in Example 3, any rule on $\mathcal{R}$ that satisfies PE and IR when restricted to $\mathcal{L}$ must select the endowment $\omega = (\{a, d\}, \{b, c\})$. If the rule is marginal, then it must also select ω at the responsive profile P' with the same marginal preferences. In particular, φ^{TTC} assigns ω at both P and P', and thus $\varphi^{\text{TTC}}(P') = \omega$ is Pareto dominated at P'. However, no trading cycle is Pareto improving at P' (under the common marginal preferences, a single-object exchange would require some agent to trade away an object for a lower-ranked one).

More generally, φ^{TTC} inherits individual-good efficiency on $\mathcal{R}$ from its Pareto efficiency on $\mathcal{L}$: if a Pareto-improving trading cycle existed at $\varphi^{\text{TTC}}(P)$ for some $P \in \mathcal{R}$, then the same cycle would be Pareto improving at the lexicographic profile with the same marginal preferences, contradicting the Pareto efficiency of φ^{TTC} on $\mathcal{L}$. In addition, Altuntaş et al. (2023) show that $\mathcal{R}$ is a maximal domain on which φ^{TTC} satisfies IR. The next fact collects these properties.

Fact 3. *On the responsive domain, φ^{TTC} satisfies IGE and IR.*

Incentive compatibility becomes even more fragile on the responsive domain because extending a rule to a larger domain preserves any manipulations that were already possible and may introduce new ones. Nevertheless, φ^{TTC} remains robust against truncation strategies on $\mathcal{R}$.

Proposition 1. *On the responsive domain, φ^{TTC} satisfies TP.*

Combining Fact 1, Fact 3, and Proposition 1, we see that φ^{TTC} satisfies BAL, IGE, WELB, and TP on $\mathcal{R}$. Within the class of marginal rules, these properties continue to identify φ^{TTC} on the larger domain $\mathcal{R}$ (cf. Theorem 1). In fact, in Theorem 3 below we show that the uniqueness holds more generally. The key observation is the following lemma, which says that BAL is superfluous: for marginal rules on $\mathcal{R}$, it is implied by WELB and TP.

Lemma 3. *On the responsive domain, if a marginal rule satisfies WELB and TP, then it satisfies BAL.*

Theorem 3. *On the responsive domain, a marginal rule satisfies IGE, WELB, and TP if and only if it equals φ^{TTC}.*

Proof. We have shown that φ^{TTC} satisfies the stated properties. Conversely, let φ be a marginal rule on $\mathcal{R}$ satisfying IGE, WELB, and TP. By Lemma 3, φ also satisfies BAL. Hence Theorem 1 implies that φ coincides with φ^{TTC} on $\mathcal{L}$.

Fix any $P' \in \mathcal{R}$, and let P be the unique lexicographic preference profile with $P|_O = P'|_O$. Since φ and φ^{TTC} are marginal, we have

$$\varphi(P') = \varphi(P) = \varphi^{\mathrm{TTC}}(P) = \varphi^{\mathrm{TTC}}(P').$$

Thus φ coincides with φ^{TTC} on $\mathcal{R}$.[24] □

Unlike on the lexicographic domain, where IR cannot generally replace WELB (see Example 9), on $\mathcal{R}$ any marginal rule satisfying IR also satisfies WELB (Lemma 1). Thus WELB can be replaced by IR in Theorem 3.

Theorem 4. *On the responsive domain, a marginal rule satisfies IGE, IR, and TP if and only if it equals φ^{TTC}.*

Proof. We have shown that φ^{TTC} satisfies the stated properties. Conversely, if φ is a marginal rule satisfying IGE, IR, and TP, then Lemma 1 implies that φ also satisfies WELB. Therefore, φ satisfies the properties in Theorem 3, and hence $\varphi = \varphi^{\mathrm{TTC}}$. □

Theorems 3 and 4 characterize φ^{TTC} on $\mathcal{R}$ using TP as the incentive requirement. The next example shows that φ^{TTC} fails the alternative requirement DSP on $\mathcal{R}$.

Example 4. Let $N = \{1, 2, 3\}$ and $\omega = (\{a, b\}, \{c\}, \{d\})$. Let P_2 and P_3 be lexicographic preferences with $P_2|_O : d, b, c, a$ and $P_3|_O : b, a, c, d$.

Let P_1 be a responsive preference with $P_1|_O : d, b, c, a$ and $\{b, c\}\ P_1\ \{a, d\}$. Then φ^{TTC} assigns $\varphi^{\mathrm{TTC}}(P) = (\{a, d\}, \{c\}, \{b\})$. Now suppose agent 1 drops object d, reporting some P_1' with the marginal preference $P_1'|_O : b, c, a, d$. Then φ^{TTC} assigns $\varphi^{\mathrm{TTC}}(P_1', P_{-1}) = (\{b, c\}, \{d\}, \{a\})$. Because $\{b, c\}\ P_1\ \{a, d\}$, agent 1 benefits from the drop strategy. ⋄

Example 4 suggests that DSP is a rather restrictive requirement on $\mathcal{R}$. Combined with our lexicographic-domain characterization (Theorem 2) and the fact that IR implies WELB for marginal rules on $\mathcal{R}$ (Lemma 1), this observation yields a sharp impossibility.

Corollary 1. *On the responsive domain, no marginal rule satisfies IGE, IR, and DSP.*

Proof. Toward contradiction, let φ be a marginal rule on $\mathcal{R}$ satisfying IGE, IR, and DSP. Then φ also satisfies BAL and WELB by Lemma 1. Hence Theorem 2 implies that φ coincides with φ^{TTC} on $\mathcal{L}$, and by marginality it must coincide with φ^{TTC} on $\mathcal{R}$. Thus φ fails DSP, a contradiction. □

[24] The argument parallels one used by Biró et al. (2022a) to extend a characterization of the "Circulation Top Trading Cycle" rule from the lexicographic to the responsive domain in a model with homogeneous objects.

Corollary 1 shows that any marginal rule on $\mathcal{R}$ that guarantees IGE and IR is necessarily manipulable—even among drop strategies.[25] Nevertheless, we show that φ^{TTC} retains some appealing incentive qualities beyond TP. In particular, it is *not obviously manipulable* according to the typology proposed by Troyan and Morrill (2020). Roughly speaking, although profitable deviations exist, no deviation is clearly beneficial when the agent compares only the best- and worst-case outcomes that can result from truthful reporting and from the deviation.

For each $i \in N$ and each $P_i \in \mathcal{R}_i$, agent i's *opportunity set* at P_i under φ^{TTC} is the set $\mathcal{O}_i(P_i) = \{\varphi_i^{\mathrm{TTC}}(P_i, P^*_{-i}) \mid P^*_{-i} \in \mathcal{P}_{-i}\}$ of all bundles that i could be assigned upon reporting P_i. For each nonempty subset $\mathcal{X} \subseteq 2^O$, $B_{P_i}(\mathcal{X})$ denotes the most-preferred bundle in $\mathcal{X}$ according to P_i, i.e., $B_{P_i}(\mathcal{X}) \in \mathcal{X}$ and $B_{P_i}(\mathcal{X})\ R_i\ X$ for each $X \in \mathcal{X}$. Similarly, $W_{P_i}(\mathcal{X})$ denotes the least-preferred bundle in $\mathcal{X}$ according to P_i.

Proposition 2. *On the responsive domain, φ^{TTC} is* not obviously manipulable*.*
More precisely, for each $P \in \mathcal{R}$, each $i \in N$, and each $P'_i \in \mathcal{R}_i$ with $\varphi_i^{\mathrm{TTC}}(P'_i, P_{-i})\ P_i\ \varphi_i^{\mathrm{TTC}}(P)$,

(i) $B_{P_i}(\mathcal{O}_i(P_i))\ R_i\ B_{P_i}(\mathcal{O}_i(P'_i))$, *and*

(ii) $W_{P_i}(\mathcal{O}_i(P_i))\ R_i\ W_{P_i}(\mathcal{O}_i(P'_i))$.

Table 2 summarizes the properties satisfied by several alternative rules on $\mathcal{R}$ (see the constructions in Appendix C). It shows that the properties used in Theorems 3 and 4, including marginality, are logically independent.

Table 2: Properties of selected rules on $\mathcal{R}$

Rule	BAL	IGE	WELB	TP	DSP	IR	MAR
φ^{TTC}	✓	✓	✓	✓		✓	✓
No-trade rule (Ex. 5)	✓		✓	✓	✓	✓	✓
Balanced Serial Dictatorship (Ex. 6)	✓	✓		✓	✓		✓
$\varphi^{\neg\mathrm{TP}}$ (Ex. 7)	✓	✓	✓			✓	✓
$\varphi^{\neg\mathrm{MAR}}$ (Ex. 10)	✓	✓	✓	✓		✓	

5 The single-object environment

We now specialize our analysis to the classic single-object environment (the "housing market") introduced by Shapley and Scarf (1974). In this environment our properties simplify, and the multi-object characterization in Theorem 1 collapses to a sharper statement.

We formalize the single-object environment as follows. Assume that $O = \{o_1, \ldots, o_n\}$ and that each agent $i \in N$ is endowed with the object $\omega_i = o_i$. Each agent i has strict preferences P_i

[25] A similar argument yields an analogous impossibility result with IR replaced by BAL and WELB: no marginal rule on $\mathcal{R}$ satisfies IGE, BAL, WELB, and DSP.

over individual objects in O, and $\mathcal{P} = \prod_{i\in N} \mathcal{P}_i$ denotes the domain of strict preference profiles.[26] An allocation is represented as a bijection $\mu : N \to O$, where μ_i denotes the object assigned to agent i.

In the single-object environment, every allocation is balanced by definition. Moreover, IR and WELB coincide because each agent must be assigned exactly one object, and PE and IGE coincide because any Pareto improvement can be realized by executing a collection of disjoint trading cycles. Truncation and drop strategies—as well as the corresponding incentive properties, TP and DSP—are defined exactly as in Section 2.2.3. In this setting, truncation strategies admit a familiar interpretation from the matching literature: a truncation strategy for P_i simply shortens agent i's list of acceptable objects by moving the "outside option" ω_i up in the ranking while preserving the ordering of objects in $O \setminus \omega_i$ (see footnote 16 and the surrounding discussion).

Theorem 5. *In the single-object environment, a rule satisfies PE, IR, and TP if and only if it equals φ^{TTC}.*

The proof closely parallels that of Theorem 1 and is therefore omitted.

Theorem 5 refines several characterizations of φ^{TTC} in housing markets. In particular, Ma (1994) shows that φ^{TTC} is the unique rule satisfying PE, IR, and SP, and Altuntaş et al. (2023) prove that SP can be replaced with either DSP or the weaker "upper invariance".[27] In the single-object environment, upper invariance together with IR implies TP; hence Theorem 5 establishes the same uniqueness under weaker criteria.

The theorem also complements characterizations using relaxed efficiency requirements. For example, Ekici (2024) replaces PE with "pair efficiency" while retaining IR and SP, and Chen et al. (2025) further weaken SP to upper invariance. Theorem 5 does not permit a similar refinement: pair efficiency together with IR and TP does not characterize φ^{TTC} (see Coreno and Feng (2025, Example 1)).

6 Extension: Conditionally lexicographic preferences

Conditionally lexicographic preferences generalize purely lexicographic preferences (Booth et al. (2010); see also Domshlak et al. (2011); Pigozzi et al. (2016)). Unlike responsive preferences, they allow the relative ranking of two objects to depend on the other objects they are obtained with: for example, an agent may prefer drinking Champagne to Bordeaux when paired with oysters, but Bordeaux to Champagne otherwise. This flexibility accommodates complementarities among

[26]In this section only, we adopt the standard convention that each agent has preferences P_i defined directly on the set of objects O. This is only a notational simplification. If instead we maintained the assumption that each agent has lexicographic preferences on 2^O, then the result below would be an immediate corollary of Theorem 1.

[27]Upper invariance requires an agent's assignment to remain unchanged when she misrepresents only the ordering of objects ranked below her assignment at the truthful profile.

objects while retaining several appealing features of lexicographic preferences. In particular, conditionally lexicographic preferences admit a compact representation via *lexicographic preference trees (LP trees)*, which makes them attractive from an implementation perspective.[28] As we shall see, PE coincides with IGE on this domain.

Loosely speaking, an agent has conditionally lexicographic preferences if, for any nonempty set of objects X disjoint from Y, there is a unique object in X which is the "lexicographically best" addition to Y from X.

Definition 9. A preference relation P_i on 2^O is *conditionally lexicographic* if, for all disjoint $X, Y \in 2^O$ with $X \neq \emptyset$, there is a unique object in X, denoted $x^* = \max_{P_i}(X \mid Y)$, such that

$$\text{for each } Z \subseteq X \setminus \{x^*\}, \quad (Y \cup x^*) \; P_i \; (Y \cup Z).$$

We call $x^* = \max_{P_i}(X \mid Y)$ agent i's most-preferred object in X conditional on already having Y.

For each $i \in N$, let $\mathcal{CL}_i$ be the set of conditionally lexicographic preferences on 2^O, and write $\mathcal{CL} \coloneqq \prod_{i \in N} \mathcal{CL}_i$ for the *conditionally lexicographic domain*. Every $P_i \in \mathcal{CL}_i$ is monotonic. Moreover, every lexicographic preference is conditionally lexicographic, and a preference is lexicographic if and only if it is both conditionally lexicographic and responsive. Formally, for each $i \in N$, $\mathcal{L}_i \subseteq \mathcal{CL}_i \subseteq \mathcal{M}_i$ (with both inclusions strict if $|O| \geq 3$), and $\mathcal{L}_i = \mathcal{CL}_i \cap \mathcal{R}_i$.

