Barcelona School of Economics

WORKING PAPERS

Balanced Exchange in a Multi-Unit Shapley-Scarf Market

BSE Working Paper 1342| April 2022

Péter Biró, Flip Klijn, Szilvia Pápai

Balanced Exchange in a Multi-Unit Shapley-Scarf Market*

Péter Biró[†] Flip Klijn[‡] Szilvia Pápai[§]

April 19, 2022

Abstract

We study markets in which each agent is endowed with multiple units of an indivisible and agent-specific good. Monetary compensations are not possible. An outcome of a market is given by a circulation which consists of a balanced exchange of goods. Agents only have (responsive) preferences over the bundles they receive.

We prove that for general capacity configurations there is no circulation rule that satisfies individual rationality, Pareto-efficiency, and strategy-proofness. We characterize the (so-called irreducible) capacity configurations for which the three properties are compatible, and show that in this case the Circulation Top Trading Cycle (cTTC) rule is the unique rule that satisfies all three properties. We also explore the incentive and efficiency properties of the cTTC rule for general capacity configurations and provide a characterization of the rule for lexicographic preferences.

Next, we introduce and study the family of so-called Segmented Trading Cycle (STC) rules. These rules are obtained by first distributing agents' endowments over a number of different smaller markets (the market segments), then applying the standard Top Trading Cycle algorithm within each market segment separately, and finally lumping together the induced circulations. We show that STC rules are individually rational, strategy-proof, and nonbossy. Even though STC rules do not satisfy group-strategy-proofness in general, they do satisfy weaker notions of group-strategy-proofness. For irreducible capacity configurations

*We thank seminar and workshop participants at Boston College, Keio University (Third International Workshop on Market Design Technologies for Sustainable Development), Université Paris Dauphine, Lausanne Economic Theory Workshop 2017, 1st Catalan Economic Society Conference, UAB Microlab, Match-Up 2019, and in particular Haris Aziz, Jérôme Lang, Julien Lesca, Jordi Massó, Hervé Moulin, Rosemarie Nagel, David Pérez-Castrillo, Hans Peters, Madhav Raghavan, Tayfun Sönmez, William Thomson, and Utku Ünver for their comments.

†Institute of Economics, Research Centre for Economic and Regional Studies, Hungarian Academy of Sciences, and Department of Operations Research and Actuarial Sciences, Corvinus University of Budapest; e-mail: peter.biro@krtk.mta.hu. P. Biró gratefully acknowledges the financial support by the Hungarian Academy of Sciences, Momentum Grant No. LP2021-2.

‡Institute for Economic Analysis (CSIC) and Barcelona GSE; e-mail: flip.klijn@iae.csic.es. F. Klijn gratefully acknowledges financial support from AGAUR–Generalitat de Catalunya (2017-SGR-1359) and the Spanish Agencia Estatal de Investigación (AEI) through grants ECO2017-88130-P and PID2020-114251GB-I00 (funded by MCIN/ AEI /10.13039/501100011033) and the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics, SEV-2015-0563 and CEX2019-000915-S).

§Department of Economics, Concordia University; e-mail: szilvia.papai@concordia.ca. S. Pápai gratefully acknowledges financial support from an FRQSC grant titled "Formation des coalitions et des réseaux dans les situations économiques et sociales avec des externalités."

the family of STC rules collapses to the cTTC rule which then is also group-strategy-proof. Finally, we characterize one particularly interesting STC rule by means of top unanimity and self-enforcing group-strategy-proofness.

Keywords: indivisible goods, circulation, top trading cycles, strategy-proofness, Pareto-efficiency
JEL classification: C71, C78

1 Introduction

In the classical "housing market" of Shapley and Scarf [29] each agent owns one indivisible good and monetary compensations are not possible. Each agent has preferences over all goods and is interested in having exactly one good. What would be a "good" redistribution of the goods? A well-studied rule in this setting is David Gale's Top Trading Cycle (TTC) rule. Roth and Postlewaite [24] showed that for any profile of strict preferences, the TTC rule yields the unique strong core element. Hence, the TTC rule is individually rational and Pareto-efficient. Roth [23] proved that the TTC rule is strategy-proof, i.e., for each agent it is a weakly dominant strategy to reveal her true preferences. Ma [14] showed in fact that on the domain of strict preferences the TTC rule is the unique rule that satisfies individual rationality, Pareto-efficiency, and strategy-proofness. But what happens if agents are endowed with multiple indivisible goods that they want to exchange in a one-for-one fashion?

Multi-unit Shapley-Scarf markets

Real-life examples of such markets abound. In the European Erasmus program, higher education institutions exchange their students for a semester or longer.[1] An important restriction is that for each involved institution the number of students it sends out equals the number of students it receives.

Another example is time banks. Each participant in a time bank offers some service (haircutting, teaching, babysitting, etc.). Hours of service are exchanged in a one-for-one fashion without the intermediation of money. Usually there are bilateral agreements or there is a credit system, where credits are earned by providing services, and which can be used later to request services, possibly from a different participant. In the latter case, balancedness can be obtained through central coordination.

Other similar barter exchange markets include timesharing networks where users offer their homes or holiday apartments in exchange for spending the same amount of time at another place. International deceased donor sharing schemes, such as Eurotransplant[2] or Scandiatransplant,[3] aim to efficiently exchange the deceased donors of participating countries. Since any available organ has to be used immediately, the exchange of grafts is dynamic. However, in the long run these schemes ensure that each country sends and receives about the same number of organs for transplantation. Thus, given a sufficiently long time period, the exchange is approximately balanced.

[1] The EuRopean community Action Scheme for the Mobility of University Students involves currently more than 4,000 higher institutions across 37 countries, see, e.g., https://ec.europa.eu/education/ and in particular https://web.archive.org/web/20160305053245/http://ec.europa.eu/education/library/statistics/ay-12-13/facts-figures_en.pdf

[2] https://www.eurotransplant.org/

[3] http://www.scandiatransplant.org/

A stylized model

In this paper, we study a stylized model of the previously described markets. As a first but non-trivial step towards a full-fledged analysis we consider the case in which each agent has multiple units of an indivisible agent-specific good. We refer to the number of units of the good an agent is endowed with as her capacity. For instance, in the setting of time banks, an agent's capacity would represent the maximum number of hours of service the agent can give. A desideratum (if not a requirement) that appears to be common to many of the markets above is that any exchange be balanced, i.e., each agent receives as many goods/service hours as she gives away. Thus, we conveniently employ the graph theoretical notion of circulation to study balanced exchanges.

In our model, agents only have preferences over the possible bundles they can receive. We assume that each agent's preferences over bundles are "responsive" with respect to her preferences over individual goods. Indeed, in some applications it is unlikely that all agents have lexicographic or even additive preferences. Responsiveness allows for a richer extension of the preferences over individual goods[4] and induces substantially different results as compared to lexicographic preferences.

In centrally coordinated markets, agents can often only reveal their ordinal preferences over the individual goods. For this reason we study circulation rules that take these ordinal preferences and the capacities of the agents as an input and return a circulation. Considering individual rationality an indispensable property of a circulation rule, we focus our attention on its compatibility with two important desiderata: Pareto-efficiency and strategy-proofness. As pointed out earlier, these properties are simultaneously achievable by the TTC rule in the classical Shapley-Scarf setting. However, we run into an impossibility in our more general model. Next, we discuss in more detail this result as well as our main findings.

Our main contributions

Our first result is that for general capacity configurations there exists no individually rational circulation rule that is both Pareto-efficient and strategy-proof (Proposition 1). The proof of this incompatibility is based on a three-agent market where just one agent has a capacity larger than 1 and all preferences are lexicographic. On the domain of responsive preferences this impossibility holds for all so-called "reducible" capacity configurations (Proposition 2). For irreducible capacity configurations a natural generalization of the TTC rule that we call the Circulation Top Trading Cycle (cTTC) rule is Pareto-efficient (Corollary 1) as well as strategy-proof (Corollary 2). In fact, for irreducible capacity configurations the TTC rule is the unique individually rational rule that satisfies both properties (Theorem 2). We complete our analysis of the cTTC rule by studying its incentive properties on the domain of responsive preferences (Propositions 5 and 6) and providing a characterization on the domain of lexicographic preferences where the cTTC rule is the unique individually rational rule that is Pareto-efficient and "dropping-proof" (Theorem 1).

Next, in Section 5, we introduce and study the family of so-called Segmented Trading Cycle (STC) rules which are obtained by extending the standard TTC rule from the classical (i.e., unit capacity) setting to our setting in a different way than the cTTC rule. More specifically, the STC rules are obtained by first distributing agents' endowments over a number of different smaller markets (the market segments), then applying the standard Top Trading Cycle algorithm within

[4] Yet, responsiveness still does not allow for certain complementarities between goods.

each market segment separately, and finally lumping together the induced circulations. We show that STC rules are individually rational, strategy-proof, and nonbossy (Proposition 9). Even though STC rules do not satisfy group-strategy-proofness in general (Proposition 10), they do satisfy a weaker notion of group-strategy-proofness (Proposition 12). Moreover, for irreducible capacity configurations the family of STC rules collapses to the cTTC rule (Proposition 8) which then is also group-strategyproof for lexicographic preferences (Proposition 14). One particularly interesting member of the class of STC rules is the Sequentially Segmented Trading Cycle (SSTC) rule which sequentially fills up the market segments with goods as much as possible, i.e., first segment 1, then segment 2, and so on. Our final main results are characterizations of the SSTC rule by top unanimity and self-enforcing group-strategy-proofness (Theorem 3) and top unanimity, strategy-proofness, and nonbossiness (Corollary 3). A rule satisfies top unanimity if for each preference profile it "respects" top trades, i.e., it is consistent with the first round of the cTTC rule. Self-enforcing group-strategy-proofness is a weaker version of group-strategy-proofness that essentially rules out all self-enforcing manipulations (where no agent would prefer to revert back unilaterally to the true preferences).

Related literature

As mentioned in the introduction, the earliest results (Roth [23] and Roth and Postlewaite [24]) on the housing market of Shapley and Scarf [29] established that on the domain of strict preferences the TTC rule is individually rational, Pareto-efficient, and strategy-proof. Ma [14] showed in fact that the TTC rule is the unique rule that satisfies these three properties.[5]

Sönmez [30] studied a general class of allocation problems that includes the class of housing markets. He proved that if there is a individually rational, Pareto-efficient, and strategy-proof rule, then the strong core correspondence is essentially single-valued and the rule is a selection from this correspondence. He also obtained an impossibility result for the model of exchange of multiple indivisible goods, namely that if at least one agent owns more than one good then there is no allocation rule that is individually rational, Pareto-efficient, and strategy-proof. Konishi et al. [11] proved a similar impossibility result in a model with market segments (cars, houses, etc.) where each agent owns exactly one good in each market segment. Todo et al. [34] also obtained an impossibility result for a particular setting where agents have multiple heterogeneous indivisible goods and lexicographic preferences. Sonoda et al. [31] considered a model with multiple heterogeneous indivisible goods where agents' preferences can have indifferences. They determined for different preference domain restrictions whether the impossibility result of Sönmez [30] still holds or fails to hold. Our impossibility result (Proposition 1) strengthens these impossibility results as we establish our result on the domain of lexicographic preferences with homogeneous goods. Moreover, our characterization of the capacity configurations for which the impossibility holds under responsive preferences (Theorem 2) appears to be a novel sort of result in this literature.

Given the above described impossibility results, several studies restored strategy-proofness by weakening the requirement of Pareto-efficiency. Pápai [19] introduced the Segmented Trading Cycle (STC) rules in a general model with multiple heterogeneous indivisible goods. In the STC rules, each trade is one-for-one and cannot cross market segments. Pápai showed that the STC rules are the only exchange rules that satisfy strategy-proofness, nonbossiness, trade sovereignty,

[5]Pycia [22] extended the result to a model that includes a network with directed edges that represent feasibility of shipment.

and strong individual rationality. Todo et al. [34] introduced variants of the STC rules for a model with multiple heterogeneous indivisible goods under different restrictions on the preference domain. The above models do not require the solutions to be balanced, but nevertheless the STC rules always produce balanced exchanges, so they are naturally suitable solution concepts for our circulation model. In Section 5, we explore the properties of the STC rules in our circulation model with homogeneous goods and responsive preferences.

Fujita et al. [8] followed up the paper of Todo et al. [34] and studied a generalized TTC algorithm, called augmented top trading cycles (ATTC), for a model with multiple heterogeneous indivisible goods under lexicographic preferences. They showed that the ATTC rule is strong core selecting (hence Pareto-efficient) but also vulnerable to strategic manipulation. On the other hand, they proved that the ATTC rule is NP-hard to manipulate, which can be interpreted as a cognitive barrier for possible manipulators. Their results do not automatically carry over to our model with homogeneous goods. However, according to Lesca [13] the latter NP-hardness reduction can be adjusted for our cTTC rule in our model.