6.1 Properties: Conditional WELB and DSP

In this section we look beyond marginal rules and consider rules that depend on agents' full preferences over bundles. Among the properties we use below, the definitions of BAL, PE, IGE, and IR carry over from Section 2.2 without modification. By contrast, WELB and DSP, which were originally defined with respect to agents' marginal preferences, must be reformulated using the conditional marginal preferences introduced below.[29]

Given $P_i \in \mathcal{CL}_i$ and a bundle $Y \subseteq O$, we define the *conditional marginal preference* $P_i(Y)|_O$ over individual objects *conditional on receiving* Y. Formally, $P_i(Y)|_O$ is the strict linear order on O such that, for all $x, y \in O$,[30]

$$x \; P_i(Y)|_O \; y \iff (Y \cup \{x\}) \setminus \{y\} \; P_i \; (Y \cup \{y\}) \setminus \{x\}.$$

Equivalently, when P_i is represented by an LP tree, $P_i(Y)|_O$ is the order in which the objects

[28]Intuitively, an LP tree is a rooted binary tree that represents conditional marginal preferences in a graphical manner. For any bundle Y, there is a unique root-to-leaf path that is consistent with Y. This path specifies the agent's preference ordering over objects conditional on receiving Y. See Appendix A for details.

[29]We do not discuss truncation-proofness on $\mathcal{CL}$ because the corresponding definition becomes rather unwieldy.

[30]Equivalently, for all $x, y \in O$, $x \; P_i(Y)|_O \; y \iff x \cup (Y \setminus \{x, y\}) \; P_i \; y \cup (Y \setminus \{x, y\})$.

are encountered along the unique path consistent with Y (see Appendix A). We define $R_i(Y)|_O$ by $x\ R_i(Y)|_O\ y$ if and only if ($x\ P_i(Y)|_O\ y$ or $x = y$). For notational convenience, we will often write $P_i(Y)$ instead of $P_i(Y)|_O$ (and similarly $R_i(Y)$ instead of $R_i(Y)|_O$). Given a nonempty subset $X \subseteq O$, let $\max_{P_i(Y)}(X)$ and $\min_{P_i(Y)}(X)$ denote the most- and least-preferred objects in X according to $P_i(Y)|_O$, respectively. Note that when X and Y are disjoint, $\max_{P_i(Y)}(X)$ coincides with $\max_{P_i}(X \mid Y)$.

Remark 2. Each conditionally lexicographic P_i induces a family $(P_i(Y)|_O)_{Y \subseteq O}$ of conditional marginal preferences. However, not every family of marginal preferences arises from a conditionally lexicographic preference: Definition 9 implies that the family $(P_i(Y)|_O)_{Y \subseteq O}$ must satisfy certain consistency constraints. For example, any two conditional marginal preferences $P_i(X)|_O$ and $P_i(Y)|_O$ must top-rank the same object. The LP-tree representation in Appendix A spells out these constraints precisely.

An allocation μ satisfies the *worst-endowment lower bound* at a preference profile $P \in \mathcal{CL}$ if, for each $i \in N$, and each $o \in \mu_i$, $o\ R_i(\mu_i) \min_{P_i(\mu_i)}(\omega_i)$. In other words, *conditional on receiving* μ_i, every object assigned to agent i is at least as good, according to $P_i(\mu_i)|_O$, as her worst endowed object.[31]

Definition 10. A rule φ satisfies the *worst-endowment lower bound (WELB)* if, for each $P \in \mathcal{CL}$, $\varphi(P)$ satisfies the worst-endowment lower bound at P.

Drop strategies extend naturally to conditionally lexicographic preferences. An agent implements a drop strategy by dropping an object she does not own to the bottom of her conditionally lexicographic preferences. Formally, given $P_i \in \mathcal{CL}_i$, we say that $P'_i \in \mathcal{CL}_i$ is a *drop strategy* for P_i if there exists $x \in O \setminus \omega_i$ such that

(i) for each nonempty $Y \subseteq O \setminus \{x\}$, $Y\ P'_i\ x$, and

(ii) for all $Y, Z \subseteq O \setminus \{x\}$, $Y\ P'_i\ Z$ if and only if $Y\ P_i\ Z$.

In this case, we say that P'_i is obtained from P_i by dropping object x.[32] Note that this definition is equivalent to the one in Section 2.2 when P_i is purely lexicographic. Let $\mathcal{D}_i(P_i)$ denote the set of all drop strategies for P_i.

Intuitively, if P'_i is obtained from P_i by dropping object x, then, for each $Y \subseteq O$, the conditional marginal preference $P'_i(Y)|_O$ is obtained from $P_i(Y \setminus \{x\})|_O$ by dropping object x. In particular, x is the worst object in every conditional marginal preference associated with P'_i. The LP-tree representation in Appendix A shows that this family of modified conditional

[31] WELB and IR are logically independent on $\mathcal{CL}$. If one adds the auxiliary assumption that, for each agent, any bundle that fails WELB is "unacceptable" (i.e., strictly worse than her endowment), then IR implies WELB on the resulting domain (see footnote 15). We do not impose this additional restriction.

[32] As in the case of purely lexicographic preferences (but not responsive preferences), there is exactly one P'_i obtained from P_i by dropping x.

marginal preferences is induced by a unique conditionally lexicographic preference $P_i' \in \mathcal{CL}_i$.

Definition 11. A rule φ satisfies *drop strategy-proofness (DSP)* if, for each $P \in \mathcal{CL}$, each $i \in N$, and each $P_i' \in \mathcal{D}_i(P_i)$, we have $\varphi_i(P)\ R_i\ \varphi_i(P_i', P_{-i})$.

6.2 Augmented Top Trading Cycles

The *Augmented Top Trading Cycles rule (ATTC)*, denoted φ^{ATTC}, is the natural extension of φ^{TTC} from lexicographic to conditionally lexicographic preferences (Fujita et al. (2018)). At each step of the ATTC algorithm, each agent points to her most-preferred unassigned object *conditional on the objects already assigned to her*, and every unassigned object points to its owner. The resulting directed graph contains at least one trading cycle, and each agent involved in a trading cycle is assigned the object to which she points. All objects involved in a trading cycle are then removed. If unassigned objects remain, then the procedure continues to the next step; otherwise, it terminates with the resulting allocation.

Formally, given a preference profile $P \in \mathcal{CL}$, φ^{ATTC} returns the allocation $\varphi^{\mathrm{ATTC}}(P)$ determined by the following *ATTC algorithm* at P. We denote this specific instance as ATTC(P).

Algorithm: ATTC(P)

Initialization: Set $\mu^0 := (\emptyset)_{i\in N}$ and $O^1 := O$.

Step $t \geq 1$: Each agent i points to $\max_{P_i}(O^t \mid \mu_i^{t-1})$, and each object in O^t points to its owner. Let $\mathcal{C}_t(P)$ be the resulting set of trading cycles. For each $C \in \mathcal{C}_t(P)$, assign to each $i \in N(C)$ the object she points to in C; this yields the partial allocation μ^t with

$$\mu_i^t = \begin{cases} \mu_i^{t-1} \cup \{\max_{P_i}(O^t \mid \mu_i^{t-1})\}, & \text{if } i \in \bigcup_{C\in\mathcal{C}_t(P)} N(C), \\ \mu_i^{t-1}, & \text{otherwise.} \end{cases}$$

Remove all objects involved in trading cycles to form the set $O^{t+1} = O^t \setminus \bigcup_{C\in\mathcal{C}_t(P)} O(C)$. Proceed to step $t+1$ if $O^{t+1} \neq \emptyset$; proceed to Termination otherwise.

Termination: Because at least one object is removed at each step, the algorithm terminates at some step T. Return the allocation $\varphi^{\mathrm{ATTC}}(P) := \mu^T$.

6.3 Characterization of ATTC

Although conditionally lexicographic preferences allow for complementarities, the domain still behaves like the lexicographic domain from the perspective of efficiency. In particular, ruling out Pareto-improving single-object exchanges is enough to guarantee full Pareto efficiency.

Proposition 3. *On the conditionally lexicographic domain, an allocation satisfies IGE if and only if it satisfies PE.*

The equivalence between IGE and PE implies that any myopic procedure which, at each step, greedily executes a Pareto-improving trading cycle whenever one exists must terminate at a PE allocation on this domain. The ATTC algorithm is one such procedure.

Beyond efficiency, Fujita et al. (2018) show that φ^{ATTC} is core selecting (and hence satisfies both PE and IR). They also prove that it is NP-hard for an agent to find a profitable manipulation, and that any successful manipulation yields only limited gains: no manipulator can obtain an object better than her most-preferred object under truthful reporting. Complementing these results, we establish that φ^{ATTC} satisfies DSP on the conditionally lexicographic domain.

Proposition 4. *On the conditionally lexicographic domain, φ^{ATTC} satisfies DSP.*

Proposition 4 extends a result of Altuntaş et al. (2023), which shows that φ^{TTC} satisfies a strong drop strategy-proofness requirement on the lexicographic domain—one that also protects against manipulations where agents drop objects they own. We work with the weaker DSP (that considers only drops outside the endowment), which suffices for our characterization result. Moreover, φ^{ATTC} fails the stronger variant on the conditionally lexicographic domain.

Theorem 6. *On the conditionally lexicographic domain, a rule satisfies BAL, IGE, WELB, and DSP if and only if it equals φ^{ATTC}.*

In light of Proposition 3, Theorem 6 says that φ^{ATTC} is the unique rule satisfying BAL, PE, WELB, and DSP. Thus, the theorem effectively extends Theorem 2 from the lexicographic domain to the broader conditionally lexicographic domain, and it addresses open questions posed by Altuntaş et al. (2023, p. 167) and Fujita et al. (2018, p. 531).

Table 3 collects the key properties that φ^{TTC} and φ^{ATTC} enjoy on the three domains we consider.

Table 3: Properties of TTC and ATTC on various domains

Domain (Rule)	BAL	WELB	IR	PE	IGE	TP	DSP
$\mathcal{L}$ (TTC)	✓	✓	✓	✓	✓	✓	✓
$\mathcal{R}$ (TTC)	✓	✓	✓		✓	✓	
$\mathcal{CL}$ (ATTC)	✓	✓	✓	✓	✓	—	✓

Notes: A ✓ indicates that the rule satisfies the corresponding property on the given domain. On $\mathcal{L}$ and $\mathcal{CL}$, PE and IGE coincide. We do not study TP on $\mathcal{CL}$. ◇

We conclude this section by showing the difficulty in obtaining positive results in more general preference domains. On any domain of monotonic preferences that strictly includes the conditionally lexicographic domain, the equivalence between PE and IGE breaks down. Thus, on larger domains, one must look beyond single-object exchanges in order to achieve efficiency.[33]

Proposition 5. *Within the domain of monotonic preferences, the conditionally lexicographic domain is a maximal domain on which IGE and PE are equivalent.*
(More precisely, if $\mathcal{CL} \subsetneq \mathcal{P} \subseteq \mathcal{M}$, then there is a set N of agents, a preference profile $P \in \mathcal{P}$, and an allocation $\mu \in \mathcal{A}$ such that μ satisfies IGE but not PE at P.)

7 Related literature and discussion

Since its introduction for housing markets in the seminal paper by Shapley and Scarf (1974), TTC has become the subject of a large body of literature; see Morrill and Roth (2024) for a survey. For housing markets with strict preferences, TTC is the canonical rule: it selects the unique core allocation (Shapley and Scarf (1974); Roth and Postlewaite (1977)), it is group strategy-proof (Roth (1982); Bird (1984)), and it is the unique rule satisfying PE, IR, and SP (Ma (1994)). Theorem 5 strengthens Ma's (1994) characterization by replacing SP with the weaker incentive requirement TP, leaving little scope for alternative rules in housing markets; see Section 5 for further discussion.

Early work on multi-object reallocation consists mostly of negative results. On general preference domains, PE, IR, and SP are incompatible once some agent owns more than one object (Sönmez (1999)), and the impossibility persists even under lexicographic preferences (Todo et al. (2014)). Moreover, the core can be empty even under additive (thus responsive) preferences (Konishi et al. (2001)).

Recent studies obtain positive findings under restricted preference domains. For instance, PE, IR, and SP are compatible when agents' bundle preferences are based on "dichotomous" (Andersson et al. (2021)) or "trichotomous" (Manjunath and Westkamp (2021, 2025)) marginal preferences over objects (see also Han et al. (2024)). In addition, under conditionally lexicographic preferences the core is guaranteed to be nonempty, and ATTC selects an allocation in the core (Fujita et al. (2018)). Since the core need not be single-valued, our characterizations under (conditionally) lexicographic preferences provide a normative justification for core selection via (augmented) TTC (see Theorems 1, 2, and 6).

In the closest paper to ours, Altuntaş et al. (2023) consider the same model but focus on the lexicographic domain. They show that TTC is characterized by PE, DSP, and a property termed the *strong endowment lower bound*, which requires each agent's assignment to "pairwise dominate" her endowment according to her marginal preferences.[34] The strong endowment lower bound implies BAL and WELB on any preference domain, and under responsive preferences it

[33] If we broaden the definition of allocations to allow agents to be assigned empty bundles, then the conditionally lexicographic domain is a maximal domain *among all strict preferences* on which PE and IGE coincide.