Dur and Ünver [6] studied balanced rules for a special exchange problem motivated by tuition and worker exchanges. Their model is two-sided; students/workers are also strategic players. Moreover, each college/firm has (responsive) preferences over the set of students/workers they may be assigned to as well as an "internal" priority order over their own students/workers which may be sent off. They studied and characterized a two-sided TTC rule which is an analogue to our cTTC rule.

For the model that we introduce in the next section, we studied in our companion paper [4] so-called single-serial and multiple-serial rules, which are the classes of rules where agents in turn choose one unit of some good or a bundle of goods, respectively. To guarantee individual rationality, at each step the set of available goods is restricted. We showed that the two resulting classes of rules are Pareto-efficient on the domain of lexicographic and responsive preferences, respectively, but they necessarily violate strategy-proofness. Cechlárová et al. [5] studied serial dictatorship rules in a related course allocation model with multiple homogeneous goods and budget constraints under lexicographic preferences. Earlier characterizations of more general serial dictatorship rules were obtained by Pápai [18] and Ehlers and Klaus [7].

Andersson et al. [1] studied time banks, where agents exchange their services/time in a one-for-one fashion. They considered dichotomous preferences over individual services and extended them to bundles in a specific way. In particular, agents have upper quotas representing their maximum needs for each service. This paper introduced so-called priority mechanisms that select individually rational and time-balanced allocations that are Pareto-efficient and maximize exchanges. They showed that priority mechanisms are strategy-proof. The same rules were also characterized by Hatfield [9] for the allocation of multiple indivisible goods under additive preferences in a model without initial endowments and where agents' capacities must be filled exactly.

Manjunath and Westkamp [15] studied the exchange of heterogeneous goods without monetary compensations. They assumed that each agent has trichotomous preferences in the sense that an agent partitions goods into desirable, undesirable but in the agent's endowment, and remaining undesirable goods. They introduced the so-called Individually Rational Priority (IRP) mechanisms and showed that IRP mechanisms are individually rational, Pareto-efficient, and strategy-proof.

Finally, we note that there are many papers that study special cases of Shapley and Scarf's [29] original setting (where each agent has one indivisible good only). This literature studies markets that range from kidney exchange (Roth et al. [25]) to timeshare exchange (Wang and Krishna [35]).

The remainder of the paper is organized as follows. In Section 2, we introduce the circulation model. In Section 3, we state and prove the first incompatibility results regarding individual rationality, Pareto-efficiency, and strategy-proofness. In Sections 4 and 5 we introduce and study the Circulation Top Trading Cycles rule and the Segmented Trading Cycle rules, respectively. Section 6 provides a final outlook.

2 The circulation model

Let N with $n = |N| \geq 2$ be the set of agents. Each agent $i \in N$ is endowed with a set e_i of indivisible, homogeneous, and agent-specific goods. The non-negative integer $q_i = |e_i| \in \mathbb{N}_+$ denotes agent i's capacity. Let $q = (q_i)_{i \in N}$. The pair (n, q) is the **capacity configuration**. Since goods are agent-specific, for any $i \in N$, we often refer to the good of agent i as good i.

Each agent i has preferences $\succ_i$ over all individual goods, i.e., preferences over receiving a unit of good $j \in N \backslash \{i\}$ and the option of retaining a unit of her good i. We assume that $\succ_i$ is a linear order on N, i.e., it is strict, complete, and transitive. For any $j, l \in N \backslash \{i\}$ with $j \neq l$, $j \succ_i l$ denotes that agent i prefers receiving one unit of good j over receiving one unit of good l. Let $\succeq_i$ denote the weak counterpart of $\succ_i$, i.e., $j \succeq_i l$ if and only if $j \succ_i l$ or $j = l$. If $j \succeq_i i$, then good j is acceptable for agent i; otherwise it is unacceptable for i. Let $A_i \equiv A_i^{\succ_i}$ denote the set of acceptable goods for agent i. For each agent $i \in N$, let $\mathcal{L}_i$ denote the set of strict preferences over individual goods for agent i. Let $\mathcal{L} = \times_{i \in N} \mathcal{L}_i$ be the set of profiles of strict preferences over individual goods.

A bundle for agent i is a vector $x_i = (x_{ij})_{j \in N \backslash \{i\}}$ with $\sum_{j \in N \backslash \{i\}} x_{ij} \leq q_i$. We often refer to the latter inequality as the **feasibility** (constraint) of the bundle. Here, $x_{ij} \in \mathbb{N}_+$ is the number of units of good j that are sent to / received by agent i. One particular bundle for agent i is the null bundle 0_i where agent i receives no good from any other agent, i.e., $0_{ij} = 0$ for all $j \neq i$. Let X_i denote the set of possible bundles for agent i.

Agent i has a linear order P_i on X_i. A bundle x_i is acceptable for i if $x_i P_i 0_i$ or $x_i = 0_i$; it is unacceptable for i otherwise. Let R_i denote the weak counterpart of P_i. So, $x_i R_i x_i'$ if either $x_i P_i x_i'$ or $x_i = x_i'$.

We assume that the preferences P_i over X_i are a **responsive** extension of the associated preferences $\succ_i$ over individual goods. Formally, P_i is a linear order that satisfies the following two properties. Let $x_i', x_i \in X_i$.

(resp-1). $x_i' P_i x_i$ if $x_i R_i 0_i$ and there is $j \in N \backslash \{i\}$ with $j \succ_i i$ such that

$$x_{ij}' = x_{ij} + 1 \text{ and } x_{ik}' = x_{ik} \text{ for all } k \in N \backslash \{i, j\};$$

(resp-2). $x_i' P_i x_i$ if $x_i R_i 0_i$ and there are $j, l \in N \backslash \{i\}$ with $j \succ_i l$ such that

$$x_{ij}' = x_{ij} + 1, x_{il}' = x_{il} - 1, \text{ and } x_{ik}' = x_{ik} \text{ for all } k \in N \backslash \{i, j, l\}.$$

Condition (resp-1) states that agent i prefers bundle x_i' to an acceptable bundle x_i if in x_i' she receives one more unit of some acceptable good than in x_i and the same number of units of all other goods. This property is also referred to as *separability* in the literature. Condition (resp-2) states that agent i prefers bundle x_i' to an acceptable bundle x_i if x_i' is obtained from x_i by replacing one unit of some good with one unit of a more preferred good (both goods being different from agent i's own good).

Remark 1. Note that if a bundle only contains acceptable goods for some agent, then, by repeated application of (resp-1), the agent finds the bundle acceptable. However, there can be acceptable bundles that contain unacceptable goods. $\diamond$

An additional assumption can also be imposed on preferences, in addition to (resp-1) and (resp-2), which rules out preferences that allow for unacceptable goods in acceptable bundles.

(resp-3). The preferences P_i of each agent $i \in N$ are such that for all $x_i \in X_i$, $0_i P_i x_i$ if for some $j \in N \backslash A_i$, $x_{ij} > 0$.

Condition (resp-3) is a property of "absolute desirability": it states that an agent finds a bundle unacceptable if it contains some good that is unacceptable for her. Together with the first part of Remark 1 this implies that a bundle is acceptable if and only if it only contains acceptable goods.

Throughout the paper, the results hold both with or without imposing condition (resp-3). We focus on the responsive preference domain in the exposition which does not require (resp-3), since it is a larger and more realistic preference domain, but we note that the impossibility results in Section 3 are stronger when condition (resp-3) also holds.

For each agent i and any responsive preferences P_i, let X_i^A denote the set of bundles that only contain goods that are acceptable for i. Formally, for $x_i \in X_i$, $x_i \in X_i^A$ if and only if for each $j \in N$ with $x_{ij} > 0$, $j \in A_i$.

We denote the set of responsive preferences for agent i by $\mathcal{P}_i$. Let $\mathcal{P} = \times_{i \in N} \mathcal{P}_i$ be the set of profiles of responsive preferences. A market is a triple (N, q, P) where $P \in \mathcal{P}$. For any responsive preferences $P_i \in \mathcal{P}_i$ of agent i, we denote the underlying preferences over individual goods by $\succ^{P_i}$. For any $P \in \mathcal{P}$, $\succ^P = (\succ^{P_i})_{i \in N}$. Whenever no confusion is possible we write $\succ_i$ for $\succ^{P_i}$ and $\succ$ for $\succ^P$.

Next, we introduce the class of additive preferences and the class of lexicographic preferences. Agent i's responsive preferences P_i are **additive** if there is a utility function $u_i : A_i \to \mathbb{R}$ such that[6] $u_i(i) = 0$ and

$$\text{for all } x_i, x_i' \in X_i^A, \quad [x_i' \, P_i \, x_i \ \text{ if and only if } \sum_{j \in A_i} x_{ij}' u_i(j) > \sum_{j \in A_i} x_{ij} u_i(j)]. \tag{1}$$

Note that we only impose conditions on bundles that solely consist of acceptable goods.[7] Moreover, responsiveness does not imply additivity.[8] We denote the set of additive preferences for agent i by $\mathcal{P}_i^A$. Let $\mathcal{P}^A = \times_{i \in N} \mathcal{P}_i^A$ be the set of profiles of additive preferences.

Agent i's responsive preferences P_i are **lexicographic** if there is a utility function $u_i : A_i \to \mathbb{R}$ such that (1) holds and

$$\text{for all } k, l \succ_i i, \quad [k \succ_i l \ \text{ if and only if } \ u_i(k) > q_i u_i(l)].$$

[6]The assumption that $u_i(i) = 0$ is without loss of generality and sets the utility of the null bundle at 0.

[7]The reason for this becomes clear in Remark 2.

[8]For instance, suppose $N = \{1, 2, 3, 4, 5, 6\}$ and $q_1 = 3$. Then, there are responsive preferences P_1 such that each good in $N \backslash \{1\}$ is acceptable for agent 1, $(1, 0, 0, 1, 0) P_1 (0, 1, 1, 0, 0)$, and $(0, 1, 1, 0, 1) P_1 (1, 0, 0, 1, 1)$. (Here, a bundle is written as a vector the entries of which indicate the quantity of goods $2, \ldots, 6$.) Any such P_1 is not additive.

By definition, all lexicographic preferences are additive. (But, obviously, not all additive preferences are lexicographic.) In the case of lexicographic preferences, the ordinal ranking over bundles *that only contain acceptable goods* is completely determined by the ordinal ranking over individual goods. For this reason we will often refer to lexicographic preferences as **lexicographic extensions** of the preferences over individual goods. We denote the set of lexicographic preferences for agent i by $\mathcal{P}_i^L$. Let $\mathcal{P}^L = \times_{i \in N} \mathcal{P}_i^L$ be the set of profiles of lexicographic preferences. Note $\mathcal{P}_i^L \subsetneq \mathcal{P}_i^A \subsetneq \mathcal{P}_i$, and hence $\mathcal{P}^L \subsetneq \mathcal{P}^A \subsetneq \mathcal{P}$.

We require the exchange of the indivisible goods be balanced. In other words, any outcome of a market should be a circulation, i.e., a vector of bundles such that each agent receives as many goods as she gives away from her initial endowment. For any vector of bundles $x = (x_i)_{i \in N} \in (X_i)_{i \in N}$, let $x_{i\cdot} = \sum_{j \in N \setminus \{i\}} x_{ij}$ denote the number of goods agent i receives from the other agents. Similarly, let $x_{\cdot i} = \sum_{j \in N \setminus \{i\}} x_{ji}$ denote the number of goods agent i sends to the other agents. A **circulation** is a vector of bundles $x = (x_i)_{i \in N} \in (X_i)_{i \in N}$ such that for each agent $i \in N$, $x_{i\cdot} = x_{\cdot i}$. We often refer to the latter equalities as the **balancedness** (constraints) of the circulation. If for agent i and circulation x, $x_{i\cdot} = q_i$, then we say that i is **filled** at x; otherwise (i.e., $x_{i\cdot} < q_i$), we say that i is **unfilled** at x. Let X denote the set of circulations. To describe a circulation x we usually only specify the strictly positive exchanges, i.e., the integers x_{ij} with $x_{ij} > 0$. Finally, at any circulation x, let $x_{ii} \equiv q_i - x_{\cdot i} \geq 0$ denote the number of goods i agent i keeps (i.e., does not send to other agents).

Circulation rules

Our aim is to study rules that can be used by a centralized clearinghouse. In practice such kind of clearinghouse often does or would collect only the ordinal preferences of the participating agents over individual goods. Moreover, given our assumption that preferences are responsive, the most important information about preferences is concisely summarized by the ranking of individual goods. For this reason we introduce the following definition of circulation rule.

Fix the set of agents N and the capacity configuration (n, q). A **circulation rule** $f : \mathcal{P} \to X$ specifies a circulation for each preference profile. For any preference profile $P \in \mathcal{P}$, $f_i(P)$ denotes agent i's bundle at P. We assume all circulation rules to be **individual-good-preference based** in the sense that for any two preference profiles, if each agent has the same underlying ordinal preferences over individual goods at both profiles, then a circulation rule yields the same circulation at both profiles. Formally,

$$\text{for all } P, P' \in \mathcal{P} \text{ with } \succ^P = \succ^{P'}, \quad f(P) = f(P'). \tag{2}$$

Below we first introduce the key desiderata that we consider in the next sections. The property that we consider indispensable is individual rationality. This standard property requires that each agent receives a bundle that is acceptable for her.

Definition 1. A circulation x is individually rational for agent $i \in N$ at $P \in \mathcal{P}$ if x_i is acceptable, i.e., $x_i R_i 0_i$. A circulation x is individually rational at $P \in \mathcal{P}$ if it is individually rational for all agents at P. A circulation rule f is **individually rational** if for all $P \in \mathcal{P}$, $f(P)$ is individually rational at P. ◇

Remark 2. As noted in Remark 1, an acceptable bundle may contain unacceptable goods. However, any individually rational circulation rule always assigns bundles that only consist of acceptable goods. To see this note, let $i \in N$. Let $P \in \mathcal{P}$. Denote $P_{-i} \equiv (P_j)_{j \neq i}$. Let $\succ_i$ be agent i's associated preferences over individual goods. Then, there is a responsive (in fact even lexicographic) extension $\tilde{P}_i$ of the preferences $\succ_i$ such that all acceptable bundles only contain acceptable goods for i. Hence, by individual rationality of f, $f_i(\tilde{P}_i, P_{-i})$ consists of acceptable goods for i. Since from (2) it follows that $f_i(P) = f_i(\tilde{P}_i, P_{-i})$, $f_i(P)$ consists of acceptable goods for i. $\diamond$

Given the relatively simple structure and the particular interest of lexicographic preferences within the class of responsive preferences we will examine two different versions of each of the axioms where applicable: "necessarily satisfied" and "possibly satisfied," indicating whether the axiom holds for every responsive extension of the underlying preferences over individual goods, or only to the lexicographic extensions that can be inferred from the ordering of individual goods. Thus, "necessarily satisfied" corresponds to the axiom being satisfied by the entire responsive preference domain and is the standard version of the axiom for responsive preferences over bundles. The "possibly satisfied" version is weaker; namely, it corresponds to the axiom being satisfied by the lexicographic extensions of any preferences over the individual goods. Henceforth, we will denote the weaker version of each axiom by adding the prefix "**ig**" (the acronym for "individual good") to the name of the standard version of the axiom.[9]

Definition 2. A circulation x is Pareto-dominated by another circulation y at $P \in \mathcal{P}$ if $y_i R_i x_i$ for all agents $i \in N$ and $y_j P_j x_j$ for some agent $j \in N$. A circulation rule f is (necessarily) **Pareto-efficient** if for all $P \in \mathcal{P}$, $f(P)$ is not Pareto-dominated by any other circulation at P. A circulation rule f is **ig-Pareto-efficient** if for all profiles of lexicographic preferences $P \in \mathcal{P}^L$, $f(P)$ is not Pareto-dominated by any other circulation at P. $\diamond$

Definition 3. Agent $i \in N$ can manipulate circulation rule f at $P \in \mathcal{P}$ if there exists a deviation $P'_i \in \mathcal{P}_i$ such that $f_i(P'_i, P_{-i}) P_i f_i(P)$. A circulation rule f is (necessarily) **strategy-proof** if no agent can manipulate f at any $P \in \mathcal{P}$. A circulation rule f is **ig-strategy-proof** if no agent can manipulate f at any profile of lexicographic preferences $P \in \mathcal{P}^L$.[10] $\diamond$

Since the three key properties above will turn out to be incompatible for a large class of capacity configurations, we will also study the following (weaker) incentive properties. A preference $\succ'_i \in \mathcal{L}_i$ is a truncation of $\succ_i \in \mathcal{L}_i$ if for any $k, l \in N$, [if $k \succeq'_i l \succeq'_i i$, then $k \succeq_i l \succeq_i i$] and [if $k \succ'_i i$ and $l \succ_i k$, then $l \succ'_i i$].

Definition 4. Agent $i \in N$ can manipulate circulation rule f at $P \in \mathcal{P}$ by means of truncation if there exists a deviation $P'_i \in \mathcal{P}_i$ such that $\succ^{P'_i}$ is a truncation of $\succ^{P_i}$ and $f_i(P'_i, P_{-i}) P_i f_i(P)$. A circulation rule f is (necessarily) **truncation-proof** if no agent can manipulate f at any $P \in \mathcal{P}$ by means of truncation.[11] $\diamond$

A preference $\succ'_i \in \mathcal{L}_i$ is a dropping of $\succ_i \in \mathcal{L}_i$ if for any $k, l \in N$, [if $k \succeq'_i l \succeq'_i i$, then $k \succeq_i l \succeq_i i$]. Obviously, any truncation is a dropping.

[9] By Remark 1, the two versions of individual rationality are equivalent.

[10] Since circulation rules are individual-good-preference based, equivalent definitions of strategy-proofness and ig-strategy-proofness are obtained by additionally demanding that the deviation P'_i be lexicographic.

[11] Kojima [10] similarly defined "non-manipulability via truncation" in the context of resource allocation with multi-unit demand. Moreover, since circulation rules are individual-good-preference based, an equivalent definition of truncation-proofness is obtained by additionally demanding that the deviation P'_i be lexicographic.

Definition 5. Agent $i \in N$ can manipulate circulation rule f at $P \in \mathcal{P}$ by means of dropping if there exists a deviation $P_i' \in \mathcal{P}_i$ such that $\succ^{P_i'}$ is a dropping of $\succ^{P_i}$ and $f_i(P_i', P_{-i}) P_i f_i(P)$. A circulation rule f is (necessarily) **dropping-proof** if no agent can manipulate f at any $P \in \mathcal{P}$ by means of dropping. A circulation rule f is **ig-dropping-proof** if no agent can manipulate f by means of dropping at any profile $P \in \mathcal{P}^L$ of lexicographic preferences.[12] $\diamond$

Note that dropping-proofness implies truncation-proofness. Finally, in the next definition, we will express the circulation outcome explicitly as a function of the capacity configuration, in addition to the preference profile. Let (n, q) be a capacity profile and let $i \in N$. We denote $q_{-i} \equiv (q_j)_{j \neq i}$.

Definition 6. A circulation rule f is **hiding-proof**[13] if for all $i \in N$, $P \in \mathcal{P}$, and $q_i' < q_i$, $f_i(P, q)\ R_i\ f_i(P, (q_{-i}, q_i'))$. $\diamond$

The underlying directed graph

The notion of circulation is well-studied in graph theory, see e.g. Schrijver [27]. To clarify the link to that literature let $D(N, A(P))$ denote the directed graph where N is the set of nodes and $A(P)$ is the set of arcs (directed edges) such that $(i, j) \in A(P)$ if and only if $j \succ^{P_i} i$.

An individually rational circulation is a non-negative function on the arcs, $x : A(P) \to \mathbb{N}_+$, where $x(i, j)$ denotes the flow from j to i, such that each node satisfies flow conservation (i.e., the incoming flow equals the outgoing flow) and incoming/outgoing flow does not exceed the capacity of the corresponding agent. Setting $x_{ij} \equiv x(i, j)$ for all i, j with $i \neq j$, gives a straightforward one-to-one correspondence between circulations as vectors of bundles and circulations as non-negative functions on the arcs.

3 Impossibility results

In this section we show that individual rationality, strategy-proofness, and Pareto-efficiency are not compatible.[14] First we prove this result for a simple market with three agents and lexicographic preferences (Proposition 1), where only one agent has capacity two and the other two agents have unit capacity. Although this result implies an incompatibility for the more general domain of responsive preferences, we will study that domain in detail as well. The reason is that the impossibility result only holds for a certain class of capacity configurations. More specifically, we will first show that for so-called *reducible* capacity configurations the impossibility result holds (Proposition 2). Next, in Section 4, we will complement this result by showing that for irreducible

[12] Since circulation rules are individual-good-preference based, equivalent definitions of dropping-proofness and ig-dropping-proofness are obtained by additionally demanding that the deviation P_i' be lexicographic.

[13] In the context of classical exchange economies, Postlewaite [21] was the first to introduce and study "non-manipulability by withholding."

[14] Note that individual rationality is compatible with Pareto-efficiency and strategy-proofness separately. For instance, Biró et al. [4] studied rules that are individually rational and Pareto-efficient. Since [4] uses the assumption of (resp-3), showing the existence of such rules without this assumption requires an adjustment to the rules in [4]. Details of this relatively straightforward adjustment are available from the authors upon request. Moreover, the Segmented Trading Cycle rules of Section 5 are individually rational and strategy-proof on both preference domains.

capacity configurations there is a unique rule (the cTTC) that satisfies all three properties (Theorem 2) and in Section 5 we revisit this theorem in order to show that this unique rule can also be derived from a different rule (the SSTC) when restricted to irreducible capacity configurations (Theorem 2 revisited).

Before stating the impossibility results in this section, recall that all the results hold with the additional assumption of (resp-3) on the preferences, which makes these impossibility results stronger, since they hold on a smaller preference domain.

Proposition 1. *There are capacity configurations for which there is no circulation rule that is individually rational, ig-Pareto-efficient, and ig-strategy-proof.*

Proof. We first note that it can be easily verified that all preferences that are considered and constructed in the proof can be assumed/taken to satisfy (resp-3).

Let $(n, q) = (3, (1, 2, 1))$. Suppose for a contradiction that there exists an individually rational rule f that is both ig-Pareto-efficient and ig-strategy-proof. Consider the market (N, q, P) where $N = \{a, b, c\}$, $(q_a, q_b, q_c) = (1, 2, 1)$, and P is a profile of lexicographic extensions of the following preferences $\succ$ over individual goods:[15]

$\succ_a$: $b \succ_a a$
$\succ_b$: $c \succ_b a \succ_b b$
$\succ_c$: $a \succ_c b \succ_c c$

The corresponding graph $D(N, A(P))$ is depicted in Figure 1. Edges correspond to acceptable goods. For instance, the edge from c to a shows that agent c finds good a acceptable. Similarly, there is no edge from a to c as agent a finds good c unacceptable. Continuous edges denote most preferred goods and discontinuous edges denote second most preferred goods.

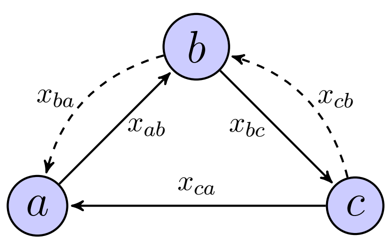

Figure 1: $\succ$ and $D(N, A(P))$

One easily verifies that by individual rationality and ig-Pareto-efficiency, $f(P) \in \{x, x'\}$ where x and x' are the circulations given by $x : x_{ab} = x_{bc} = x_{ca} = 1$ and $x' : x'_{ab} = x'_{ba} = 1, x'_{bc} = x'_{cb} = 1$, respectively.

Suppose $f(P) = x$. Now let b report lexicographic preferences P'_b such that $\succ'_b \equiv \succ^{P'_b}$ is given by $a \succ'_b c \succ'_b b$. At the new preference profile $P' = (P'_b, P_{-b})$ the only two individually rational and Pareto-efficient circulations are still x and x'. If $f(P') = x'$ then b can manipulate at P via P'_b. If $f(P') = x$ then b can manipulate at P' by reporting lexicographic preferences P''_b where $\succ''_b \equiv \succ^{P''_b}$ is given by $a \succ''_b b$, because at the resulting profile $P'' = (P''_b, P_{-b})$ the unique individually rational and Pareto-efficient circulation is $x'' : x''_{ab} = x''_{ba} = 1$. This contradicts ig-strategy-proofness of f.

[15]Because of individual rationality we always omit unacceptable goods from the description of the preferences.

Suppose now $f(P) = x'$. Then c can manipulate at P by reporting lexicographic preferences P_c^* such that $\succ_c^* \equiv \succ^{P_c^*}$ is given by $a \succ_c^* c$, since the only individually rational and Pareto-efficient circulation at $P^* = (P_c^*, P_{-c})$ is x. This contradicts again ig-strategy-proofness of f. Hence, there is no individually rational rule that is both ig-Pareto-efficient and ig-strategy-proof. □

In the proof of Proposition 1 we made use of a particular capacity configuration, namely $(n, q) = (3, (1, 2, 1))$. The incompatibility of Pareto-efficiency and strategy-proofness holds for "many" other capacity configurations as well. In fact, our next result gives a sufficient condition for this incompatibility– the necessity of this condition is established in the next section.