[34] Formally, an allocation μ satisfies the *strong endowment lower bound* at P if, for each agent i, there is a bijection $f_i : \omega_i \to \mu_i$ such that, for each $o \in \omega_i$, $f_i(o)\ R_i\ o$.

also implies IR (and Pápai's (2003) "strong individual rationality"). Consequently, we establish the same uniqueness under substantially weaker criteria—replacing their strong endowment lower bound with BAL and WELB (Theorems 1 and 2), and substituting DSP with TP (Theorem 1). We also obtain characterizations on richer domains—namely, the responsive domain (where TTC fails PE and DSP but satisfies the meaningful compromises IGE and TP) and the conditionally lexicographic domain (where ATTC satisfies PE, IR, and DSP)—addressing open questions posed by Altuntaş et al. (2023, p. 167).

We prove the uniqueness of TTC with a minimal-counterexample argument, using a novel minimality condition that combines the size criterion of Sethuraman (2016) and the similarity criterion of Ekici (2024). Our proofs of Theorems 1 and 6 (and the omitted proof of Theorem 5) are minimal in the sense that they invoke each axiom exactly once, further illustrating how the general approach of Sethuraman (2016) and Ekici (2024) can be quite fruitful when paired with a well-chosen minimality criterion. Our proof of Proposition 4 is based on a modified ATTC algorithm that is outcome-equivalent to the original ATTC algorithm; a similar modification is used by Gonczarowski et al. (2023) to "expose" the strategy-proofness of TTC in single-object environments.

Related studies consider adaptations of TTC to multi-object reallocation problems with additional structure. Biró et al. (2022a,b) consider a multiple-copy Shapley-Scarf model, where each agent owns multiple copies of a homogeneous, agent-specific object (equivalently, every agent is indifferent between all objects with the same initial owner). Focusing on marginal rules for responsive preferences, Biró et al. (2022a) characterize a variant of TTC using subset-drop strategy-proofness, and they determine the capacity configurations under which this variant satisfies PE, IR, and SP. Similarly, Feng et al. (2024) characterize another variant of TTC for Moulin's (1995) multiple-type housing markets, where objects are partitioned into types and each agent consumes exactly one object of each type.

Our characterizations of TTC under responsive preferences complement characterizations of strategy-proof alternatives on the responsive domain. Strategy-proofness sharply restricts the admissible rules, leading to rather extreme rules under mild auxiliary axioms. *Sequential Dictatorships* are characterized by SP, PE, and "non-bossiness" (e.g., Pápai (2001); Ehlers and Klaus (2003); Hatfield (2009); Monte and Tumennasan (2015)), whereas *Segmented Trading Cycles rules* are characterized by SP, strong individual rationality, non-bossiness, and "trade sovereignty" (Pápai (2003, 2007)). Our characterizations identify TTC as the sensible middle ground: though it fails two of the three ideals on the responsive domain (PE and SP), it delivers meaningful guarantees along all three dimensions (IGE, IR/WELB, and TP).

We relax SP by ruling out restricted classes of simple manipulations: subset-drop, drop, and truncation strategies. This focus is consistent with computational results emphasizing the difficulty of manipulating TTC in practice (e.g., Fujita et al. (2018); Phan and Purcell (2022)). Drop strategies and their variants are studied by Altuntaş et al. (2023) for multi-object reallocation under lexicographic preferences. Truncation has roots in the literature on matching markets, where it corresponds to shortening a list of acceptable partners (Mongell and Roth (1991); Roth and Vande Vate (1991); Roth and Rothblum (1999); Ehlers (2008)). Truncation strategies also appear in matching with contracts (e.g., Afacan (2016); Hatfield et al. (2021)) and school-choice problems (Shirakawa (2025)), though our notion is closest to properties studied by Kojima (2013) for multi-object assignment and Biró et al. (2022a,b) for multi-object reallocation.

Finally, when relaxing marginality and considering more complex reporting languages, competitive-equilibrium approaches become natural benchmarks. Existence issues are typically avoided either by seeking equilibria in pseudomarkets for probability shares (e.g., Hylland and Zeckhauser (1979); Echenique et al. (2021, 2023)) or by accepting only approximate market clearing (e.g., Budish (2011); Nguyen and Vohra (2024); Jantschgi et al. (2025)). These procedures offer compelling efficiency and fairness properties on general preference domains, but they typically suffer from onerous reporting requirements and considerable computational complexity. Our study of ATTC under conditionally lexicographic preferences provides a practical alternative: its reporting language comprises LP-trees which are expressive but compact, and it delivers strong welfare guarantees using a procedure that is computationally tractable.

A Lexicographic preference trees

This appendix introduces lexicographic preference trees, which provide a compact representation of conditionally lexicographic preferences.

Definition 12. A *lexicographic preference tree (LP tree)* on a nonempty subset $X \subseteq O$ is a directed binary tree τ_i such that

1. each vertex v is labeled with an object $o(v) \in X$.
2. every object in X appears exactly once on any path from the root to a leaf.
3. every internal (non-leaf) vertex v has two outgoing edges: (i) an "in edge" labeled $o(v)$, and (ii) a "not-in edge" labeled $\neg o(v)$.

An LP tree on O is referred to more succinctly as an *LP tree*.

Given an LP tree τ_i and a bundle $X \subseteq O$, let $\tau_i(X) = (v_1, v_2, \ldots, v_{|O|})$ be the unique root-to-leaf path in τ_i containing only edges consistent with X: for each $k = 1, \ldots, |O| - 1$, traverse the "in edge" from v_k if $o(v_k) \in X$, and the "not-in edge" from v_k if $o(v_k) \notin X$. Given bundles A and B with $A \neq B$, let $\tau_i(A, B)$ denote the last vertex common to both $\tau_i(A)$ and $\tau_i(B)$. Equivalently, $\tau_i(A, B)$ is the first vertex v on $\tau_i(A)$ and $\tau_i(B)$ with $o(v) \in (A \setminus B) \cup (B \setminus A)$. Figures 1 and 2 provide a graphical representation of two LP trees on $O = \{a, b, c, d\}$, together with the paths consistent with $\{a, b\}$ and $\{a, c\}$.

Definition 13. The preference relation P_{τ_i} associated with an LP tree τ_i on O is defined by:

1. for all $A, B \subseteq O$ with $A \neq B$, $[A \; P_{\tau_i} \; B \iff o(\tau_i(A, B)) \in A \setminus B]$.
2. for all $A, B \subseteq O$, $[A \; R_{\tau_i} \; B \iff (A = B \text{ or } A \; P_{\tau_i} \; B)]$.

A preference relation P_i on 2^O is conditionally lexicographic if and only if there exists an LP tree τ_i (on O) such that $P_i = P_{\tau_i}$. In fact, there is a one-to-one correspondence between LP trees and conditionally lexicographic preferences on O, and we denote the unique LP tree associated with P_i by τ_{P_i}. A lexicographic preference corresponds to an LP tree where all root-to-leaf paths use the same object order (i.e., all vertices at the same depth are labeled with the same object).

Worst-endowment lower bound. Given an LP tree τ_i and a vertex v of τ_{P_i}, let $a(v)$ denote the set of objects labeling the ancestors of v, including v itself. That is, if $(v_1, v_2, \ldots, v_k = v)$ is the unique path from the root of τ_i to v, then $a(v) = \{o(v_1), o(v_2), \ldots, o(v_k)\}$. Given a preference relation $P_i \in \mathcal{CL}_i$ and a bundle $X \in 2^O$, define $w_{P_i}(\omega_i \mid X)$ as the last vertex on the path $\tau_{P_i}(X)$ whose label belongs to ω_i. Equivalently, $w_{P_i}(\omega_i \mid X)$ is the unique vertex on $\tau_{P_i}(X)$ such that $o(w_{P_i}(\omega_i \mid X)) \in \omega_i$ and $\omega_i \subseteq a(w_{P_i}(\omega_i \mid X))$. We illustrate this construction graphically in figures 3 and 4.

An allocation μ satisfies the *worst-endowment lower bound* at a preference profile $P \in \mathcal{CL}$ if, for each $i \in N$, $\mu_i \subseteq a(w_{P_i}(\omega_i \mid \mu_i))$. This means every object in μ_i labels a vertex on $\tau_{P_i}(\mu_i)$

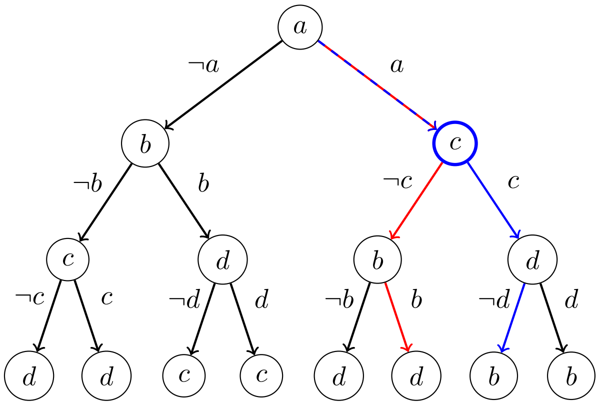

Figure 1: An LP tree τ_i

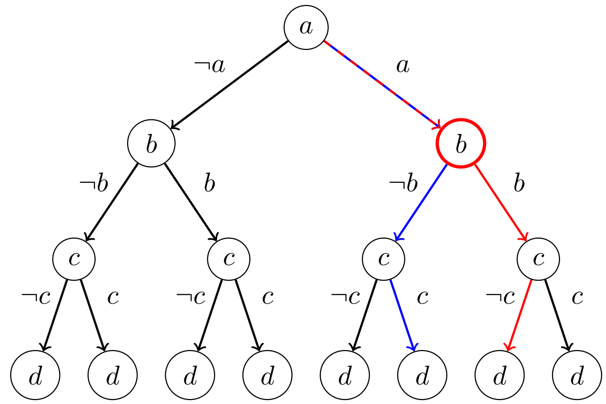

Figure 2: An LP tree τ_i^*

Notes: In both figures, the paths corresponding to the bundles $\{a, b\}$ and $\{a, c\}$ are highlighted in red and blue, respectively. In Figure 1, the last common vertex, $\tau_i(\{a, b\}, \{a, c\})$, is highlighted in blue. Because $o(\tau_i(\{a, b\}, \{a, c\})) = c$ belongs to $\{a, c\} \setminus \{a, b\}$, we have $\{a, c\}\ P_{\tau_i}\ \{a, b\}$. In Figure 2, the last common vertex, $\tau_i^*(\{a, b\}, \{a, c\})$, is highlighted in red. Because $o(\tau_i^*(\{a, b\}, \{a, c\})) = b$ belongs to $\{a, b\} \setminus \{a, c\}$, we have $\{a, b\}\ P_{\tau_i^*}\ \{a, c\}$. Since all paths in τ_i^* use the same object order, $P_{\tau_i^*}$ is purely lexicographic. ⋄

that appears before (or equals) vertex $w_{P_i}(\omega_i \mid \mu_i)$. Intuitively, conditional on receiving μ_i, no object assigned to agent i is worse than every object in her endowment ω_i. When agents have purely lexicographic preferences, this definition aligns with the original WELB.

Drop strategies. Given $P_i \in \mathcal{CL}_i$, let P_i' be obtained from P_i by dropping object x. Then P_i' is represented by the LP tree $\tau_{P_i'}$, obtained from τ_{P_i} by moving x to the bottom of every root-to-leaf path. Formally, let $\tau_{P_i}^{-x}$ denote the LP tree on $O \setminus \{x\}$ representing the restriction of P_i to subsets of $O \setminus \{x\}$.[35] To construct $\tau_{P_i'}$ from $\tau_{P_i}^{-x}$, at each leaf of $\tau_{P_i}^{-x}$ append two child nodes labeled with x, connected by an "in edge" and a "not-in edge" accordingly. This modification ensures that x is the least desirable addition to any bundle. Figures 5 and 6 illustrate this construction.

[35] We can construct $\tau_{P_i}^{-x}$ as follows. For each vertex v of τ_{P_i}, let T_v denote the maximal subtree of τ_{P_i} consisting of a vertex v of τ_{P_i} together with all of its successors, and let v' be the child of v whose incoming edge (v, v') is labeled with $\neg o(v)$. For each vertex v of τ_{P_i} such that $o(v) = x$, simply replace T_v with the subtree $T_{v'}$ (or the empty tree if v is a leaf of τ_{P_i}).