A capacity configuration (n, q) is called **reducible** if it satisfies one of the following two conditions:

(rd1) $n \geq 3$ and there are three different agents $a, b, c \in N$ with $q_b > q_c \geq q_a$;
(rd2) $n \geq 4$ and there are four different agents $a, b, c, d \in N$ with $q_a = q_b = q_c > q_d$.

A capacity configuration is irreducible if it is not reducible.

Note that the capacity configuration in the proof of Proposition 1 is reducible as it satisfies condition (rd1).

Proposition 2. *If the capacity configuration is reducible, then there is no circulation rule that is individually rational, Pareto-efficient, and strategy-proof.*

A direct proof of Proposition 2 is relegated to Appendix A. A much shorter proof is provided at the end of Section 5, which relies on the characterizations of the two main rules that we propose, cTTC (Theorem 1) and SSTC (Theorem 3).

In view of the results in the upcoming sections it is convenient to note (and not difficult to verify) that a capacity configuration (n, q) is **irreducible** if and only if it satisfies one of the following three conditions:

(ird1) $n = 2$;
(ird2) $n = 3$ and there are three different agents $a, b, c \in N$ with $q_b = q_c > q_a$;
(ird3) $n \geq 3$ and all agents have identical capacity.

We call irreducible capacity configurations "irreducible" because, if we start with an irreducible capacity configuration, when we reduce the available units for any subset of the agents $S \subseteq N$ with $|S| \geq 2$ by the same number and at least one of these agents reaches exactly zero units, we still get an irreducible capacity configuration. This holds only when the capacity configuration is irreducible.

4 The Circulation Top Trading Cycle rule

In this section we introduce the Circulation Top Trading Cycle (cTTC) rule which is a circulation rule based on David Gale's Top Trading Cycle algorithm. We will show that this rule satisfies individual rationality, Pareto-efficiency, and strategy-proofness for irreducible capacity configurations. In view of Proposition 2, this also shows that irreducibility is a necessary and sufficient condition on the capacity configuration for the compatibility of the three properties. Our second main result in this section is a characterization: on the domain of lexicographic preferences, the cTTC rule is the unique individually rational rule that is both Pareto-efficient and dropping-proof (Theorem 1). As a consequence, the cTTC rule is the *unique* rule that satisfies individual rationality, Pareto-efficiency, and strategy-proofness on the domain of responsive preferences when capacity configurations are irreducible (Theorem 2).

4.1 Description of cTTC

In the first step of this algorithm let each agent point to her most preferred good, or equivalently, potential trading partner. If there is no such trading partner then the agent points to herself. There is at least one "top trading" cycle. Let the agents in each top trading cycle send as many goods as possible such that the number of goods exchanged within the cycle is the same for each agent (i.e., maximize the flow subject to capacity restrictions). If the agent points to herself then she receives her own remaining endowments. We decrease the capacity of each agent in a cycle according to the number of goods traded in that cycle (i.e., the flow). In each top trading cycle, there is at least one agent that exhausts all of her endowment (or in other words, at least one node in the graph becomes saturated). All such agents are now removed from the market (i.e., all saturated nodes are removed from the graph), and we repeat the same process in the remaining market (reduced graph) until all agents are removed.

We will refer to this algorithm as the cTTC algorithm and denote the corresponding **cTTC rule** by τ. Note that the cTTC algorithm can be equivalently executed and described by restricting the flow of each cycle in each step to one unit. In the one-unit-per-cycle (version of the) algorithm, it is not necessarily true that in each step some agent is removed; multiple steps with identical cycles may be required to exhaust the capacity of some agent.

Note that the cTTC rule is individually rational. Moreover, when all agents have unit capacity, i.e., $q_i = 1$ for all agents i, the cTTC algorithm is the original top trading cycles algorithm. In fact, when all agents have the same capacity, say q^*, the cTTC algorithm gives each agent the bundle that consists of q^* units of the good she receives at the circulation obtained from the TTC algorithm applied to the economy with unit capacities and the same preferences over individual goods.

4.2 Pareto-efficiency

We now turn to Pareto-efficiency, which is the first main property we consider.

Proposition 3. *If $n \leq 3$ or all agents have the same capacity, then the cTTC rule is Pareto-efficient. If $n \geq 4$, then there are capacity configurations for which the cTTC rule is not Pareto-efficient. For any capacity configuration, the cTTC rule is ig-Pareto-efficient.*

Proof. We prove the result in four steps.

STEP 1. If $n \leq 3$, then the cTTC rule is Pareto-efficient.

If $n = 2$ the result follows trivially. Let $n = 3$ and let $N = \{a, b, c\}$. Assume, without loss of generality, that $q_a \leq q_b \leq q_c$. Let $P \in \mathcal{P}$. If $\tau(P)$ can also be obtained by a serial dictatorship following some ordering of the agents (where each agent in the ordering sequentially picks the bundle she prefers most among the bundles that can be formed using all still available goods), then it is Pareto-efficient at P. So, it is sufficient to show that at *any* preference profile, the cTTC outcome can also be obtained by *some* serial dictatorship (not necessarily the same serial dictatorship for all preference profiles).

Fix a preference profile. Suppose in the first round of the cTTC algorithm there is a top trading cycle that consists of a single agent. Then, this agent, say (without loss of generality) agent a, gets her favorite bundle. Then, since the other two agents do not get good a, it can be easily verified that the serial dictatorship based on ordering (a, b, c) or (a, c, b) gives the cTTC outcome.

Now suppose that in the first round of the cTTC algorithm no top trading cycle consists of a single agent. Then, there is a unique top trading cycle in the first round of the cTTC algorithm. We consider the following three cases.

Case 1: The top trading cycle in the first round consists of either a, b, and c, or of a and b.
Then, in the first round they trade q_a units and a gets her most preferred bundle. Suppose $q_a = q_b$. Then b also gets her most preferred bundle and the serial dictatorship with ordering (a, b, c) gives the same circulation as cTTC. Now suppose $q_a < q_b$. If both $c \succ_b b$ and $b \succ_c c$, then b and c trade $q_b - q_a$ units in the second round of cTTC and the serial dictatorship with ordering (a, b, c) gives the same circulation as cTTC. Otherwise b and c do not trade in the second round of cTTC but still b or c gets her most preferred bundle.[16] Thus, the serial dictatorship with (a, b, c) or (a, c, b) gives the same circulation as cTTC.

Case 2: The top trading cycle in the first round consists of a and c.
Then, in the first round they trade q_a units and a gets her most preferred bundle. Suppose $q_a = q_c$. Then c also gets her most preferred bundle and the serial dictatorship with ordering (a, c, b) gives the same circulation as cTTC. Now suppose $q_a < q_c$. If both $c \succ_b b$ and $b \succ_c c$, then b and c trade $\min\{q_b, q_c - q_a\}$ units in the second round of cTTC and the serial dictatorship with ordering (a, c, b) gives the same circulation as cTTC. Otherwise b and c do not trade in the second round of cTTC but still b or c gets her most preferred bundle.[17] Thus, the serial dictatorship with (a, b, c) or (a, c, b) gives the same circulation as cTTC.

Case 3: The top trading cycle in the first round consists of b and c.
Then, in the first round they trade q_b units and b gets her most preferred bundle. Suppose $q_b = q_c$. Then c also gets her most preferred bundle and the serial dictatorship with ordering (b, c, a) gives the same circulation as cTTC. Now suppose $q_b < q_c$. If both $c \succ_a a$ and $a \succ_c c$, then a and c trade $\min\{q_a, q_c - q_b\}$ units in the second round of cTTC and the serial dictatorship with ordering (b, c, a) gives the same circulation as cTTC. Otherwise a and c do not trade in the second round of cTTC but still a or c gets her most preferred bundle.[18] Thus, the serial dictatorship with (b, a, c) or (b, c, a) gives the same circulation as cTTC.

STEP 2. If all agents have the same capacity, then the cTTC rule is Pareto-efficient.

Suppose the agents have the same capacity, say q^*. Let $P \in \mathcal{P}$. Suppose that $x = \tau(P)$ is not Pareto-efficient at P. Let x' be a circulation that Pareto-dominates x. Note that when cTTC is applied to P each agent trades in only one round and receives exactly q^* units of some good (or receives the null bundle). Consider the earliest round of cTTC, say r, in which some agent i receives a bundle $x_i \neq x'_i$. Let j be the good of which i receives q^* units at x. Since preferences are responsive and i prefers x'_i to x_i there is a good l such that $l \succ_i j$ and $x'_{il} > 0$. By definition of cTTC, all q^* units of l were assigned to some agent k in some round $1, \ldots, r-1$ of cTTC (or

[16] More precisely, if agent $i \in \{b, c\}$ finds good $j \in \{b, c\}\backslash\{i\}$ unacceptable, then i gets her most preferred bundle.
[17] More precisely, if agent $i \in \{b, c\}$ finds good $j \in \{b, c\}\backslash\{i\}$ unacceptable, then i gets her most preferred bundle.
[18] More precisely, if agent $i \in \{a, c\}$ finds good $j \in \{a, c\}\backslash\{i\}$ unacceptable, then i gets her most preferred bundle.

$k = l$ and agent k left the market with all q^* units of her own good). But then $x_k \neq x'_k$. This contradicts the fact that round r is the earliest round in which some agent receives a bundle that differs from the one she receives at x'. Hence, $x = \tau(P)$ is Pareto-efficient at P.

STEP 3. Let $n \geq 4$. Then, there are capacity configurations for which the cTTC rule is not Pareto-efficient.

Without loss of generality assume that $N = \{a, b, c, d\}$. (Otherwise, let any additional agent find all goods (but her own good) unacceptable.) Suppose $q_a < q_c \leq q_d \leq q_b$. Consider the class of preferences $\mathcal{P}' \subset \mathcal{P}$ such that the preferences over individual goods are as follows:

$\succ_a$: $c \succ_a a$
$\succ_b$: $d \succ_b b$
$\succ_c$: $a \succ_c b \succ_c c$
$\succ_d$: $a \succ_d b \succ_d c \succ_d d$

Let $P \in \mathcal{P}'$. Let $x = \tau(P)$. In the first round of the cTTC algorithm, there is a unique top trading cycle which consists of a and c, which yields $x_{ac} = x_{ca} = q_a$. In the second round, there is a unique top trading cycle which consists of b and d, which yields $x_{bd} = x_{db} = q_d$. In the final round(s), the remaining goods are kept by their owners.

Let $x' \in X$ be the circulation defined by $x'_{ac} = x_{ac} = q_a$ and $x'_{bd} = x_{bd} = q_d$, but $x'_{cb} = q_c$, $x'_{da} = q_a$, $x'_{db} = q_d - q_c$, and $x'_{dc} = q_c - q_a$. Note that a and b receive the same bundles at x and x'. Since $q_c > q_a$, there exist two numbers $u_c(a)$ and $u_c(b)$ such that $0 < u_c(b) < u_c(a)$ and $q_c u_c(b) > q_a u_c(a)$. Let $P \in \mathcal{P}'$ be such that d has lexicographic preferences and c has additive preferences with $q_c u_c(b) > q_a u_c(a)$. Then c prefers x'_c to x_c and d prefers x'_d to x_d. Hence, x' Pareto-dominates x at P. Hence, the cTTC rule does not yield a Pareto-efficient circulation.

STEP 4. For any capacity configuration, the cTTC rule is ig-Pareto-efficient.

Suppose the cTTC rule does not yield a Pareto-efficient circulation for some market (N, q, P), where P is a profile of lexicographic preferences. Let $x = \tau(P)$. Then, there is a circulation x' that Pareto-dominates x. Consider the one-unit-per-cycle version of the cTTC algorithm for x. Let i be an agent in a top trading cycle in the first round. Among the goods that i receives at x', let j be the good she prefers most. Since agent i weakly prefers x'_i to x_i and since she has lexicographic preferences, j is the good to which agent i points in the first round of the cTTC algorithm for x. Therefore, j is also in a top trading cycle in the first round (for x). By repeating the same arguments, it follows that each top trading cycle in the first round (for x) is part of circulation x'. More formally, let i be any agent in a top trading cycle in the first round for x, and let j be the good of which she then (i.e., in the first round) receives one unit. Then, $x'_{ij} > 0$.

Consider the market that is obtained by 1) reducing with one unit the capacity of each agent that was in a top trading cycle in the first round, and 2) removing the agents that as a consequence have capacity 0. Let the remaining agents have the same lexicographic preferences as before (if necessary, update preferences by omitting goods/agents that are no longer present). Let y and y' be the circulations obtained from x and x', respectively, by 1) setting $y = x$ and $y' = x'$, 2) updating $y_{ij} = y_{ij} - 1$ and $y'_{ij} = y'_{ij} - 1$ for any agent i in a top trading cycle of the first round, j being the good agent i points to, and[19] 3) removing any y_{kl} and y'_{kl} if good l is no longer present,

[19] Since x and x' are both feasible circulations, it follows that for each good l it holds that for all k, $y_{kl} = 0$ if and only if for all k, $y'_{kl} = 0$.

i.e., if $y_{kl} = y'_{kl} = 0$ for all k.

Since x' Pareto-dominates x and preferences are lexicographic, y' Pareto-dominates y with respect to the updated lexicographic preferences. Since y is also the circulation obtained from the cTTC algorithm applied to the updated lexicographic preferences, we can use the same arguments as before to subtract again the first top trading cycles from y and y', update the market, etc. By repeating this procedure a finite number of times it follows that $x = x'$, which contradicts the fact that x' Pareto-dominates x. □

Corollary 1. *If the capacity configuration is irreducible, then the cTTC rule is Pareto-efficient.*

4.3 Strategy-proofness and other incentive properties

When all agents have unit capacity the cTTC rule boils down to the classical top trading cycles mechanism and hence the rule is strategy-proof. In fact, by considering the one-unit-per-cycle version of the cTTC algorithm, we immediately obtain the following more general result.

Proposition 4. *For any capacity configuration, the cTTC rule cannot be manipulated by any agent with unit capacity, i.e., any agent i with $q_i = 1$.*

We will show that there are reducible capacity configurations for which Proposition 4 cannot be generalized to agents with larger capacity (Proposition 6). However, our next result shows that it can be generalized for all irreducible capacity configurations.

Corollary 2. *If the capacity configuration is irreducible, then the cTTC rule is strategy-proof.*

We defer the proof of Corollary 2 to Section 5, as it is a direct corollary to Propositions 8 and 9.

Next, we study if and when the cTTC rule satisfies other incentive properties.

Proposition 5. *For any capacity configuration, the cTTC rule is hiding-proof and truncation-proof.*

Proof. We only show that τ is hiding-proof; using similar arguments one easily proves that τ is also truncation-proof. Suppose $q'_i < q_i$. Then, the (one-unit-per-cycle) cTTC algorithm applied to $(P, (q_{-i}, q'_i))$ proceeds precisely in the same way as the (one-unit-per-cycle) cTTC algorithm applied to (P, q) until agent i exhausts her capacity. Therefore the bundle $f_i(P, (q_{-i}, q'_i))$ contains a weakly smaller number of units of each good compared to $f_i(P, q)$. By responsiveness of P_i, $f_i(P, q) \, R_i \, f_i(P, (q_{-i}, q'_i))$. This shows that τ is hiding-proof. □

The formulation of the following result is very similar to that of Proposition 3.

Proposition 6. *If $n \leq 3$ or all agents have the same capacity, then the cTTC rule is dropping-proof. If $n \geq 4$, then there are capacity configurations for which the cTTC rule is not dropping-proof. For any capacity configuration, the cTTC rule is ig-dropping-proof.*

Proof. We prove the result in four steps.

STEP 1. If $n \leq 3$, then the cTTC rule is dropping-proof.

If $n = 2$ the result follows trivially. Let $n = 3$ and let $N = \{a, b, c\}$. Fix a preference profile. We show that a cannot manipulate the cTTC rule by dropping. We can assume that each agent finds at least one other good acceptable (otherwise we are essentially back in the case $n = 2$). Then, there is a unique top trading cycle in the first round of cTTC.

Suppose a is not involved in the first top trading cycle. Suppose a plays a dropping strategy. Then the first top trading cycle remains unchanged. But then, as in the second round we are left with at most two agents, the result follows.

Now suppose a is involved in the first top trading cycle with only one other agent, say with b. Then a or b gets saturated in the first round, i.e., fills her capacity, and leaves the market. If agent a gets saturated then she received q_a units of her favorite good, so she cannot be better off in any other circulation. If it is agent b (but not a) who becomes saturated in the first round, then a receives q_b units of her favorite good and she is left with c in the second round of cTTC, where she gets $\min\{q_a - q_b, q_c\}$ units of c, provided that c finds a acceptable (and a finds c acceptable; otherwise, a cannot be better off by dropping b). It is not difficult to see that a cannot benefit from dropping b or c, since if she drops c then she just receives q_b units of b as a final bundle; if she drops b then only when a and c are mutually acceptable she gets ($\min\{q_a, q_c\}$ units of) good c, which by responsiveness, is worse than her final bundle when she tells the truth.

Finally, suppose all three agents are involved in the first top trading cycle, say in order (a, b, c). Suppose that a is saturated in the first round, i.e., $q_a \leq \min\{q_b, q_c\}$. Then, she gets as many units of her favorite good as her capacity. Hence, she cannot benefit from any dropping strategy. Suppose now that b is saturated in the first round, i.e., $q_b \leq \min\{q_a, q_c\}$. In this case, the final allocation of a is q_b of b and if a finds c acceptable, also $\min\{q_a, q_c\} - q_b$ units of good c. If a drops c then a gets q_b of b only. If a drops b then a gets $\min\{q_a, q_c\}$ units of good c only (if a finds c acceptable), which by responsiveness is worse than her bundle under truth-telling. Finally, suppose that c is saturated in the first round, i.e., $q_c \leq \min\{q_a, q_b\}$. If b finds a unacceptable, then a cannot manipulate by dropping for sure. Suppose now that b finds a acceptable. Then, since a gets the bundle that consists of $\min\{q_a, q_b\}$ units of good b, dropping c does not yield a more preferred bundle. By dropping b, a would get the bundle that consists q_c units of c (assuming c is acceptable for a), which by responsiveness does not yield a more preferred bundle.

STEP 2. If all agents have the same capacity, then the cTTC rule is dropping-proof.

If all agents have the same capacity, then the capacity configuration is irreducible. Then, by Corollary 2, the cTTC rule is strategy-proof and hence dropping-proof.

STEP 3. Let $n \geq 4$. Then, there are capacity configurations for which the cTTC rule is not dropping-proof.

Without loss of generality assume that $N = \{a, b, c, d\}$. (Otherwise, let any additional agent find all goods (but her own good) unacceptable.) Let $q_a = q_b = q_c = 1$ and $q_d = 2$. Consider the class of preferences $\mathcal{P}' \subset \mathcal{P}$ such that the preferences over individual goods are as follows:

$\succ_a$: $b \succ_a a$
$\succ_b$: $c \succ_b d \succ_b b$
$\succ_c$: $d \succ_c c$
$\succ_d$: $a \succ_d c \succ_d b \succ_d d$

Let $P \in \mathcal{P}'$. Let $x = f(P)$. In the first round of the cTTC algorithm there is a top trading cycle that consists of a, b, c, and d, which yields $x_{ab} = x_{bc} = x_{cd} = x_{da} = 1$. In the second (and final) round, d keeps the remaining unit of her good.

Now let agent d report preferences such that

$$\succ_d\colon\ c \succ'_d b \succ'_d d,$$

which is a dropping of the original preferences of d. Then, in the first round of the cTTC algorithm there is a (unique) top trading cycle that consists of c and d, which yields $x'_{cd} = x'_{dc} = 1$. In the second round there is a (unique) top trading cycle that consists of b and d, which yields $x'_{bd} = x'_{db} = 1$. In the third (and final) round, a keeps her good. Let $P \in \mathcal{P}'$ be such that d has additive preferences with $u_d(b) + u_d(c) > u_d(a)$. Then d strictly prefers x'_d to x_d. Hence, the cTTC rule is not dropping-proof.

Step 4. For any capacity configuration, the cTTC rule is ig-dropping-proof.

Suppose the cTTC rule is not ig-dropping-proof. Then there is a market (N, q, P) such that for some $i \in N$, P_i is lexicographic and there is P'_i with $\tau_i(P'_i, P_{-i})\, P_i\, \tau_i(P)$ and $\succ'_i$ is a dropping of $\succ_i$ (these being the preferences over individual goods induced by P'_i and P_i, respectively). Let $x = \tau(P)$ and $x' = \tau(P'_i, P_{-i})$.

If $\succ'_i$ is obtained from $\succ_i$ by only dropping one or several goods k with $x_{ik} = 0$, then $x = x'$, which contradicts $x'_i P_i x_i$. So, $\succ'_i$ is obtained from $\succ_i$ by dropping some good k^* with $x_{ik^*} > 0$ (and possibly dropping one or several goods k with $x_{ik} = 0$). Let k^{**} be the most preferred good with $x_{ik^{**}} > 0$ and that is dropped in $\succ'_i$. Then, for all goods l that are strictly preferred to k^{**}, $x_{il} = x'_{il}$. Then, since $x_{ik^{**}} > 0 = x'_{ik^{**}}$ and preferences P_i are lexicographic, $x_i P_i x'_i$, which contradicts $x'_i P_i x_i$. □

4.4 A characterization of cTTC and an (almost) impossibility result

We are now ready to characterize the cTTC rule.

Theorem 1. *For any capacity configuration, the cTTC rule is the* unique *circulation rule that is individually rational, ig-Pareto-efficient, and ig-dropping-proof.*

Proof. The cTTC rule is individually rational by construction, ig-Pareto-efficient by Proposition 3, and ig-dropping-proof by Proposition 6.

It remains to prove that the cTTC rule is the unique rule that satisfies the three properties. Let f be a circulation rule that is individually rational, ig-Pareto-efficient, and ig-dropping-proof. We will show that for any profile of *lexicographic* preferences P, $f(P) = \tau(P)$. Since f and τ are individual-good-preference based, then for *any* profile of preferences P, $f(P) = \tau(P)$.

We will show by induction that for any integer $l \in \mathbb{N}_{++}$, if P is a profile of lexicographic preferences such that cTTC requires $k \geq l$ rounds, then the circulation induced by the union of the top trading cycles (and their flows) in rounds $1, \ldots, l$ are contained in circulation $f(P)$.

We first prove the statement for $l = 1$. Let P be a profile of lexicographic preferences and let $x = f(P)$. Let $\succ \equiv \succ^P$ be the induced preferences over individual goods. Let $Z \subseteq N$ be a top trading cycle coalition[20] in the first round of the cTTC algorithm applied to P. Let q be

[20]More precisely, a top trading cycle coalition is a set of agents that constitute some cycle in the cTTC algorithm.

the minimum capacity among the members of Z. Suppose that there exists *some* agent $a \in Z$ such that for a's most preferred good, say good b, we have $x_{ab} < q$. Let P'_a be any lexicographic preferences such that the induced preferences over individual goods are given by dropping $\succ'_a$: b, a. Let $x' = f(P'_a, P_{-a})$. Individual rationality of f implies that $x'_{a\bullet} = x'_{ab}$. Since $\succ'_a$ is a dropping of $\succ_a$, it follows from ig-dropping-proofness of f that a does not strictly prefer x'_a to x_a. Hence, $x'_{ab} \leq x_{ab}$ and thus $x'_{ab} < q$. Then, from the balancedness of circulation x' at a it follows that for agent $c \in Z$ that points in the top trading cycle to a, $x'_{ca} \leq x'_{a\bullet} = x'_{ab} < q$.

Note that by construction P' is a profile of lexicographic preferences such that $Z \subseteq N$ is also a top trading cycle coalition in the first round of the cTTC algorithm applied to P' and, as we have just shown, $x'_{ca}, x'_{ab} < q$. Hence, we can repeat the same arguments to the next predecessors in the top trading cycle coalition Z until we are back in a. At that moment we have constructed a profile of lexicographic preferences, say P^*, in which each agent of Z only finds acceptable the good she was pointing to in the top trading cycle associated with Z (and her own good). By construction, at $f(P^*)$ the agents in Z (only) circulate their goods suboptimally, since the flow is strictly below the maximum q, i.e., all agents in Z are unfilled at $f(P^*)$. But then $f(P^*)$ is not Pareto-efficient at P^*: all agents in Z would be strictly better off if the flow in the associated cycle would be increased. This shows that for any profile of lexicographic preferences P, if Z is a top trading cycles coalition with flow q in the first round of cTTC applied to P, then *each* agent $a \in Z$, who points to $b \in Z$ say, obtains at $f(P)$ a quantity of good b that is weakly larger than q. So, for any profile of lexicographic preferences P, the circulation induced by the union of top trading cycles (and their flows) that are generated in round 1 of cTTC applied to P are contained in circulation $f(P)$.