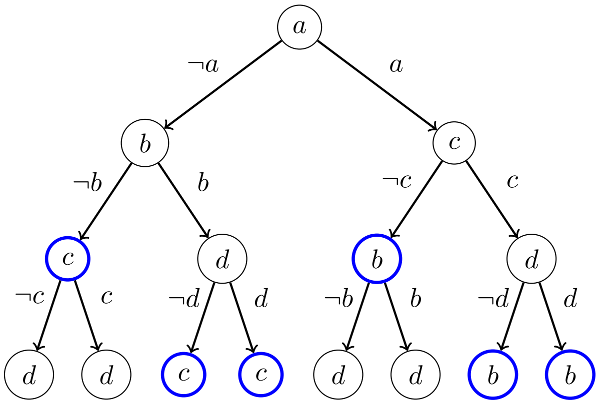

Figure 3: The LP tree τ_i

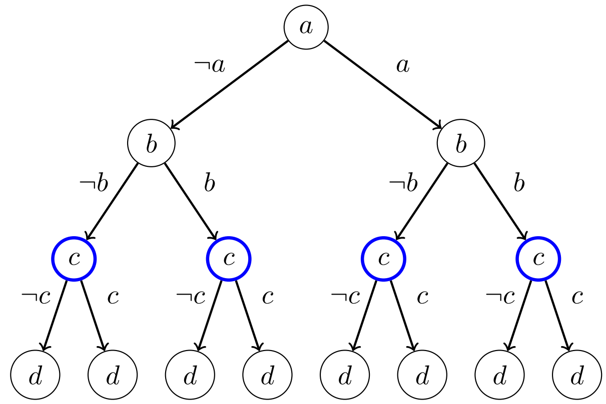

Figure 4: The LP tree τ_i^*

Notes: Figures 3 and 4 illustrate the construction of $w_{P_i}(\omega_i \mid X)$. Suppose agent i's endowment is $\omega_i = \{b, c\}$. In both figures, the vertices corresponding to $w_{P_i}(\omega_i \mid X)$ for various $X \in 2^O$ are highlighted in blue. Notice that on any path from the root to a leaf, there is exactly one such vertex. Under τ_i, WELB precludes assigning agent i object d together with any of the bundles $\emptyset, \{a\}, \{a, b\}$, or $\{c\}$; that is, she does not receive a bundle in $\{\{d\}, \{a, d\}, \{a, b, d\}, \{c, d\}\}$. Under τ_i^*, WELB precludes any bundle containing object d. $\diamond$

B Proofs

B.1 Proof of Lemma 1

Suppose φ is a marginal rule on $\mathcal{R}$ satisfying IR. Let $P \in \mathcal{R}$ and denote $\mu \coloneqq \varphi(P)$.

Toward contradiction, suppose that μ violates BAL at P. Then there exists an agent i with $|\mu_i| > |\omega_i|$. Because responsiveness imposes no restrictions on comparisons among bundles with different cardinalities (and it does not require monotonicity), there exists $P_i^* \in \mathcal{R}_i$ with $P_i^*|_O = P_i|_O$ such that, for each bundle $X \subseteq \omega_i$, $X \; P_i^* \; \mu_i$.[36] Then marginality of φ implies that $\varphi(P_i^*, P_{-i}) = \varphi(P) = \mu$. By construction of P_i^*, we have $\omega_i \; P_i^* \; \mu_i = \varphi_i(P_i^*, P_{-i})$, which contradicts IR. Thus φ satisfies BAL.

Now suppose for a contradiction that μ violates WELB at P. Then there is some agent $i \in N$ such that $\min_{P_i}(\omega_i) \; P_i \; \min_{P_i}(\mu_i)$. By BAL, we also have $|\mu_i| = |\omega_i|$. Consequently, there exists $P_i' \in \mathcal{R}_i$ with $P_i'|_O = P_i|_O$ and $\omega_i \; P_i' \; \mu_i$. However, marginality of φ implies that $\varphi_i(P_i', P_{-i}) = \varphi_i(P) = \mu_i$. This contradicts IR. ■

B.2 Proof of Lemma 2

Suppose φ satisfies DSP and WELB. Let P_i' be obtained from P_i by dropping some subset $X \subseteq O \setminus \omega_i$. Suppose that $X = \{x_1, x_2, \ldots, x_k\}$, where $x_1 \; P_i \; x_2 \; P_i \cdots P_i \; x_k$. Then P_i' is obtained

[36]For instance, let P_i^* order bundles first by cardinality (with smaller bundles preferred to larger bundles), and then order bundles with the same cardinality using the lexicographic order associated with $P_i|_O$. For example, if $O = \{a, b, c\}$ and $P_i|_O : a, b, c$, then $P_i^* : \emptyset, \{a\}, \{b\}, \{c\}, \{a, b\}, \{a, c\}, \{b, c\}, \{a, b, c\}$.

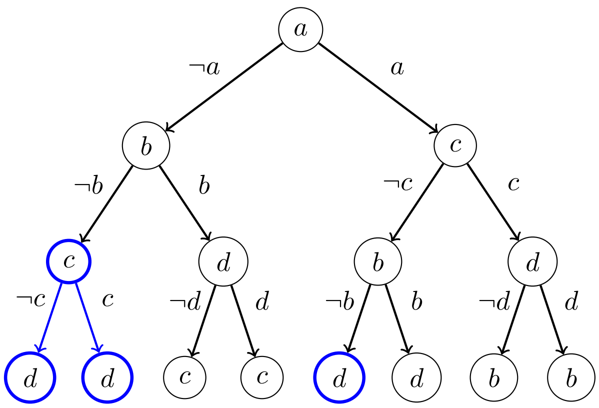

Figure 5: The LP tree τ_i

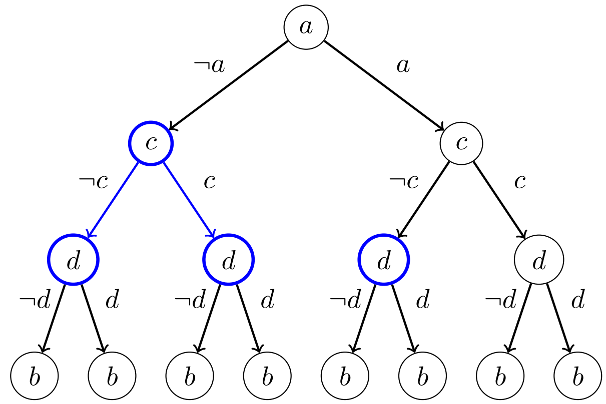

Figure 6: The LP tree τ_i' obtained by dropping b

Notes: Figure 5 displays the original LP tree from Figures 1 and 3. To illustrate the drop strategy where agent i drops object b, we construct the LP tree τ_i' shown in Figure 6. At each internal vertex v labeled b, we do the following: (i) remove the subtree consisting of v and all of its descendants, and (ii) append the subtree, shown in blue, consisting of the child of v reached via the "not-in edge" together with all of its descendants. Then, at each leaf of this modified tree, we append two child nodes labeled b, connected via an "in edge" and a "not-in edge" accordingly. $\diamond$

from P_i via a sequence of k drop strategies. That is, $P_i' = P_i^k$, where $P_i^0 = P_i$ and $P_i^1, \ldots, P_i^k$ are such that, for each $\ell \in \{1, \ldots, k\}$, P_i^ℓ is obtained from $P_i^{\ell-1}$ by dropping object x_ℓ.

Claim 1. For each $\ell = 1, \ldots, k$, $\varphi_i(P_i^{\ell-1}, P_{-i})\ R_i\ \varphi_i(P_i^\ell, P_{-i})$.

Proof. The proof is by induction on ℓ. Clearly, $\varphi_i(P) = \varphi_i(P_i^0, P_{-i})\ R_i\ \varphi_i(P_i^1, P_{-i})$ by DSP. For the inductive step, suppose that $\ell \in \{1, \ldots, k-1\}$ is such that

$$\varphi_i(P) = \varphi_i(P_i^0, P_{-i})\ R_i\ \varphi_i(P_i^1, P_{-i})\ R_i \cdots R_i\ \varphi_i(P_i^\ell, P_{-i}).$$

It suffices to show that $\varphi_i(P_i^\ell, P_{-i})\ R_i\ \varphi_i(P_i^{\ell+1}, P_{-i})$. By DSP, we have $\varphi_i(P_i^\ell, P_{-i})\ R_i^\ell\ \varphi_i(P_i^{\ell+1}, P_{-i})$. Moreover, WELB implies that

$$\varphi_i(P_i^\ell, P_{-i}) \subseteq O \setminus \{x_1, \ldots, x_\ell\} \quad \text{and} \quad \varphi_i(P_i^{\ell+1}, P_{-i}) \subseteq O \setminus \{x_1, \ldots, x_{\ell+1}\} \subseteq O \setminus \{x_1, \ldots, x_\ell\}.$$

Because $P_i^\ell|_{O\setminus\{x_1,\ldots,x_\ell\}} = P_i|_{O\setminus\{x_1,\ldots,x_\ell\}}$, and P_i and P_i^ℓ are lexicographic, we have $\varphi_i(P_i^\ell, P_{-i})\ R_i\ \varphi_i(P_i^{\ell+1}, P_{-i})$, as desired. $\square$

It follows from Claim 1 that

$$\varphi_i(P) = \varphi_i(P_i^0, P_{-i})\ R_i\ \varphi_i(P_i^1, P_{-i})\ R_i \cdots R_i\ \varphi_i(P_i^k, P_{-i}) = \varphi_i(P_i', P_{-i}).$$

Therefore, $\varphi_i(P)\ R_i\ \varphi_i(P_i', P_{-i})$. $\blacksquare$

B.3 Proof of Lemma 3

Suppose φ is a marginal rule on $\mathcal{R}$ satisfying WELB and TP. Let $P \in \mathcal{R}$ and denote $\mu \coloneqq \varphi(P)$.

Toward contradiction, suppose that $\mu \coloneqq \varphi(P)$ violates BAL at P. Then there is an agent $i \in N$ such that $|\mu_i| > |\omega_i|$. Because responsiveness imposes no restrictions on comparisons among bundles with different cardinalities (and it does not require monotonicity), there exists $P_i^* \in \mathcal{R}_i$ such that $P_i^*|_O = P_i|_O$ and, for each bundle $X \subseteq \omega_i$, $X \; P_i^* \; \mu_i$ (see footnote 36). Then marginality of φ implies that $\varphi(P_i^*, P_{-i}) = \varphi(P) = \mu$.

Now let $P_i' \in \mathcal{R}_i$ be a truncation of P_i^* obtained by dropping the set $O \setminus \omega_i$ (thus $P_i'|_O$ ranks every object in ω_i above all others). Then WELB implies that $\varphi_i(P_i', P_{-i}) \subseteq \omega_i$. By construction of P_i^*, we have $\varphi_i(P_i', P_{-i}) \; P_i^* \; \mu = \varphi_i(P_i^*, P_{-i})$. Thus φ fails TP, a contradiction. ■

B.4 Proof of Theorem 1

Toward contradiction, suppose that φ satisfies the properties but $\varphi \neq \varphi^{\text{TTC}}$. We will select a preference profile according to a particular minimality criterion, and then show that the minimality of the profile leads to a contradiction. We start by introducing some notation.

For each $P \in \mathcal{L}$ and each $t \in \mathbb{N}$, recall that $\mathcal{C}_t(P)$ denotes the set of trading cycles that arise at step t of TTC(P). (We assume that $\mathcal{C}_t(P) = \emptyset$ if TTC(P) terminates before step t.) Define the "size" function $s : \mathcal{L} \to \mathbb{N}$ such that, for each $P \in \mathcal{L}$,

$$s(P) = \sum_{i \in N} \left| \left\{ o \in O \mid o \; R_i \; \min_{P_i} (\omega_i) \right\} \right| .$$

The "similarity" function $\rho : \mathcal{L} \to \mathbb{N} \cup \{\infty\}$ measures the earliest point of divergence between φ and TTC(P). Formally, for each $P \in \mathcal{L}$,

$$\rho(P) = \inf\{t \in \mathbb{N} \mid \varphi(P) \text{ does not execute each trading cycle in } \mathcal{C}_t(P)\},$$

where it is understood that $\rho(P) = \infty$ precisely when $\varphi(P) = \varphi^{\text{TTC}}(P)$.

Among all profiles that minimize the similarity function ρ, let P be one that further minimizes the size function s. Thus, for any profile $P' \in \mathcal{L}$, either (i) $\rho(P) < \rho(P')$, or (ii) $\rho(P) = \rho(P')$ and $s(P) \leq s(P')$.

Let $t \coloneqq \rho(P)$. Then $\varphi(P)$ executes each trading cycle in $\bigcup_{\tau=1}^{t-1} \mathcal{C}_\tau(P)$, but it does not execute some trading cycle $C \in \mathcal{C}_t(P)$. Suppose that

$$C = (o_1, i_1, o_2, i_2, \ldots, o_k, i_k, o_{k+1} = o_1) .$$

Because $\varphi(P)$ does not execute C, there is an agent $i_\ell \in N(C)$ who is not assigned the object she points to within C, i.e., $o_{\ell+1} \notin \varphi_{i_\ell}(P)$. Without loss of generality, let $i_\ell = i_k$. Thus,

$o_{k+1} = o_1 \notin \varphi_{i_k}(P)$.

Let O^t denote the set of objects available at the beginning of Step t of TTC(P), i.e., $O^t \coloneqq O \setminus \left(\bigcup_{\tau=1}^{t-1} \bigcup_{C' \in \mathcal{C}_\tau(P)} O(C')\right)$. Then $\varphi_{i_k}^{\mathrm{TTC}}(P) \setminus O^t$ is the set of objects assigned to agent i_k before step t of TTC(P); moreover $\varphi_{i_k}(P) \setminus O^t = \varphi_{i_k}^{\mathrm{TTC}}(P) \setminus O^t$.

Claim 2. $\varphi_{i_k}(P) \cap O^t = \omega_{i_k} \cap O^t$.

Proof. Suppose otherwise. By BAL and the fact that $|\varphi_{i_k}(P) \setminus O^t| = |\omega_{i_k} \setminus O^t|$, we must have $|\varphi_{i_k}(P) \cap O^t| = |\omega_{i_k} \cap O^t|$. Consequently, $\varphi_{i_k}(P) \cap O^t \neq \omega_{i_k} \cap O^t$ implies there is an object $o' \in \varphi_{i_k}(P) \cap O^t$ with $o' \notin \omega_{i_k}$.