We now prove the statement for $l = 2$. Let P be a profile of lexicographic preferences such that cTTC requires $k \geq l$ rounds. Let $x = f(P)$. The result for $l = 1$ implies that if we take x and subtract the circulation induced by the union of the top trading cycles from round 1 we obtain a (feasible) circulation y. Let $Z \subseteq N$ be a top trading cycle coalition in the second round of the cTTC algorithm applied to P. Let q be the minimum remaining capacity among the members of Z at the beginning of the second round. (Some or all members of Z can have been members of top trading cycles in the first round, but Z itself cannot have been a top trading cycle coalition in the first round.)

Suppose that there exists *some* agent $a \in Z$ such that for the good to which a now points, say good b, we have $y_{ab} < q$. Let P'_a be any lexicographic preferences such that the induced preferences over individual goods are given by the dropping $\succ'_a$: b, a if a did not receive any good in round 1 of cTTC or received some units of good b in round 1; and the dropping $\succ'_a$: d, b, a if a received some units of good $d \neq b, a$ in round 1. Let $x' = f(P')$ where $P' = (P'_a, P_{-a})$. Note that by definition of cTTC the top trading cycles in rounds 1 and 2 of cTTC are the same at P and P'. The result for $l = 1$ implies that if we take x' and subtract the circulation induced by the union of the top trading cycles from round 1 we obtain a (feasible) circulation y'. Next we show that $y'_{a\bullet} = y'_{ab}$ and $y'_{ab} < q$.

Suppose that in round 1 of cTTC at P and P' agent a did not receive any good (zero amount at both P and P') or received some units of good b (the same amount at P and P'). Then, by individual rationality of f, $y'_{a\bullet} = y'_{ab}$, and so by ig-dropping-proofness of f, $y'_{ab} \leq y_{ab} < q$.

Suppose now that in round 1 of cTTC at P and P' agent a received some units of good $d \neq b, a$. Since a received some units of good $d \neq b$ in the first round of cTTC and some units of good b in the second round of cTTC, it follows that d was filled in the first round of cTTC. Then, since the top trading cycles from round 1 are the same for x and x', it follows that $x_{ad} = x'_{ad}$. Then, by individual rationality and ig-dropping-proofness of f, $x'_{ab} \leq x_{ab}$. Hence, $y'_{ab} = x'_{ab} \leq x_{ab} = y_{ab} < q$. From the capacity constraint for d, $y'_{ad} = 0$. Then, from the individual rationality of f, $y'_{a\bullet} = y'_{ab}$.

Let $c \in Z$ be the agent that points to a in the top trading cycle Z (in the second round, which is common to P and P'). Then, from the balancedness of circulation y' at a, $y'_{a\bullet} = y'_{ab}$, and $y'_{ab} < q$ we have that $y'_{ca} \leq y'_{a\bullet} = y'_{ab} < q$.

Note that by construction P' is a profile of lexicographic preferences such that $Z \subseteq N$ is also a top trading cycle coalition in the second round of the cTTC algorithm applied to P' and, as we have just shown, $y'_{ca}, y'_{ab} < q$. Hence, we can repeat the same arguments to the next predecessors in the top trading cycle coalition Z until we are back in a. Let P^* be the profile of lexicographic preferences that has been sequentially constructed until this moment. Let y^* be the circulation obtained from $f(P^*)$ by subtracting the circulation induced by the union of the top trading cycles from round 1. By construction, $Z \subseteq N$ is also a top trading cycle coalition in the second round of the cTTC algorithm applied to P^*, and for each agent $z \in Z$, $y^*_{zw} < q$ where w is the good that agent z points to in the top trading cycle defined by Z. But then $f(P^*)$ is not Pareto-efficient at P^*: all agents in Z would be strictly better off if the flow in the associated cycle would be increased. This, together with the statement for $l = 1$, shows that for any profile of lexicographic preferences P, the circulation induced by the union of top trading cycles (and their flows) that are generated in rounds 1 and 2 of cTTC applied to P are contained in circulation $f(P)$.

The statement for $l > 2$ can be proved with similar arguments. □

Remark 3. The properties in Theorem 1 are logically independent. First, multi-serial-IR rules in Biró et al. [4] are individually rational and Pareto-efficient, and thus also ig-Pareto-efficient. Second, multi-serial-rules in Biró et al. [4] are Pareto-efficient and strategy-proof (see footnote 14 on dispensing with (resp-3)). Finally, all Segmented Trading Cycle rules of Section 5 are individually rational and strategy-proof, and thus also ig-dropping-proof. ◇

We now state and prove our (almost) impossibility theorem.

Theorem 2. *There is a circulation rule that is individually rational, Pareto-efficient, and strategy-proof if and only if the capacity configuration is irreducible. For irreducible capacity configurations, the cTTC rule is the* unique *circulation rule that satisfies all three properties.*

Proof. The "only if" part of the first statement follows from Proposition 2. The "if" part of the first statement follows from Corollaries 1 and 2.

Regarding the second statement, Corollaries 1 and 2 show that the cTTC rule satisfies the three properties. The uniqueness follows from Theorem 1. □

In light of the results that we present in the next section, see also the revisited version of Theorem 2 in subsection 5.3.

5 Segmented Trading Cycle rules

The cTTC rule generalizes the classical Top Trading Cycle (TTC) rule and allows for exchange in the overall market. While not fully efficient, it is ig-Pareto-efficient, but it lacks strategy-proofness. We now consider a different approach in order to find strategy-proof rules, following Pápai [19]:

we first distribute the agents' endowments over a number of smaller markets (the so-called market segments), then apply the classical Top Trading Cycle algorithm within each market segment, and finally aggregate the circulations. This class of "Segmented Trading Cycle" (STC) rules are strategy-proof but in general do not satisfy ig-Pareto-efficiency, unlike cTTC. We also identify one notable member of this class of rules, which we refer to as the SSTC rule, as the main alternative to the cTTC rule, and characterize it in order to shed some light on the tradeoffs between the cTTC rule and the SSTC rule. Both the cTTC and SSTC rule boil down to the classical TTC rule in the original Shapley-Scarf market, and we will show that more generally in multi-unit Shapley-Scarf markets with an irreducible capacity configuration the two rules coincide and are characterized by individual rationality, Pareto-efficiency, and strategy-proofness.

5.1 Market Segmentation and Uniqueness

Let $q_{\max} \equiv \max_{i \in N} q_i$. We create $q_{\max}$ **market segments** such that each agent has at most one unit of her good in each market segment. For all $t = 1, \ldots, q_{\max}$, let W_t denote the set of goods in market segment t. Then, for all $i \in N$ and all $t = 1, \ldots, q_{\max}$, $|W_t \bigcap e_i| \leq 1$ and $\bigcup_{t=1}^{q_{\max}} W_t = \bigcup_{i \in N} e_i$. For each market segment $t = 1, \ldots, q_{\max}$ and each preference profile $P \in \mathcal{P}$, let $D_t = (N_t, A_t)$ be the directed graph such that $N_t = \{i \in N : |W_t \bigcap e_i| = 1\}$, and for all $i, j \in N_t$, $(i, j) \in A_t$ if and only if $(i, j) \in A(P)$.

A **Segmented Trading Cycle (STC) rule** is obtained by applying the following algorithm to each preference profile. In each market segment t, carry out the TTC algorithm, given the restrictions imposed by D_t. Thus, for each given market segment t, in the first step of the algorithm let each agent point to her most preferred potential trading partner in D_t. If there is no such trading partner then the agent points to herself. We get at least one top trading cycle. Let the agents in each top trading cycle trade their goods (one unit for each agent). If an agent points to herself then she receives her own endowment in this market segment. We remove each agent in top trading cycles from the market (i.e., we remove the corresponding nodes from the graph D_t), and repeat the same process in the remaining market until all agents are removed. The aggregation of the circulations obtained in all $q_{\max}$ market segments yields the final circulation at the preference profile in question.

As we will see, a Segmented Trading Cycle rule can only achieve full Pareto-efficiency for irreducible capacity configurations. In order to avoid further and unnecessary efficiency losses, we require "compactness" of the market segments by restricting the number of segments to at most $q_{\max}$, which is the minimum number of possible market segments. Note that, in general, not all Pareto-efficient circulations can be reached under this restriction, but this is already true for the original Shapley-Scarf housing markets.[21] However, subject to the trade restrictions imposed by the trading segments, a Segmented Trading Cycle rule is as efficient as possible. That is, there is no Pareto-improving trade that can be carried out within any market segment, since the classical TTC rule is Pareto-efficient. We can think of this as an "efficiency-on-the-range" property, which

[21] There are capacity configurations and preference profiles such that some Pareto-efficient circulation can only be reached by allowing a market segmentation with strictly more than $q_{\max}$ segments. For instance, consider the Shapley-Scarf market given by $N = \{a, b, c\}$, $q = (1, 1, 1)$, and preferences $b \succ_a c \succ_a a$, $a \succ_b b$, and $a \succ_c c$. If there is only $q_{\max} = 1$ market segment, then the unique Segmented Trading Cycle rule is equivalent to the classical TTC rule and yields the circulation where a and b swap their single units. The circulation where a and c swap is also Pareto-efficient, but it requires two market segments: $\{a, c\}$ and $\{b\}$.

requires that the circulation outcome at any preference profile is not Pareto-dominated by another circulation that is obtained by the rule at a different preference profile. Since "efficiency-on-the-range" becomes more demanding when more trade can take place, this property of Segmented Trading Cycle rules explains why keeping the number of market segments to a minimum helps to improve the efficiency of these rules.

Observe that a Segmented Trading Cycle rule depends on how the initial endowments are partitioned into market segments, and consequently there is a class of Segmented Trading Cycle rules in which each member is determined by $W = (W_t)_{t=1,\ldots,q_{\max}}$.

We will refer to the following set of capacity configurations as **quasi-irreducible**: the capacity configuration is either irreducible or satisfies

(rd2*) $n \geq 4$, $N = \{i_1, i_2, \ldots, i_n\}$, and $q_{i_1} < q_{i_2} = q_{i_3} = \ldots = q_{i_n}$.

Thus, quasi-irreducible capacity configurations include all irreducible capacity configurations and a subclass of (rd2) reducible capacity configurations,[22] as defined in Section 3.

Proposition 7. *There is a unique Segmented Trading Cycle rule if and only if the capacity configuration is quasi-irreducible.*

Proof. This statement is obvious for $n = 2$, and in the case where all agents have the same capacities. Thus, we only have to consider the case where $q_{i_1} < q_{i_2} = q_{i_3} = \ldots = q_{i_n}$ with $N = \{i_1, i_2, \ldots, i_n\}$ and $n = |N| \geq 3$. In this case there is a unique Segmented Trading Cycle rule as the market segmentation is unique up to isomorphism (i.e., permutations of complete market segments).

Suppose that the capacity configuration does not fall under any of the previous cases. Then there are at least two agents whose capacities are strictly less than the maximum capacity $q_{\max}$ among all agents, which implies that the market segmentation is not unique (not even up to isomorphism). One easily verifies that in this case there are at least two different Segmented Trading Cycle rules. □

In the next proposition we show that the uniqueness of the Segmented Trading Cycle rule does not necessarily make it identical to the cTTC rule for all quasi-irreducible capacity configurations. In fact, the two rules only coincide for irreducible capacity configurations but not for (rd2*) configurations. This makes intuitive sense, since the irreducible capacity configurations are the exceptions in the impossibility theorem (Theorem 2).

Proposition 8. *If the capacity configuration is irreducible, then the unique Segmented Trading Cycle rule is identical to the cTTC rule. If the capacity configuration satisfies (rd2*), then the unique Segmented Trading Cycle Frule is* not *identical to the cTTC rule.*

Proof. The first statement is obvious for $n = 2$ and in the case where all agents have the same capacities. Thus, we only have to consider the case where $q_a < q_b = q_c$ with $N = \{a, b, c\}$ and $|N| = 3$. Here, the unique segmentation is to have q_a segments with all agents and $q_b - q_a$ segments with $\{b, c\}$ only. If the first round of the cTTC involves a cycle with all agents then they trade in the same way in cTTC as in the 3-agent segments of the Segmented Trading Cycle rule, and later

[22]Notice that (rd2) only requires that there are at least 3 (i.e., not necessary $n - 1$) different agents that have a common capacity that is strictly larger than the capacity of some other agent.

the agents $\{b, c\}$ trade in the cTTC in the same way as in the 2-agent segments of the Segmented Trading Cycle rule. Similar arguments apply if in the first round of the cTTC there is no top trading cycle that involves all agents. Hence, the unique Segmented Trading Cycle rule is identical to the cTTC rule.

To prove the second statement, let $N = \{a, b, c, d\}$ with $|N| = 4$ and $q_a < q_b = q_c = q_d$ (for $|N| > 4$, let any additional agent find all goods (but her own good) unacceptable). Consider any preference profile P such that the preferences over individual goods $\succ = \succ^P$ are as follows:

$\succ_a$: $b \succ_a a$
$\succ_b$: $c \succ_b d \succ_b b$
$\succ_c$: $a \succ_c d \succ_c c$
$\succ_d$: $c \succ_d b \succ_d d$

One can easily verify that the unique STC rule yields at P the circulation given by $x_{ab} = x_{bc} = x_{ca} = q_a$ and $x_{cd} = x_{dc} = q_c - q_a$. However, the cTTC rule yields at P the circulation given by $x_{ab} = x_{bc} = x_{ca} = q_a$, $x_{cd} = x_{dc} = q_c - q_a$, and $x_{bd} = x_{db} = \min\{q_a, q_b - q_a\} > 0$. Hence, the unique Segmented Trading Cycle rule is *not* identical to the cTTC rule. □

5.2 Properties of Segmented Trading Cycle rules

Definition 7. A circulation rule f is **nonbossy**[23] if for all $i \in N$, $P_i, P'_i \in \mathcal{P}_i$, and $P_{-i} \in \times_{j \neq i} \mathcal{P}_j$, $f_i(P'_i, P_{-i}) = f_i(P)$ implies $f(P'_i, P_{-i}) = f(P)$. ◇

Since the classical TTC rule is strategy-proof and nonbossy[24] for each market segment and since preferences are responsive, we obtain the following result.

Proposition 9. *For any capacity configuration, Segmented Trading Cycle rules are strategy-proof and nonbossy.*

Proof. Since preferences are responsive, it is easy to verify that Segmented Trading Cycle rules are strategy-proof, given that the classical TTC rule is strategy-proof. In order to verify nonbossiness, we will first show that if an agent is assigned the same bundle of goods when the agent reports different preferences then each individual good in the bundle has to come from the same market segment. Suppose, to the contrary, that agent i is assigned the same set of goods but some of the assigned goods were obtained in different market segments depending on agent i's reported preferences. Specifically, assume without loss of generality that in market segment t_v agent i gets good a_v when reporting $\succ_i$, and good a_{v+1} (modulo k) when reporting $\succ'_i$, that is, we have a simple cycle permutation of the assigned goods (which may include i's own good) considering the market segments that they come from when we compare two preference profiles that only differ in i's report. Then, by the strategy-proofness of the classical TTC rule, given that both goods a_v and a_{v+1} can be assigned to agent i in segment t_v depending on i's report, we have that $a_1 \succ_i a_2 \succ_i \ldots \succ_i a_k \succ_i a_1$. This is a contradiction since $\succ_i$ is transitive. Therefore, if $f_i(P) = f_i(P'_i, P_{-i})$ then i is assigned the same good in each segment at both profile P and (P'_i, P_{-i}), and thus the nonbossiness of the classical TTC rule implies that the Segmented Trading Cycle rules are nonbossy. □

[23] Satterthwaite and Sonnenschein [26] were the first to introduce and study nonbossiness. We refer to Thomson [33] for a comprehensive overview and discussion on nonbossiness.
[24] See Ma [14], Svensson [32] and Miyagawa [16].

As a corollary to Propositions 8 and 9 we obtain Corollary 2: the cTTC rule is strategy-proof (and nonbossy) for irreducible capacity configurations.

In view of the result above, we will consider more demanding notions of strategy-proofness. For any $P \in \mathcal{P}$ and any $S \subseteq N$, denote $P_S \equiv (P_i)_{i \in S}$ and $P_{-S} \equiv (P_i)_{i \in N \setminus S}$.

Definition 8. A **coalition** $S \subseteq N$ can **manipulate** circulation rule f at $P \in \mathcal{P}$ via deviation $P'_S \in \times_{i \in S} \mathcal{P}_i$ if, for all $i \in S$, $f_i(P'_S, P_{-S}) R_i f_i(P)$ and, for some $j \in S$, $f_j(P'_S, P_{-S}) P_j f_j(P)$. A circulation rule f is (necessarily) **group-strategy-proof** if no coalition can manipulate f at any $P \in \mathcal{P}$. A circulation rule f is **ig-group-strategy-proof** if no coalition can manipulate f at any profile of lexicographic preferences $P \in \mathcal{P}^L$.[25] $\diamond$

When allocating a single unit of an indivisible good to each agent (e.g., in the original Shapley-Scarf housing markets), strategy-proofness and nonbossiness are equivalent to group-strategy-proofness (see Lemma 1 in Pápai [17]). In our context where agents can receive multiple types and units of indivisible goods, strategy-proofness and nonbossiness do not imply group-strategy-proofness. We will show that in fact a weaker property, which we call *self-enforcing group-strategy-proofness,* is equivalent to the two axioms. Before establishing this result, we demonstrate first that there exist capacity configurations for which all Segmented Trading Cycle rules violate group-strategy-proofness. Specifically, we show that even for (ird2) irreducible capacity configurations the unique Segmented Trading Cycle rule is not group-strategy-proof. On the other hand, for (ird1) and (ird3) capacity configurations, that is, when $n = 2$ and when all agents have identical capacities, respectively, it is clear that the unique Segmented Trading Cycle rule (or, equivalently, the cTTC rule) is group-strategy-proof.

Proposition 10. *Even for some irreducible capacity configurations, the unique Segmented Trading Cycle rule (or, equivalently, the cTTC rule) is not group-strategy-proof.*

Proof. To prove the statement, consider the market (N, q, P) where $N = \{a, b, c\}$ with $q_a < q_b = q_c$ and preferences P are such that the underlying preferences over individual goods are as follows:

$\succ_a$: $c \succ_a a$
$\succ_b$: $a \succ_b b$
$\succ_c$: $a \succ_c b \succ_c c$

Assume also that preferences are additive such that $q_a u_b(a) + (q_b - q_a) u_b(c) > 0$ and $q_c u_c(b) > q_a u_c(a)$.

Consider the unique Segmented Trading Cycle rule induced by q_a segments of $\{a, b, c\}$ and $q_b - q_a$ segments of $\{b, c\}$. Then, at profile P, agent b keeps the q_b units of her own good, while agent c receives q_a units of good a and keeps $q_c - q_a$ units of her own good. Suppose agents b and c now report preferences P'_b and P'_c such that the underlying preferences over individual goods are:

$\succ'_b$: $a \succ_b c \succ_b b$
$\succ'_c$: $b \succ_c c$

[25] Since circulation rules are individual-good-preference based, equivalent definitions of group-strategy-proofness and ig-group-strategy-proofness are obtained by additionally demanding that the deviation P'_S be a profile of lexicographic preferences.

Then, at profile $(P_a, P'_{\{b,c\}})$, the unique Segmented Trading Cycle rule gives q_a units of good a and $q_b - q_a$ units of good c to agent b, while it gives q_c units of good b to agent c. Given the assumptions on the preferences of agents b and c, the rule is not group-strategy-proof. □

Definition 9. A **coalition** $S \subseteq N$ can **manipulate** circulation rule f at $P \in \mathcal{P}$ **in a self-enforcing manner** if S can manipulate f at $P \in \mathcal{P}$ via deviation P'_S and, in addition, for all $i \in S$, $f_i(P'_S, P_{-S}) R_i f_i(P_i, P'_{S\setminus\{i\}}, P_{-S})$. If coalition S can manipulate f at $P \in \mathcal{P}$ via deviation P'_S then S is a **minimal manipulating coalition** at P via P'_S if there is no $\bar{S} \subsetneq S$ such that $\bar{S}$ can manipulate f at P via deviation $P'_{\bar{S}}$. A circulation rule f is **self-enforcing group-strategy-proof** if no minimal manipulating coalition can manipulate f in a self-enforcing manner at any $P \in \mathcal{P}$.[26] ◇

A manipulation by coalition S is self-enforcing if members of S have no incentive to revert back unilaterally to the true preferences at the manipulation profile (P'_S, P_{-S}). These manipulations are self-enforcing in the sense that no member of the manipulating coalition is strictly better off by reporting his true preferences when all other members of the coalition report their manipulation preferences. Self-enforcing group-strategy-proofness is a weaker version of group-strategy-proofness.[27] It does not rule out all manipulations by coalitions, but it rules out any self-enforcing manipulation by minimal manipulating coalitions.

Proposition 11. *For any capacity configuration, a circulation rule is strategy-proof and nonbossy if and only if it is self-enforcing group-strategy-proof.*

Proof. It is easy to see that self-enforcing group-strategy-proofness implies strategy-proofness. We will show that self-enforcing group-strategy-proofness also implies nonbossiness. Suppose that a circulation rule f is self-enforcing group-strategy-proof and bossy. Then there exist $P \in \mathcal{P}$, $i, j, \in N$, and $P'_i \in \mathcal{P}_i$ such that $f_i(P) = f_i(P'_i, P_{-i})$ and $f_j(P) \neq f_j(P'_i, P_{-i})$. Let $S = \{i, j\}$ and let $P'_j = P_j$. Then $(P'_S, P_{-S}) = (P'_i, P_{-i})$ and $f_i(P'_S, P_{-S}) R_i f_i(P)$ holds. Moreover, since $f_j(P'_S, P_{-S}) \neq f_j(P)$, we can assume without loss of generality that $f_j(P'_S, P_{-S}) P_j f_j(P)$. This proves that $S = \{i, j\}$ is a manipulating coalition at P via deviation (P'_i, P'_j). Since f is self-enforcing group-strategy-proof, it is also strategy-proof, which implies that $S = \{i, j\}$ is a minimal manipulating coalition at P via deviation (P'_i, P'_j). Finally, note that $f_i(P'_S, P_{-S}) R_i f_i(P_i, P'_j, P_{-\{i,j\}})$ is equivalent to $f_i(P'_S, P_{-S}) R_i f_i(P)$, and $f_j(P'_S, P_{-S}) R_j f_j(P'_i, P_j, P_{-\{i,j\}})$ is equivalent to $f_j(P'_S, P_{-S}) R_j f_j(P'_S, P_{-S})$, both of which are satisfied. Thus, $S = \{i, j\}$ is a minimal manipulating coalition that can manipulate in a self-enforcing manner at P, which contradicts our assumption that f is self-enforcing group-strategy-proof.

We will show next the converse statement: if f is strategy-proof and nonbossy then it is self-enforcing group-strategy-proof. Suppose that there exists a coalition $S \subseteq N$ that can manipulate f at $P \in \mathcal{P}$ via deviation $P'_S \in \times_{i\in S} \mathcal{P}_S$ in a self-enforcing manner such that S is a minimal manipulating coalition at P via P'_S. Let $j \in S$ such that $f_j(P'_S, P_{-S}) P_j f_j(P)$. Note that $S \setminus \{j\} \neq \emptyset$ since f is strategy-proof. Let $i \in S \setminus \{j\}$. Then $f_i(P'_S, P_{-S}) R_i f_i(P_i, P'_{S\setminus\{i\}}, P_{-S})$. Since f

[26] Since circulation rules are individual-good-preference based, an equivalent definition of self-enforcing group-strategy-proofness is obtained by additionally requiring that the deviation P'_S be a profile of lexicographic preferences.

[27] Related concepts are studied by Barberà et al. [3] and Serizawa [28].

is strategy-proof, this implies that $f_i(P'_S, P_{-S}) = f_i(P_i, P'_{S\setminus\{i\}}, P_{-S})$ and thus, by nonbossiness, $f(P'_S, P_{-S}) = f(P_i, P'_{S\setminus\{i\}}, P_{-S})$. Let $T = S \setminus \{i\}$. Then, since $j \in T$, coalition T is a manipulating coalition at P via P'_T. Therefore, S is not a minimal manipulating coalition at P via P'_S, which is a contradiction. □

Proposition 11 provides a counterpart to the equivalence shown in Lemma 1 in Pápai [17], one that also applies to markets with multiple units of each type of good.

Proposition 12. *For any capacity configuration, Segmented Trading Cycle rules are self-enforcing group-strategy-proof.*

Proof. As stated in Proposition 9, Segmented Trading Cycle rules are strategy-proof and nonbossy. Then Proposition 11 implies that Segmented Trading Cycle rules are self-enforcing group-strategy-proof. □

We can see that in the market specified in the proof of Proposition 10 the manipulation by coalition $\{b, c\}$ is due to preferences over bundles and, in particular, agent c's preferences are not lexicographic. For an illustration that Segmented Trading Cycle rules are self-enforcing group-strategy-proof, one can check that the manipulation by coalition $\{b, c\}$ in the same market is not self-enforcing. In fact, it is similar to a Prisoner's Dilemma situation, since each of agents b and c is strictly better off by reporting her true preferences rather than her manipulation preferences, regardless of whether the other one reports truthfully or goes ahead with reporting the manipulation preferences.[28]

Another relaxation of group-strategy-proofness is ig-group-strategy-proofness, which requires group-strategy-proofness for lexicographic preferences (see Definition 8). The next result demonstrates that Segmented Trading Cycle rules are not ig-group-strategy-proof in general.