By the definition of TTC(P), we know that $o_1 = \max_{P_{i_k}}(O^t)$. Because $o' \in \varphi_{i_k}(P) \cap O^t$ and $o_1 \notin \varphi_{i_k}(P)$, it follows that $o_1 \; P_{i_k} \; o'$. Furthermore, WELB implies that $o' \; P_{i_k} \; \min_{P_{i_k}}(\omega_{i_k})$.

Let P'_{i_k} be the truncation of P_{i_k} at o_1. Thus, P'_{i_k} is obtained from P_{i_k} by dropping the set $\{o \in O \setminus \omega_{i_k} \mid o_1 \; P_{i_k} \; o\}$.

Let $P' \coloneqq (P'_{i_k}, P_{-i_k})$. Then $o_1 \; P_{i_k} \; o' \; P_{i_k} \; \min_{P_{i_k}}(\omega_{i_k})$ implies that $s(P') < s(P)$. Consequently, the choice of P implies that $\rho(P') > \rho(P) = t$.

By definition of P', the two runs TTC(P) and TTC(P') execute precisely the same trading cycles at every step $\tau \leq t$, i.e., $\mathcal{C}_\tau(P) = \mathcal{C}_\tau(P')$ for $\tau = 1, \ldots, t$. Thus, $\rho(P') > t$ implies that $\varphi(P')$ executes all trading cycles in $\bigcup_{\tau=1}^{t} \mathcal{C}_\tau(P)$. In particular,

$$(\varphi_{i_k}(P) \setminus O^t) \cup \{o_1\} \subseteq \varphi_{i_k}(P').$$

Since P_{i_k} is lexicographic and $\max_{P_{i_k}}(O^t) = o_1 \notin \varphi_{i_k}(P)$, we must have $\varphi_{i_k}(P') \; P_{i_k} \; \varphi_{i_k}(P)$. This contradicts TP. □

Next, we show that at least two agents are involved in C.

Claim 3. $|N(C)| > 1$.

Proof. Toward contradiction, suppose that $|N(C)| = 1$. Thus $i_1 = i_k$ and $C = (o_1, i_1, o_1)$. Then $o_1 \in \omega_{i_k}$ and Claim 2 implies that $o_1 \in \varphi_{i_k}(P)$, a contradiction. □

By Claim 2 and Claim 3, agent i_{k-1} points to o_k on C, yet she does not receive o_k. An argument similar to the proof of Claim 2 implies that $\varphi_{i_{k-1}}(P) \cap O^t = \omega_{i_{k-1}} \cap O^t$.

Proceeding by induction, we conclude that for each agent i_ℓ involved in C, we have $\varphi_{i_\ell}(P) \cap O^t = \omega_{i_\ell} \cap O^t$. Thus, C is a Pareto-improving trading cycle at P. It follows that φ is not IGE, a contradiction. ■

B.5 Proof of Proposition 1

Fix $P \in \mathcal{R}$ and $i \in N$, and let P'_i be a truncation of P_i at an object $x \in O$.[37] Denote $P' := (P'_i, P_{-i})$ and $U := \{o \in O \mid o \; R_i \; x\}$, and observe that $P_i|_U = P'_i|_U$. Consider two cases.

Case 1. Suppose $\varphi_i^{\mathrm{TTC}}(P) \subseteq U$. Then, by the definition of TTC(P) and TTC(P'), the allocations $\varphi^{\mathrm{TTC}}(P)$ and $\varphi^{TTC}(P')$ execute precisely the same trading cycles. Thus $\varphi_i^{\mathrm{TTC}}(P) = \varphi_i^{\mathrm{TTC}}(P')$.

Case 2. Suppose $\varphi_i^{\mathrm{TTC}}(P) \nsubseteq U$. Then agent i receives at least one object in $O \setminus U$ at $\varphi^{\mathrm{TTC}}(P)$. Let t be the earliest step of TTC(P) at which agent i points to an object in $O \setminus U$. Let O^t be the set of objects remaining at the beginning of step t of TTC(P). Since $P_i|_U = P'_i|_U$, we know that TTC(P) and TTC(P') execute the same trading cycles at steps $1, \ldots, t-1$. Thus, O^t is the set of objects remaining at step t of TTC(P'). Moreover, we have

$$\varphi_i^{\mathrm{TTC}}(P) \setminus O^t = \varphi_i^{\mathrm{TTC}}(P') \setminus O^t. \tag{1}$$

By construction, P'_i ranks every object in $\omega_i \cap O^t$ above every object in $O^t \setminus \omega_i$. Thus, at each step $\tau \geq t$ of TTC(P') where agent i is involved in a trading cycle, she is assigned an object in ω_i. Consequently,

$$\varphi_i^{\mathrm{TTC}}(P') \cap O^t = \omega_i \cap O^t. \tag{2}$$

On the other hand, at each step $\tau \geq t$ of TTC(P) where agent i is involved in a trading cycle, she is assigned an object weakly better than the object she relinquishes from her endowment. In other words, there is a bijection $\sigma : \omega_i \cap O^t \to \varphi_i^{\mathrm{TTC}}(P) \cap O^t$ such that, for each $o \in \omega_i \cap O^t$, $\sigma(o) \; R_i \; o$.

Together with (1) and (2), we conclude that there is a bijection $\pi : \varphi_i^{\mathrm{TTC}}(P') \to \varphi_i^{\mathrm{TTC}}(P)$ such that, for each $o \in \varphi_i^{\mathrm{TTC}}(P')$, $\pi(o) \; R_i \; o$. Since P_i is responsive, we must have $\varphi_i^{\mathrm{TTC}}(P) \; R_i \; \varphi_i^{\mathrm{TTC}}(P')$.[38] ∎

B.6 Proof of Proposition 2

Let $P \in \mathcal{R}$, $i \in N$, and $P'_i \in \mathcal{R}_i$ be such that $\varphi_i^{\mathrm{TTC}}(P'_i, P_{-i}) \; P_i \; \varphi_i^{\mathrm{TTC}}(P)$.

(i) Let $P_i|_O : o_1, \ldots, o_{|\omega_i|}, \ldots$. In particular, $P_i|_O$ ranks every object in $\{o_1, \ldots, o_{|\omega_i|}\}$ above every object in $O \setminus \{o_1, \ldots, o_{|\omega_i|}\}$. For each $j \in N \setminus \{i\}$, let $P^*_j|_O : \ldots, o_1, \ldots, o_{|\omega_i|}$. In particular, $P^*_j|_O$ ranks every object in $O \setminus \{o_1, \ldots, o_{|\omega_i|}\}$ above every object in $\{o_1, \ldots, o_{|\omega_i|}\}$. Then $\varphi_i^{\mathrm{TTC}}(P_i, P^*_{-i}) = \{o_1, \ldots, o_{|\omega_i|}\}$ (otherwise φ^{TTC} would not satisfy IGE). Since P_i is re-

[37] If P'_i is obtained from P_i by dropping $O \setminus \omega_i$, then trivially $\varphi_i^{TTC}(P) \; R_i \; \varphi_i^{TTC}(P'_i, P_{-i}) = \omega_i$.

[38] Starting from $\varphi_i^{\mathrm{TTC}}(P')$, replace each object $o \in \varphi_i^{\mathrm{TTC}}(P')$ with object $\pi(o)$, one at a time, and apply the definition of responsiveness.

sponsive, $\{o_1, \ldots, o_{|\omega_i|}\}$ is the most-preferred $|\omega_i|$-subset of O according to P_i. Consequently, $B_{P_i}(\mathcal{O}_i(P_i))\ R_i\ B_{P_i}(\mathcal{O}_i(P'_i))$.

(ii) For each $j \in N \setminus \{i\}$, let $\omega_j = \{o^j_1, \ldots, o^j_{|\omega_j|}\}$ and $P^*_j|_O : o^j_1, \ldots, o^j_{|\omega_j|}, \ldots$. In particular, $P^*_j|_O$ ranks every object in ω_j above every object in $O \setminus \omega_j$. Then $\varphi^{\text{TTC}}_i(P_i, P^*_{-i}) = \varphi^{\text{TTC}}_i(P'_i, P^*_{-i}) = \omega_i$. Consequently, since φ^{TTC} satisfies IR on $\mathcal{R}$, we have $W_{P_i}(\mathcal{O}_i(P_i)) = \omega_i\ R_i\ W_{P_i}(\mathcal{O}_i(P'_i))$. ■

B.7 Proof of Proposition 4

Fix an agent $i \in N$. To prove that i cannot manipulate using drop strategies, we consider an alternative procedure that proceeds exactly as in the ATTC algorithm except that, at each step in which multiple trading cycles arise, it executes only the cycles that do not involve i. Thus any trading cycle involving agent i is temporarily deferred, and it is executed only when it is the unique remaining cycle. It is easy to see that the order in which cycles are executed does not affect the final allocation; hence this modified procedure yields the same allocation as the original ATTC algorithm.[39]

Formally, given a preference profile $P \in \mathcal{CL}$, the allocation $\varphi^{\text{ATTC}}(P)$ is equivalently obtained via the following algorithm, which we denote $\text{ATTC}^i(P)$. The algorithm defines, for each step t, several ingredients that will be used throughout the proof: (i) the remaining set of objects $O^t(P)$, (ii) the directed graph $G^t(P)$, (iii) the full set $\widetilde{\mathcal{C}}_t(P)$ of cycles that *arise* at step t, (iv) the subset $\overline{\mathcal{C}}_t(P)$ of trading cycles that are *executed* at step t, and (v) the partial allocation $\mu^t(P)$ constructed by the end of step t.

Algorithm: $\text{ATTC}^i(P)$

Initialization: Set $\mu^0(P) \coloneqq (\emptyset)_{j \in N}$ and $O^1(P) \coloneqq O$.

Step $t \geq 1$: Construct a directed graph $G^t(P)$ on $N \cup O^t(P)$ as follows. Each agent j points to $\max_{P_j}(O^t(P) \mid \mu^{t-1}_j(P))$, and each object in $O^t(P)$ points to its owner. Let $\widetilde{\mathcal{C}}_t(P)$ be the resulting set of trading cycles, and let

$$\overline{\mathcal{C}}_t(P) = \begin{cases} \{C \in \widetilde{\mathcal{C}}_t(P) \mid i \notin N(C)\}, & \text{if } |\widetilde{\mathcal{C}}_t(P)| \geq 2 \\ \widetilde{\mathcal{C}}_t(P), & \text{if } |\widetilde{\mathcal{C}}_t(P)| = 1. \end{cases}$$

[39] A similar modification is used by Gonczarowski et al. (2023) to "expose" the strategy-proofness of TTC in single-object environments.

Execute all cycles in $\overline{\mathcal{C}}_t(P)$, yielding the partial allocation $\mu^t(P)$ with

$$\mu_j^t(P) = \begin{cases} \mu_j^{t-1}(P) \cup \{\max_{P_j}(O^t(P) \mid \mu_j^{t-1}(P))\}, & \text{if } j \in \bigcup_{C \in \overline{\mathcal{C}}_t(P)} N(C) \\ \mu_j^{t-1}(P), & \text{otherwise.} \end{cases}$$

Remove all objects involved in executed trading cycles to form the set $O^{t+1}(P) = O^t(P) \setminus \bigcup_{C \in \overline{\mathcal{C}}_t(P)} O(C)$. Proceed to step $t+1$ if $O^{t+1}(P) \neq \emptyset$; proceed to Termination otherwise.

Termination: Let T be the earliest step with $O^{T+1}(P) = \emptyset$. Return $\varphi^{\text{ATTC}}(P) \coloneqq \mu^T(P)$.

Let $P \in \mathcal{CL}$. Let $\tilde{P}_i \in \mathcal{D}_i(P_i)$ be obtained from P_i by dropping an object $x \in O \setminus \omega_i$, and denote $\tilde{P} \coloneqq (\tilde{P}_i, P_{-i})$. We prove that $\varphi_i^{\text{ATTC}}(P) \; R_i \; \varphi_i^{\text{ATTC}}(\tilde{P})$ by comparing the two runs of the modified ATTC algorithm.[40]

If $\overline{\mathcal{C}}_t(P) = \overline{\mathcal{C}}_t(\tilde{P})$ for every step t, then the two runs execute exactly the same cycles; hence $\varphi^{\text{ATTC}}(P) = \varphi^{\text{ATTC}}(\tilde{P})$ and the claim holds.

Otherwise, let t be the earliest step such that $\overline{\mathcal{C}}_t(P) \neq \overline{\mathcal{C}}_t(\tilde{P})$. Then $\overline{\mathcal{C}}_s(P) = \overline{\mathcal{C}}_s(\tilde{P})$ for each step $s < t$, which implies that the partial allocations and the remaining objects coincide at the beginning of step t in the two runs, i.e.,

$$\mu^{t-1}(P) = \mu^{t-1}(\tilde{P}) \quad \text{and} \quad O^t(P) = O^t(\tilde{P}).$$

Consider the directed graphs $G^t(P)$ and $G^t(\tilde{P})$ constructed at step t in the two runs. Because $P_{-i} = \tilde{P}_{-i}$, $\mu^{t-1}(P) = \mu^{t-1}(\tilde{P})$, and $O^t(P) = O^t(\tilde{P})$, every agent $j \in N \setminus \{i\}$ points to the same object in both graphs. Hence, the sets of cycles not involving i coincide:

$$\{C \in \widetilde{\mathcal{C}}_t(P) \mid i \notin N(C)\} = \{C \in \widetilde{\mathcal{C}}_t(\tilde{P}) \mid i \notin N(C)\}. \tag{3}$$

If the common set in (3) were nonempty, then by definition of $\text{ATTC}^i(P)$ and $\text{ATTC}^i(\tilde{P})$ we would have

$$\overline{\mathcal{C}}_t(P) = \{C \in \widetilde{\mathcal{C}}_t(P) \mid i \notin N(C)\} = \{C \in \widetilde{\mathcal{C}}_t(\tilde{P}) \mid i \notin N(C)\} = \overline{\mathcal{C}}_t(\tilde{P}),$$

contradicting the choice of t. Thus

$$\{C \in \widetilde{\mathcal{C}}_t(P) \mid i \notin N(C)\} = \{C \in \widetilde{\mathcal{C}}_t(\tilde{P}) \mid i \notin N(C)\} = \emptyset.$$

[40] A direct step-by-step comparison of $\text{ATTC}(P)$ and $\text{ATTC}(\tilde{P})$ is awkward because agent i may receive the same object at different steps in the two runs, creating a timing mismatch. The modified algorithm is used to synchronize the two runs $\text{ATTC}^i(P)$ and $\text{ATTC}^i(\tilde{P})$ up to the earliest step at which i receives a different object.

Since agent i can be involved in at most one cycle in $G^t(P)$ and in $G^t(\tilde{P})$, and each graph contains a cycle, we must have $\overline{\mathcal{C}}_t(P) = \tilde{\mathcal{C}}_t(P) = \{C\}$ and $\overline{\mathcal{C}}_t(\tilde{P}) = \tilde{\mathcal{C}}_t(\tilde{P}) = \{\tilde{C}\}$ for some trading cycles $C, \tilde{C}$ with $i \in N(C) \cap N(\tilde{C})$. Moreover, $C \neq \tilde{C}$ because $\overline{\mathcal{C}}_t(P) \neq \overline{\mathcal{C}}_t(\tilde{P})$. Thus agent i must point to different objects in $G^t(P)$ and $G^t(\tilde{P})$ (since all other agents point to the same object). Because $Y \coloneqq \mu_i^{t-1}(P) = \mu_i^{t-1}(\tilde{P})$ and $\tilde{P}_i$ is obtained from P_i by dropping object x, the conditional marginal preferences $P_i(Y)|_O$ and $\tilde{P}_i(Y)|_O$ agree on $O \setminus \{x\}$; hence if i points to different objects in the two graphs, then she necessarily points to x in C but not in $\tilde{C}$. Since C is the unique cycle executed at step t of $\text{ATTC}^i(P)$, we have $x \in \varphi_i^{\text{ATTC}}(P)$.

Because $x \notin \omega_i$ and $\tilde{P}_i$ is obtained from P_i by dropping x, the object x is ranked below $\min_{\tilde{P}_i(Z)}(\omega_i)$ according to every conditional marginal preference $\tilde{P}_i(Z)|_O$ associated with $\tilde{P}_i$. Since φ^{ATTC} satisfies WELB, we must have $x \notin \varphi_i^{\text{ATTC}}(\tilde{P})$.

Finally, since agent i points to x at step t of $\text{ATTC}^i(P)$, we have $x = \max_{P_i}\Big(O^t(P) \mid \mu_i^{t-1}(P)\Big)$. Since P_i is conditionally lexicographic and $x \in \varphi_i^{\text{ATTC}}(P) \setminus \varphi_i^{\text{ATTC}}(\tilde{P})$, agent i strictly prefers to complete her partial assignment $\mu_i^{t-1}(P)$ by adding the set $\varphi_i^{\text{ATTC}}(P) \cap O^t(P)$ instead of $\varphi_i^{\text{ATTC}}(\tilde{P}) \cap O^t(\tilde{P})$. Since

$$\begin{aligned} \varphi_i^{\text{ATTC}}(P) &= \mu_i^{t-1}(P) \cup \Big(\varphi_i^{\text{ATTC}}(P) \cap O^t(P)\Big) \\ \text{and} \quad \varphi_i^{\text{ATTC}}(\tilde{P}) &= \mu_i^{t-1}(P) \cup \Big(\varphi_i^{\text{ATTC}}(\tilde{P}) \cap O^t(\tilde{P})\Big), \end{aligned}$$

we conclude that $\varphi_i^{\text{ATTC}}(P)\ P_i\ \varphi_i^{\text{ATTC}}(\tilde{P})$. ■

B.8 Proof of Theorem 6

Toward contradiction, suppose that φ satisfies the properties but $\varphi \neq \varphi^{\text{ATTC}}$. The argument here mirrors the proof of Theorem 1 for the lexicographic domain (see the sketch in Section 3).

For each $P \in \mathcal{CL}$ and each $t \in \mathbb{N}$, recall that $\mathcal{C}_t(P)$ denotes the set of trading cycles that arise at step t of $\text{ATTC}\,(P)$. (We assume that $\mathcal{C}_t(P) = \emptyset$ if $\text{ATTC}(P)$ terminates before step t.) The "size" function $s : \mathcal{CL} \to \mathbb{N}$ is defined, for each $P \in \mathcal{CL}$, by

$$s(P) = \sum_{i \in N} \sum_{Y \in 2^O} \left| \{o \in O \mid o\ R_i(Y) \min_{P_i(Y)}(\omega_i)\} \right|,$$

The "similarity" function $\rho : \mathcal{CL} \to \mathbb{N} \cup \{\infty\}$ is such that, for each $P \in \mathcal{CL}$,

$$\rho(P) = \inf\{t \in \mathbb{N} \mid \varphi(P) \text{ does not execute each trading cycle in } \mathcal{C}_t(P)\},$$

where it is understood that $\rho(P) = \infty$ precisely when $\varphi(P) = \varphi^{\text{ATTC}}(P)$.

Among all profiles that minimize the similarity function ρ, let P be one that further minimizes the size function s. Thus, for any profile $P' \in \mathcal{CL}$, either (i) $\rho(P) < \rho(P')$, or (ii) $\rho(P) = \rho(P')$

and $s(P) \leq s(P')$.

Let $t \coloneqq \rho(P)$. Then $\varphi(P)$ executes each trading cycle in $\bigcup_{\tau=1}^{t-1} \mathcal{C}_\tau(P)$, but it does not execute some trading cycle $C \in \mathcal{C}_t(P)$. Suppose that

$$C = (o_1, i_1, o_2, i_2, \dots, o_k, i_k, o_{k+1} = o_1)\,.$$

Because $\varphi(P)$ does not execute C, there is an agent $i_\ell \in N(C)$ who is not assigned the object she points to within C, i.e., $o_{\ell+1} \notin \varphi_{i_\ell}(P)$. Without loss of generality, let $i_\ell = i_k$. Thus, $o_{k+1} = o_1 \notin \varphi_{i_k}(P)$.

Let O^t denote the set of objects available at the beginning of Step t of ATTC(P), i.e., $O^t \coloneqq O \setminus \left(\bigcup_{\tau=1}^{t-1} \bigcup_{C' \in \mathcal{C}_\tau(P)} O(C')\right)$. Then $\varphi_{i_k}^{\text{ATTC}}(P) \setminus O^t$ is the set of objects assigned to agent i_k before step t of ATTC(P); moreover $\varphi_{i_k}(P) \setminus O^t = \varphi_{i_k}^{\text{ATTC}}(P) \setminus O^t$.

Claim 4. $\varphi_{i_k}(P) \cap O^t = \omega_{i_k} \cap O^t$.

Proof. Suppose otherwise. By BAL and the fact that $|\varphi_{i_k}(P) \setminus O^t| = |\omega_{i_k} \setminus O^t|$, we must have $|\varphi_{i_k}(P) \cap O^t| = |\omega_{i_k} \cap O^t|$. Consequently, $\varphi_{i_k}(P) \cap O^t \neq \omega_{i_k} \cap O^t$ implies there is an object $o' \in \varphi_{i_k}(P) \cap O^t$ with $o' \notin \omega_{i_k}$.

By the definition of ATTC(P), we know that $o_1 = \max_{P_{i_k}}(O^t \mid \varphi_{i_k}(P) \setminus O^t)$. Because $o' \in \varphi_{i_k}(P) \cap O^t$ and $o_1 \notin \varphi_{i_k}(P)$, it follows that $o_1 \; P_{i_k}(\varphi_{i_k}(P) \setminus O^t) \; o'$ (thus $o_1 \neq o'$). Furthermore, WELB implies that $o' \; P_{i_k}(\varphi_{i_k}(P)) \; \min_{P_{i_k}(\varphi_{i_k}(P))}(\omega_{i_k})$.

Let P'_{i_k} be the drop strategy obtained from P_{i_k} by dropping object o'. Let $P' \coloneqq (P'_{i_k}, P_{-i_k})$. Then, for each $Y \subseteq O$,

$$\left|\{o \in O \mid o \; R'_i(Y) \min_{P'_i(Y)}(\omega_i)\}\right| \leq \left|\{o \in O \mid o \; R_i(Y) \min_{P_i(Y)}(\omega_i)\}\right|,$$

with strict inequality for $Y = \varphi_{i_k}(P)$ because $o' \; P_{i_k}(\varphi_{i_k}(P)) \; \min_{P_{i_k}(\varphi_{i_k}(P))}(\omega_{i_k})$. Thus $s(P') < s(P)$. Consequently, the choice of P implies that $\rho(P') > \rho(P) = t$.

By definition of P'_{i_k}, the two conditional marginal preferences $P'_{i_k}(\varphi_{i_k}(P) \setminus O^t)$ and $P_{i_k}(\varphi_{i_k}(P) \setminus O^t)$ agree for all objects at least as good as o_1. In particular, $\max_{P'_{i_k}}(X \mid Y) = \max_{P_{i_k}}(X \mid Y)$ whenever $Y \subseteq \varphi_{i_k}(P) \setminus O^t$ and $X \supseteq O^t$. Consequently, ATTC(P) and ATTC(P') execute precisely the same trading cycles at every step $\tau \leq t$, i.e., $\mathcal{C}_\tau(P) = \mathcal{C}_\tau(P')$ for $\tau = 1, \dots, t$. Thus, $\rho(P') > t$ implies that $\varphi(P')$ executes all trading cycles in $\bigcup_{\tau=1}^{t} \mathcal{C}_\tau(P)$. In particular,

$$(\varphi_{i_k}(P) \setminus O^t) \cup \{o_1\} \subseteq \varphi_{i_k}(P').$$

Since P_{i_k} is conditionally lexicographic and $\max_{P_{i_k}}(O^t \mid \varphi_{i_k}(P) \setminus O^t) = o_1 \notin \varphi_{i_k}(P)$, Definition 9 implies that $\varphi_{i_k}(P') \; P_{i_k} \; \varphi_{i_k}(P)$. This contradicts DSP. $\square$

Next, we show that at least two agents are involved in C.

Claim 5. $|N(C)| > 1$.

Proof. Toward contradiction, suppose that $|N(C)| = 1$. Thus $i_1 = i_k$ and $C = (o_1, i_1, o_1)$. Then $o_1 \in \omega_{i_k}$ and Claim 4 implies that $o_1 \in \varphi_{i_k}(P)$, a contradiction. □

By Claim 4 and Claim 5, agent i_{k-1} points to o_k on C, yet she does not receive o_k. An argument similar to the proof of Claim 4 implies that $\varphi_{i_{k-1}}(P) \cap O^t = \omega_{i_{k-1}} \cap O^t$.

Proceeding by induction, we conclude that for each agent i_ℓ involved in C, we have $\varphi_{i_\ell}(P) \cap O^t = \omega_{i_\ell} \cap O^t$. Thus, C is a Pareto-improving trading cycle at P. It follows that φ is not IGE, a contradiction. ■

B.9 Proof of Proposition 5

Without loss of generality, $P_1 \in \mathcal{M}_1 \setminus \mathcal{CL}_1$. Since P_1 is not conditionally lexicographic, there exist disjoint subsets $X, Y \in 2^O$ with $X \neq \emptyset$ such that

$$\text{for all } x \in X, \text{ there exists } Z_x \subseteq X \setminus x \text{ such that } (Y \cup Z_x)\ P_1\ (Y \cup x).$$

Because P_1 is *monotone*, $Z_x \subseteq X \setminus x$ implies $[Y \cup (X \setminus x)]\ R_1\ (Y \cup Z_x)$. Thus, we have

$$\text{for all } x \in X, \quad [Y \cup (X \setminus x)]\ P_1\ (Y \cup x).$$

Note also that $|X| \geq 3$.[41]

Let $x^* \in X$ be such that $(Y \cup x^*)\ R_1\ (Y \cup x)$ for all $x \in X$. Then, since $|X| \geq 3$ and P_1 is monotonic, it follows that

$$\text{for all } x \in X, \quad [Y \cup (X \setminus x)]\ P_1\ (Y \cup x^*)\ R_1\ (Y \cup x).$$

Case 1. Suppose $X \cup Y = O$. Let $N = \{1, 2\}$, and let $P_2 \in \mathcal{L}$ be such that (i) $\max_{P_2}(O) = x^*$, and (ii) for all $x \in X$, $x\ P_2\ Y$. Then, in particular,

$$x^*\ P_2\ [Y \cup (X \setminus x^*)]\ R_2\ (X \setminus x^*).$$

Consider the allocations

$$\mu := (Y \cup x^*, X \setminus x^*) \quad \text{and} \quad \overline{\mu} := (Y \cup (X \setminus x^*), x^*).$$

[41] If X is a singleton, say $X = \{x\}$, then $Y\ P_1\ (Y \cup x)$, a violation of monotonicity. If X contains two objects, say $X = \{x, y\}$, we would have $(Y \cup x)\ P_1\ (Y \cup y)\ P_1\ (Y \cup x)$, a violation of transitivity.

Then μ is not PE at P because it is Pareto dominated by $\overline{\mu}$ at P.

We claim that μ satisfies IGE at P. Indeed, consider any trading cycle $C = (x, 2, y, 1, x)$, where $x \in \mu_2$ and $y \in \mu_1$. Executing C yields the allocation μ', where

$$\mu'_1 = (Y \cup \{x^*, x\}) \setminus y \quad \text{and} \quad \mu'_2 = (X \setminus \{x^*, x\}) \cup y.$$

If $y \in Y$, then this exchange harms agent 2, i.e., $\mu_2 \; P_2 \; \mu'_2$. On the other hand, if $y = x^*$, then this exchange harms agent 1 because $\mu_1 = (Y \cup x^*) \; P_1 \; (Y \cup x) = \mu'_1$. Hence, μ satisfies IGE.

Case 2. Suppose $X \cup Y \subsetneq O$. Let $N = \{1, 2, 3\}$ and denote $\overline{O} \coloneqq O \setminus (X \cup Y)$. Let $P_2 \in \mathcal{L}$ be such that (i) $\max_{P_2}(O) = x^*$, and (ii) for all $x \in X$, $x \; P_2 \; (O \setminus X)$. Let $P_3 \in \mathcal{L}$ be such that (i) for all $z \in \overline{O}$, $z \; P_3 \; (X \cup Y)$.

Consider the allocations

$$\mu \coloneqq (Y \cup x^*, X \setminus x^*, \overline{O}) \quad \text{and} \quad \overline{\mu} \coloneqq (Y \cup (X \setminus x^*), x^*, \overline{O}).$$

Then μ is not PE at P because it is Pareto dominated by $\overline{\mu}$ at P.

We claim that μ satisfies IGE at P. Because μ_3 is agent 3's top-ranked $|\overline{O}|$-subset of O, any trading cycle at μ involving agent 3 must harm agent 3. Thus, any Pareto-improving trading cycle at μ must involve only agents 1 and 2. However, there is no such exchange by the argument in Case 1. Hence, μ satisfies IGE. ■

C Examples

The examples below illustrate the logical independence of the properties used in Theorems 1–4 (see Tables 1 and 2 for a summary of which properties each example satisfies on $\mathcal{L}$ and $\mathcal{R}$). Since $\mathcal{L} \subseteq \mathcal{R}$, any rule defined on $\mathcal{R}$ is also viewed as a rule on $\mathcal{L}$ by restriction.

Concretely, for each property appearing in one of these theorems, we exhibit a rule (defined on the relevant domain) that violates this property while satisfying all other properties in the theorem. To keep notation light, some examples fix convenient values for the set N of agents and the endowment ω. Analogous constructions can be given on $\mathcal{CL}$; we omit them for brevity.

Example 5. *No-trade rule* (fails IGE; satisfies MAR, BAL, WELB, IR, and SP)

Let φ^{NT} be the *no-trade rule* on $\mathcal{R}$, which selects ω for each problem. Then φ^{NT} satisfies MAR, BAL, WELB, IR, and SP on $\mathcal{R}$, but it fails IGE on $\mathcal{L}$ (and thus on $\mathcal{R}$). ◇

Example 6. *Balanced Serial Dictatorship* (fails WELB and IR; satisfies MAR, BAL, IGE, and SP)

Let φ^{BSD} be the *Balanced Serial Dictatorship* on $\mathcal{R}$, defined recursively as follows: for each

$P \in \mathcal{R}$, $\varphi_1^{\text{BSD}}(P) = \max_{P_1}(O, |\omega_1|)$ and, for each $i = 2, \ldots, n$,

$$\varphi_i^{\text{BSD}}(P) = \max_{P_i}\left(O \setminus \bigcup_{j=1}^{i-1} \varphi_j^{\text{BSD}}(P), |\omega_i|\right),$$

where $\max_{P_i}(X, k)$ denotes the most-preferred k-subset of X at P_i. Then φ^{BSD} satisfies MAR, BAL, IGE, and SP on $\mathcal{R}$, but it fails WELB and IR on $\mathcal{L}$ (and thus on $\mathcal{R}$). $\diamond$

Example 7. *A rule* $\varphi^{\neg TP}$ (fails TP and DSP; satisfies MAR, BAL, IGE, WELB, and IR)

Let $N = \{1, 2\}$, $O = \{a, b, c\}$, and $\omega = (\{a, b\}, \{c\})$. Fix $P^* \in \mathcal{L}$ with $P_1^*|_O : c, a, b$ and $P_2^*|_O : a, b, c$, and let $\mu^* = (\{a, c\}, \{b\})$. Define the rule $\varphi^{\neg\text{TP}}$ on $\mathcal{R}$ by

$$\varphi^{\neg\text{TP}}(P) = \begin{cases} \mu^*, & \text{if } P|_O = P^*|_O \\ \varphi^{\text{TTC}}(P), & \text{otherwise.} \end{cases}$$

It is easy to see that $\varphi^{\neg\text{TP}}$ satisfies MAR, BAL, IGE, WELB, and IR on $\mathcal{R}$.

We show that $\varphi^{\neg\text{TP}}$ fails TP and DSP on $\mathcal{L}$. Consider $P_2' \in \mathcal{L}_2$ with $P_2'|_O : a, c, b$. Then $\varphi_2^{\neg\text{TP}}(P_1^*, P_2') = \{a\}\ P_2^*\ \{b\} = \varphi_2^{\neg\text{TP}}(P^*)$, so P_2' is a profitable manipulation for agent 2 at P^*. Since P_2' is both a truncation strategy and a drop strategy for P_2^* (it is obtained by dropping the singleton tail subset $\{o \in O \setminus \omega_2 \mid b\ R_2^*\ o\} = \{b\}$), $\varphi^{\neg\text{TP}}$ fails TP and DSP on $\mathcal{L}$ (and thus on $\mathcal{R}$). $\diamond$

Example 8. *A rule* $\varphi^{\neg BAL}$ *on* $\mathcal{L}$ (fails BAL; satisfies PE, WELB, IR, and SDSP)

Let $N = \{1, 2\}$, $O = \{a, b, c\}$, and $\omega = (\{a, b\}, \{c\})$. Consider the lexicographic profile P^* with $P_1^* : c, a, b$ and $P_2^* : a, b, c$, and let $\mu^* \coloneqq (\{c\}, \{a, b\})$. Define the rule $\varphi^{\neg\text{BAL}}$ on $\mathcal{L}$ by

$$\varphi^{\neg\text{BAL}}(P) = \begin{cases} \mu^*, & \text{if } P = P^* \\ \varphi^{\text{TTC}}(P), & \text{otherwise.} \end{cases}$$

It is easy to see that $\varphi^{\neg\text{BAL}}$ satisfies PE, WELB, and IR.

We show that $\varphi^{\neg\text{BAL}}$ satisfies DSP (and thus SDSP by Lemma 2). Fix $P \in \mathcal{L}$, agent $i \in N$, and a drop strategy $P_i' \in \mathcal{D}_i(P_i) \setminus \{P_i\}$. Denote $P' \coloneqq (P_i', P_{-i})$ and consider three cases.

Case 1. Suppose $P \neq P^*$ and $P' \neq P^*$. Then $\varphi^{\neg\text{BAL}}(P) = \varphi^{\text{TTC}}(P)$ and $\varphi^{\neg\text{BAL}}(P') = \varphi^{\text{TTC}}(P')$, so the deviation is not profitable because φ^{TTC} satisfies SDSP on $\mathcal{L}$.

Case 2. Suppose $P = P^*$. If $i = 2$, then at P_2 the assignment $\varphi^{\neg\text{BAL}}(P) = \mu_2^*$ is agent 2's most-preferred feasible bundle (i.e., among all nonempty, proper subsets of O), so no deviation can be profitable. If $i = 1$, then any proper drop strategy for P_1 must drop object $c \in O \setminus \omega_1$, giving $P_1' : a, b, c$. Thus $\varphi_1^{\neg\text{BAL}}(P) = \{c\}\ P_1\ \{a, b\} = \varphi_1^{\neg\text{BAL}}(P')$.

Case 3. The case $P' = P^*$ is impossible. Indeed, under P_1^* the bottom object is $b \in \omega_1$, so P_1^* is not a proper drop strategy for any $P_1 \in \mathcal{L}$. Similarly, under P_2^* the bottom object is $c \in \omega_2$, so P_2^* is not a proper drop strategy for any $P_2 \in \mathcal{L}$. $\diamond$

Example 9. *A rule $\varphi^{\neg WELB}$ on $\mathcal{L}$* (fails WELB; satisfies BAL, IGE, IR, and SDSP)

Let $N = \{1, 2\}$, $O = \{a, b, c, d\}$, and $\omega = (\{a, b\}, \{c, d\})$. Consider the lexicographic profile P^* with $P_1^* : c, a, b, d$ and $P_2^* : a, b, c, d$, and let $\mu^* = (\{c, d\}, \{a, b\})$. Define $\varphi^{\neg \text{WELB}}$ on $\mathcal{L}$ by

$$\varphi^{\neg \text{WELB}}(P) = \begin{cases} \mu^*, & \text{if } P = P^* \\ \varphi^{\text{TTC}}(P), & \text{otherwise.} \end{cases}$$

It is easy to see that $\varphi^{\neg \text{WELB}}$ satisfies BAL, IGE, and IR. Furthermore, $\varphi^{\neg \text{WELB}}$ fails WELB at P^* because $d \in \varphi_1^{\neg \text{WELB}}(P^*)$ although $\min_{P_1^*}(\omega_1) = b \; P_1^* \; d$.

We now show that $\varphi^{\neg \text{WELB}}$ satisfies SDSP (hence TP and DSP). Fix $P \in \mathcal{L}$, agent $i \in N$, and a subset-drop strategy $P_i' \in \mathcal{S}_i(P_i) \setminus \{P_i\}$. Denote $P' \coloneqq (P_i', P_{-i})$ and consider three cases.

Case 1. Suppose $P = P^*$. If $i = 2$, then $\varphi_2^{\neg \text{WELB}}(P) = \mu_2^*$ is agent 2's most-preferred two-object bundle at P_2^*, so no deviation can be profitable. If $i = 1$, then any proper subset-drop strategy for P_1 must drop object $c \in O \setminus \omega_1$ (possibly together with d), which means that P_1' is either $P_1^1 : a, b, d, c$ or $P_1^2 : a, b, c, d$. In each case, φ^{TTC} assigns $\varphi^{\text{TTC}}(P') = \omega$. Thus $\varphi_1^{\neg \text{WELB}}(P) = \{c, d\} \; P_1 \; \{a, b\} = \varphi_1^{\neg \text{WELB}}(P')$.

Case 2. Suppose $P' = P^*$. Then necessarily $i = 1$ and $P_2 = P_2^*$ (since P_2^* ranks an object $d \in \omega_2$ at the very bottom, it is not a proper subset-drop strategy for any $P_2 \in \mathcal{L}_2$). Moreover, $P_1' = P_1^*$ can be obtained from P_1 only by dropping object d, which means that P_1 is either $P_1^1 : d, c, a, b$, $P_1^2 : c, d, a, b$, or $P_1^3 : c, a, d, b$. In each case, φ^{TTC} assigns $\varphi^{\text{TTC}}(P) = \mu^* = \varphi^{\neg \text{WELB}}(P')$, so the deviation is not profitable.

Case 3. Suppose $P \neq P^*$ and $P' \neq P^*$. Then $\varphi^{\neg \text{WELB}}(P) = \varphi^{\text{TTC}}(P)$ and $\varphi^{\neg \text{WELB}}(P') = \varphi^{\text{TTC}}(P')$, so the deviation is not profitable because φ^{TTC} satisfies SDSP on $\mathcal{L}$. $\diamond$

Example 10. *A rule $\varphi^{\neg MAR}$ on $\mathcal{R}$* (fails MAR; satisfies IGE, WELB, IR, and TP)

Let $N = \{1, 2\}$ and $\omega = (\{a, b, d\}, \{c, e\})$. Fix an allocation $\mu^* \coloneqq (\{b, c, e\}, \{a, d\})$. Let

$$\mathcal{R}^* = \{P^* \in \mathcal{R} \mid P_1^*|_O : a, e, b, c, d \text{ and } P_2^*|_O : a, c, b, d, e\}.$$

Since φ^{TTC} is marginal, for each $P \in \mathcal{R}^*$ we have $\varphi^{\text{TTC}}(P) = (\{a, d, e\}, \{b, c\})$.

Define the rule $\varphi^{\neg \text{MAR}}$ on $\mathcal{R}$ by

$$\varphi^{\neg \text{MAR}}(P) = \begin{cases} \mu^*, & \text{if } P \in \mathcal{R}^* \text{ and } \mu^* \text{ Pareto dominates } \varphi^{\text{TTC}}(P) = (\{a, d, e\}, \{b, c\}), \\ \varphi^{\text{TTC}}(P), & \text{otherwise.} \end{cases}$$

One easily checks that $\varphi^{\neg\text{MAR}}$ satisfies IGE, WELB, and IR (and BAL).

We show that $\varphi^{\neg\text{MAR}}$ satisfies TP. Fix $P \in \mathcal{R}$ and $i \in N$. Let $P'_i \in \mathcal{T}_i(P_i)$ and denote $P' \coloneqq (P'_i, P_{-i})$. Denote the unique agent in $N \setminus \{i\}$ by j. We consider two cases.

Case 1. Suppose $P \in \mathcal{R}^*$. Consider two further subcases.

Subcase i. Suppose $\varphi_j^{\text{TTC}}(P)\ P_j\ \mu_j^*$. Then μ^* does not Pareto dominate $\varphi^{\text{TTC}}(P)$ at P, so the definition of $\varphi^{\neg\text{MAR}}$ implies that $\varphi_i^{\neg\text{MAR}}(P) = \varphi^{\text{TTC}}(P)$ and $\varphi_i^{\neg\text{MAR}}(P') = \varphi^{\text{TTC}}(P')$. Thus TP for φ^{TTC} implies that P'_i is not a profitable manipulation.

Subcase ii. Suppose $\mu_j^*\ R_j\ \varphi_j^{\text{TTC}}(P)$. Then the definition of $\varphi^{\neg\text{MAR}}$ implies that $\varphi_i^{\neg\text{MAR}}(P)$ is the most-preferred bundle in $\{\mu_i^*, \varphi_i^{\text{TTC}}(P)\}$ according to P_i. If $P' \in \mathcal{R}^*$, then $\varphi_i(P') \in \{\mu_i^*, \varphi_i^{\text{TTC}}(P)\}$, which implies $\varphi_i^{\neg\text{MAR}}(P)\ R_i\ \varphi_i^{\neg\text{MAR}}(P')$. If $P' \notin \mathcal{R}^*$, then TP for φ^{TTC} implies that

$$\varphi_i^{\neg\text{MAR}}(P)\ R_i\ \varphi_i^{\text{TTC}}(P)\ R_i\ \varphi_i^{\text{TTC}}(P') = \varphi_i^{\neg\text{MAR}}(P').$$

Case 2. Suppose $P \in \mathcal{R} \setminus \mathcal{R}^*$. If $P' \in \mathcal{R} \setminus \mathcal{R}^*$, then TP for φ^{TTC} implies that

$$\varphi_i^{\neg\text{MAR}}(P) = \varphi_i^{\text{TTC}}(P)\ R_i\ \varphi_i^{\text{TTC}}(P') = \varphi_i^{\neg\text{MAR}}(P').$$

Furthermore, it is not possible that $P' \in \mathcal{R}^*$. Indeed, $P' \in \mathcal{R}^*$ would imply $P'_i|_O \neq P_i|_O$, in which case $P'_i|_O$ is obtained from $P_i|_O$ by dropping a *nonempty* tail subset $X \subseteq O \setminus \omega_i$. In particular, this means $\min_{P'_i}(O) \notin \omega_i$, a contradiction. $\diamond$

References

Afacan, M. O. (2016): "Characterizations of the cumulative offer process," *Social Choice and Welfare*, 47, 531–542.

Altuntaş, A., W. Phan, and Y. Tamura (2023): "Some characterizations of generalized top trading cycles," *Games and Economic Behavior*, 141, 156–181.

Andersson, T., Á. Cseh, L. Ehlers, and A. Erlanson (2021): "Organizing time exchanges: Lessons from matching markets," *American Economic Journal: Microeconomics*, 13, 338–373.

Aziz, H., P. Biró, J. Lang, J. Lesca, and J. Monnot (2019): "Efficient reallocation under additive and responsive preferences," *Theoretical Computer Science*, 790, 1–15.

Bichler, M., A. Hammerl, T. Morrill, and S. Waldherr (2021): "How to Assign Scarce Resources Without Money: Designing Information Systems that are Efficient, Truthful, and (Pretty) Fair," *Information Systems Research*, 32, 335–355.

Bird, C. G. (1984): "Group incentive compatibility in a market with indivisible goods," *Economics Letters*, 14, 309–313.

Biró, P., F. Klijn, and S. Pápai (2022a): "Balanced Exchange in a Multi-Unit Shapley-Scarf Market," Barcelona School of Economics Working Paper 1342.

——— (2022b): "Serial rules in a multi-unit Shapley-Scarf market," *Games and Economic Behavior*, 136, 428–453.

Bloch, F., D. Cantala, and D. Gibaja (2020): "Matching through institutions," *Games and Economic Behavior*, 121, 204–231.

Booth, R., Y. Chevaleyre, J. Lang, J. Mengin, and C. Sombattheera (2010): "Learning conditionally lexicographic preference relations," in *ECAI 2010*: IOS Press, 269–274.

Budish, E. (2011): "The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes," *Journal of Political Economy*, 119, 1061–1103.

Budish, E., and E. Cantillon (2012): "The multi-unit assignment problem: Theory and evidence from course allocation at Harvard," *American Economic Review*, 102, 2237–2271.

Caspari, G. (2020): "Booster draft mechanism for multi-object assignment," ZEW-Centre for European Economic Research Discussion Paper 20-074.

Cechlárová, K., P. Eirinakis, T. Fleiner, D. Magos, I. Mourtos, and E. Potpinková (2014): "Pareto optimality in many-to-many matching problems," *Discrete Optimization*, 14, 160–169.

Chen, Y., Z. Jiao, C. Zhang, and L. Zhang (2025): "Some new characterizations of top trading cycles," *Journal of Mathematical Economics*, 121, 103190.

Coreno, J., and I. Balbuzanov (2022): "Axiomatic Characterizations of Draft Rules," arXiv preprint arXiv:2204.08300.

Coreno, J., and D. Feng (2025): "Characterizing TTC via endowments-swapping-proofness and truncation-proofness," *Economics Letters*, 247, 112159.

De Keijzer, B., S. Bouveret, T. Klos, and Y. Zhang (2009): "On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences," in *International Conference on Algorithmic Decision Theory*, 98–110, Springer.

Domshlak, C., E. Hüllermeier, S. Kaci, and H. Prade (2011): "Preferences in AI: An overview," *Artificial Intelligence*, 175, 1037–1052.

Dur, U. M., and M. U. Ünver (2019): "Two-sided matching via balanced exchange," *Journal of Political Economy*, 127, 1156–1177.

Echenique, F., A. Miralles, and J. Zhang (2021): "Fairness and efficiency for allocations with participation constraints," *Journal of Economic Theory*, 195, 105274.

——— (2023): "Balanced equilibrium in pseudo-markets with endowments," *Games and Economic Behavior*, 141, 428–443.

Ehlers, L. (2008): "Truncation strategies in matching markets," *Mathematics of Operations Research*, 33, 327–335.

Ehlers, L., and B. Klaus (2003): "Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems," *Social Choice and Welfare*, 21, 265–280.

Ekici, Ö. (2024): "Pair-efficient reallocation of indivisible objects," *Theoretical Economics*, 19, 551–564.

Ergin, H., T. Sönmez, and M. U. Ünver (2017): "Dual-donor organ exchange," *Econometrica*, 85, 1645–1671.

——— (2020): "Efficient and incentive-compatible liver exchange," *Econometrica*, 88, 965–1005.

Feng, D., B. Klaus, and F. Klijn (2024): "Characterizing the Typewise Top-Trading-Cycles Mechanism for Multiple-Type Housing Markets," *Games and Economic Behavior*, Forthcoming.

Fujita, E., J. Lesca, A. Sonoda, T. Todo, and M. Yokoo (2018): "A complexity approach for core-selecting exchange under conditionally lexicographic preferences," *Journal of Artificial Intelligence Research*, 63, 515–555.

Gigerenzer, G., and D. G. Goldstein (1996): "Reasoning the fast and frugal way: models of bounded rationality.," *Psychological review*, 103, 650.

Gigerenzer, G., and P. M. Todd (1999): *Simple heuristics that make us smart*: Oxford University Press, USA.

Gonczarowski, Y. A., O. Heffetz, and C. Thomas (2023): "Strategyproofness-exposing mechanism descriptions," NBER Working paper No. 31506.

Han, X., O. Kesten, and M. U. Ünver (2024): "Blood Allocation with Replacement Donors: Theory and Application," *Available at SSRN 3858158*.

Hatfield, J. W. (2009): "Strategy-proof, efficient, and nonbossy quota allocations," *Social*

Choice and Welfare, 33, 505–515.

Hatfield, J. W., S. D. Kominers, and A. Westkamp (2021): "Stability, Strategy-Proofness, and Cumulative Offer Mechanisms," *The Review of Economic Studies*, 88, 1457–1502.

Hylland, A., and R. Zeckhauser (1979): "The Efficient Allocation of Individuals to Positions," *Journal of Political Economy*, 87, 293–314.

Jantschgi, S., A. Teytelboym, and T. Nguyen (2025): "Competitive Combinatorial Exchange," *Available at SSRN 5283955*.

Klaus, B., and E. Miyagawa (2002): "Strategy-proofness, solidarity, and consistency for multiple assignment problems," *International Journal of Game Theory*, 30, 421–435.

Kojima, F. (2013): "Efficient resource allocation under multi-unit demand," *Games and economic behavior*, 82, 1–14.

Konishi, H., T. Quint, and J. Wako (2001): "On the Shapley–Scarf economy: the case of multiple types of indivisible goods," *Journal of Mathematical Economics*, 35, 1–15.

Ma, J. (1994): "Strategy-proofness and the strict core in a market with indivisibilities," *International Journal of Game Theory*, 23, 75–83.

Manjunath, V., and A. Westkamp (2021): "Strategy-proof exchange under trichotomous preferences," *Journal of Economic Theory*, 193, 105197.

——— (2025): "Marginal mechanisms for balanced exchange," arXiv preprint arXiv:2502.06499.

Milgrom, P. (2009): "Assignment messages and exchanges," *American Economic Journal: Microeconomics*, 1, 95–113.

——— (2011): "Critical issues in the practice of market design," *Economic Inquiry*, 49, 311–320.

Mongell, S., and A. E. Roth (1991): "Sorority rush as a two-sided matching mechanism," *The American Economic Review*, 441–464.

Monte, D., and N. Tumennasan (2015): "Centralized allocation in multiple markets," *Journal of Mathematical Economics*, 61, 74–85.

Morrill, T., and A. E. Roth (2024): "Top trading cycles," *Journal of Mathematical Economics*, 112, 102984.

Moulin, H. (1995): *Cooperative Microeconomics: A Game-Theoretic Introduction*: Princeton, NJ: Princeton University Press.

Nguyen, T., and R. Vohra (2024): "(Near-)Substitute Preferences and Equilibria with Indivisibilities," *Journal of Political Economy*, 132, 4122–4154.

Pápai, S. (2001): "Strategyproof and nonbossy multiple assignments," *Journal of Public Economic Theory*, 3, 257–271.

——— (2003): "Strategyproof exchange of indivisible goods," *Journal of Mathematical Economics*, 39, 931–959.

——— (2007): "Exchange in a general market with indivisible goods," *Journal of Economic*

Theory, 132, 208–235.

Phan, W., and C. Purcell (2022): "The parameterized complexity of manipulating Top Trading Cycles," *Autonomous Agents and Multi-Agent Systems*, 36, 51.

Pigozzi, G., A. Tsoukias, and P. Viappiani (2016): "Preferences in artificial intelligence," *Annals of Mathematics and Artificial Intelligence*, 77, 361–401.

Roth, A. E. (1982): "Incentive compatibility in a market with indivisible goods," *Economics Letters*, 9, 127–132.

——— (1985): "The college admissions problem is not equivalent to the marriage problem," *Journal of economic Theory*, 36, 277–288.

——— (2002): "The economist as engineer: Game theory, experimentation, and computation as tools for design economics," *Econometrica*, 70, 1341–1378.

——— (2015): "Experiments in Market Design," in *The Handbook of Experimental Economics, Volume 2* ed. by Roth, A. E., and Kagel, J. H.: Princeton University Press, 290–346.

Roth, A. E., and E. Peranson (1999): "The redesign of the matching market for American physicians: Some engineering aspects of economic design," *American economic review*, 89, 748–780.

Roth, A. E., and A. Postlewaite (1977): "Weak versus strong domination in a market with indivisible goods," *Journal of Mathematical Economics*, 4, 131–137.

Roth, A. E., and U. G. Rothblum (1999): "Truncation strategies in matching markets—in search of advice for participants," *Econometrica*, 67, 21–43.

Roth, A. E., T. Sönmez, and M. U. Ünver (2004): "Kidney exchange," *The Quarterly Journal of Economics*, 119, 457–488.

——— (2005): "Pairwise kidney exchange," *Journal of Economic Theory*, 125, 151–188.

Roth, A. E., and J. H. Vande Vate (1991): "Incentives in two-sided matching with random stable mechanisms," *Economic Theory*, 1, 31–44.

Sethuraman, J. (2016): "An alternative proof of a characterization of the TTC mechanism," *Operations Research Letters*, 44, 107–108.

Shapley, L. S., and H. Scarf (1974): "On cores and indivisibility," *Journal of Mathematical Economics*, 1, 23–37.

Shirakawa, R. (2025): "Simple manipulations in school choice mechanisms," *American Economic Journal: Microeconomics*, 17, 202–233.

Sönmez, T. (1999): "Strategy-proofness and essentially single-valued cores," *Econometrica*, 67, 677–689.

Todo, T., H. Sun, and M. Yokoo (2014): "Strategyproof exchange with multiple private endowments," in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 805–811.

Troyan, P., and T. Morrill (2020): "Obvious manipulations," *Journal of Economic Theory*, 185, 104970.