Proposition 13. *There are capacity configurations for which some Segmented Trading Cycle rules are not ig-group-strategy-proof.*

Proof. Consider the market (N, q, P) where $N = \{a, b, c, d, e\}$, $q_a = q_b = q_c = 2$ and $q_d = q_e = 3$. Let preferences P be lexicographic such that the underlying preferences over individual goods are as follows:

$\succ_a$: $c \succ_a e \succ_a b \succ_a a$
$\succ_b$: $a \succ_b d \succ_b c \succ_b b$
$\succ_c$: $b \succ_c d \succ_c a \succ_c c$
$\succ_d$: $b \succ_d c \succ_d d$
$\succ_e$: $a \succ_e e$

Let the Segmented Trading Cycle rule be given by $W_1 = \{a, b, d, e\}$, $W_2 = \{b, c, d, e\}$, $W_3 = \{a, c, d, e\}$. Then this Segmented Trading cycle rule yields circulation x at profile P with

[28] For an illustration, let $q_a = 1 < 2 = q_b = q_c = q_d$. Assume agents a and d report their preferences truthfully. If b tells the truth, then c gets bundle $\{a, c\}$ if she tells the truth and bundle $\{b, c\}$ if she lies. If b lies, then c gets bundle $\{a, b\}$ if she tells the truth and bundle $\{b, b\}$ if she lies. If c tells the truth, then b gets bundle $\{d, d\}$ if she tells the truth and bundle $\{c, d\}$ if she lies. If c lies, then b gets bundle $\{a, d\}$ if she tells the truth and bundle $\{a, c\}$ if she lies.

$x_{ae} = 2, x_{bd} = 2$, $x_{cd} = 1$, and $x_{cc} = 1$. Suppose agents a, b and c now report preferences P'_a, P'_b and P'_c such that the underlying preferences over individual goods are:

$\succ'_a$: $c \succ_a b \succ_a a$
$\succ'_b$: $a \succ_b c \succ_b b$
$\succ'_c$: $b \succ_c a \succ_c c$

Then, at profile $(P'_{\{a,b,c\}}, P_{\{d,e\}})$, the above Segmented Trading Cycle rule yields circulation x' with $x'_{ac} = 1$, $x'_{ba} = 1$, and $x'_{cb} = 1$. Since preferences are lexicographic, this rule is not ig-group-strategy-proof. □

Remark 4. If a circulation rule is group-strategy-proof then it is both ig-group-strategy-proof and self-enforcing group-strategy-proof. Since Segmented Trading Cycle rules are self-enforcing-group-strategy-proof (Proposition 12) but not ig-group-strategy-proof in general (Proposition 13), self-enforcing-group-strategy-proofness does not imply ig-group-strategy-proofness. We can also verify that ig-group-strategy-proofness does not imply self-enforcing group-strategy-proofness. This can be demonstrated by the following variation on the cTTC rule, which we will refer to as the *Single-cTTC* rule: run the cTTC algorithm, but remove agents from the market after they participated in any one round of trading (which may involve multiple units). As a result, each agent may obtain at most one other agent's goods (but possibly more than one unit). The Single-cTTC rule is ig-group-strategy-proof, which follows from the group-strategy-proofness of the classical TTC rule. To see this, observe that if a coalition successfully manipulates the Single-cTTC rule then each manipulating agent who strictly gains would have to get either some units of a higher-ranked good or more units of the same good that she is receiving without manipulation, provided that these agents' preferences are lexicographic. Obtaining any units of a higher-ranked good is ruled out by group-strategy-proofness of the classical TTC rule, and receiving more units of the same good via a single trading cycle is not feasible, since the number of units exchanged in each trading cycle is bounded by the lowest-capacity agent(s) in the cycle. However, the Single-cTTC rule is manipulable when preferences are not lexicographic, since more units of a lower-ranked good may be preferred to fewer units of the good that is received. Therefore, the Single-cTTC rule is not strategy-proof, and hence it is not self-enforcing group-strategy-proof, given Proposition 11. ⋄

Note that it follows from Propositions 8 and 12 that for irreducible capacity configurations the cTTC rule is self-enforcing group-strategy-proof. While the cTTC rule is not necessarily strategy-proof for reducible capacity configurations (Proposition 6), the cTTC rule is strategy-proof for irreducible capacity configurations (Corollary 2), in which case it coincides with the unique Segmented Trading Cycle rule (Proposition 8). The cTTC rule is also ig-group-strategy-proof for irreducible capacity configurations, which is proved below.

Proposition 14. *For any irreducible capacity configuration, the cTTC rule (or, equivalently, the unique Segmented Trading Cycle rule) is ig-group-strategy-proof.*

Proof. First note that, by Proposition 7, there is a unique Segmented Trading Cycle rule when the capacity configuration is irreducible and, by Proposition 8, it is the same as the cTTC rule. The statement on ig-group-strategy-proofness is obvious for $n = 2$ and in the case where all agents have the same capacities, that is, for (ird1) and (ird3) irreducible capacity configurations, respectively. Indeed, for these irreducible capacity configurations the cTTC rule (or, equivalently, the unique

Segmented Trading Cycle rule) is group-strategy-proof, which implies ig-group-strategy-proofness. Thus, we only need to consider (ird2), the case where $n = 3$ with $N = \{a, b, c\}$ and $q_a < q_b = q_c$. Suppose that there exists some (ird2) capacity configuration for which the cTTC rule (or the unique Segmented Trading cycle rule) f is not ig-group-strategyproof. Note that the unique segmentation is given by q_a segments of $\{a, b, c\}$ and $q_b - q_a$ segments of $\{b, c\}$. We will refer to the identical market segments with a as W_a segments, and the identical market segments without a as W_{-a} segments. Since f is not ig-group-strategy-proof, there exists a coalition $S \subseteq N$ that can manipulate f at $P \in \mathcal{P}^L$ via deviation $P'_S \in \times_{i \in S} \mathcal{P}_i$. Without loss of generality, let S be a manipulating coalition. Since the cTTC rule is Pareto-efficient for irreducible capacity configurations (Corollary 1), the coalition of all three agents cannot be a manipulating coalition, so we only need to check coalitions of two agents, given that individual manipulations are ruled out by Corollary 2. Since the classical TTC rule is group-strategy-proof, there exists $i \in S$ who gets a strictly worse assignment in W_a segments at (P'_S, P_{-S}) than at P, which implies that i gets a strictly better assignment in W_{-a} segments at (P'_S, P_{-S}) than at P, since $f_i(P'_S, P_{-S}) R_i f_i(P)$. Similarly, there exists $j \in S \setminus \{i\}$ who gets a strictly worse assignment in W_{-a} segments at (P'_S, P_{-S}) than at P, which implies that j gets a strictly better assignment in W_a segments at (P'_S, P_{-S}) than at P. Since agent a does not participate in W_{-a} segments, this implies that $S = \{b, c\}$.

Assume without loss of generality that b is strictly worse off in W_a segments and strictly better off in W_{-a} segments at (P'_S, P_{-S}) than at P, which implies that c is strictly worse off in W_{-a} segments and strictly better off in W_a segments at (P'_S, P_{-S}) than at P, given the true preferences $P_b \in \mathcal{P}_b^L$ and $P_c \in \mathcal{P}_c^L$. Let $\succ \equiv \succ^P$ and $\succ' \equiv \succ^{(P'_S, P_{-S})}$. Since W_{-a} segments consist of agents b and c and the cTTC rule only assigns acceptable goods to agents, b can only be strictly better off in these market segments at the manipulation profile if $c \succ_b b$, $c \succ_c b$, and $b \succ'_c c$. Then, given that c is strictly better off in W_a segments at (P'_S, P_{-S}) than at P, it must be the case that $a \succ_c c$, since otherwise c has no acceptable good according to $\succ_c$, and thus the cTTC rule cannot assign any other good but her own at P, and agent c cannot be strictly better off at (P'_S, P_{-S}) than at P in any segment. Therefore, $a \succ_c c \succ_c b$. Therefore, in order for c to be strictly better off in W_a segments at (P'_S, P_{-S}) than at P, c has to obtain a in W_a segments at (P'_S, P_{-S}) but not at P. Then, given that $\succ_c$ ranks a first, it cannot be that case that $\succ_a$ ranks c first, so it must rank b first. The only other option would be that a ranks her own good first, but then c could not obtain any units of a at the manipulation profile. Thus, $b \succ_a c$ and $b \succ_a a$. Then, since c does not obtain a at P, $\succ_b$ cannot rank c first. Given $c \succ_b b$, this implies that $\succ_b$ ranks a first and thus $a \succ_b c \succ_b b$. This implies that $x_{ba} = q_a$ and $x'_{ba} = 0$, where $x \equiv f(P)$ and $x' \equiv (P'_S, P_{-S})$. Since preferences P_b are lexicographic, it follows that $f_b(P) P_b f_b(P'_S, P_{-S})$, which is a contradiction. □

There is a unique Segmented Trading Cycle rule if and only if the capacity configuration is quasi-irreducible (Proposition 7). If the capacity configuration is irreducible, then the unique Segmented Trading Cycle rule coincides with the cTTC rule (Proposition 8), which in this case is also ig-group-strategy-proof (Proposition 14).

For the quasi-irreducible but not irreducible capacity configurations the cTTC rule is not strategy-proof. This is because in these cases the cTTC may still reduce the capacity configuration in some round of the procedure to a configuration that is not quasi-irreducible. For example, if $|N| = 4$ with $N = \{a, b, c, d\}$, $q_a = 1$ and $q_b = q_c = q_d = 2$, then a top trading cycle of coalition $\{a, b, c\}$ reduces the capacity configuration to $q_a = 0, q_b = q_c = 1$, and $q_d = 2$, which is not a quasi-irreducible capacity configuration. This explains intuitively why the irreducible capacity configurations are the only exceptions in the impossibility theorem (Theorem 2) whenever the market segmentation is unique both for the entire problem and for all of its "reduced" problems that can be reached via top trading by the cTTC, the unique Segmented Trading Cycle rule coincides with the cTTC rule (see Proposition 8), which allows for combining strategy-proofness with Pareto-efficiency. As explained earlier, this happens only if the capacity configuration is such that when we reduce the available units for any non-singleton subset of the agents by the same number and at least one of these agents reaches exactly zero units, we still get an irreducible configuration. This holds only in the restricted cases when the capacity configuration is irreducible, and does not hold for quasi-irreducible configurations in general. Therefore, in the quasi-irreducible cases that are not irreducible, even though the Segmented Trading Cycle rule is unique, it is not identical to the cTTC rule, and this precludes the reconciliation of strategy-proofness and Pareto-efficiency.

5.3 The Sequentially Segmented Trading Cycle rule

We now define a specific member of the class of Segmented Trading Cycle rules which sequentially fills up the market segments with objects as much as possible, first segment 1, then segment 2, and so on. We call this rule the Sequentially Segmented Trading Cycle rule. More formally, the **Sequentially Segmented Trading Cycle (SSTC) rule** is obtained if market segments are such that W_1 to W_{q_1} contain one unit of each good, where q_1 denotes the minimum capacity among all agents, W_{q_1+1} to W_{q_2} contains one unit of each good that is still available, where q_2 denotes the minimum capacity among the remaining agents who can participate in segment $q_1 + 1$, and so on. Note that Proposition 10 also proves that the SSTC rule is not group-strategy-proof in general, while Proposition 14 shows that it is ig-group-strategy-proof for irreducible capacity configurations.

We will characterize the SSTC rule using the axiom of *top unanimity*, which ensures that whenever there is an agreement among a set of agents regarding their respective first choices, in the sense of a feasible trade as in the classical TTC rule, this trade is carried out with the highest possible flow. Top unanimity calls for trading goods in an intuitive manner, and while it is not implied by Pareto-efficiency, it guarantees the most obvious individually rational and Pareto-improving trades corresponding to top-ranked choices. This is in contrast to the incentive properties of strategy-proofness, nonbossiness, and the three discussed versions of group-strategy-proofness, all of which are compatible with not trading any units, even if trade is mutually desired by all respective parties.

Definition 10. A circulation rule f is **top unanimous** if for all $P \in \mathcal{P}$ and $S \subseteq N$ such that S is a top trading cycle coalition at P (that is, $S = \{i_1, i_2, \ldots, i_k\}$ such that for all $t = 1, \ldots, k$, modulo k, i_t's top-ranked good is i_{t+1}), $f(P)$ reflects the corresponding trade for S with the maximum flow (that is, the volume of the exchange corresponds to the capacity of the minimum-capacity agent in S). More precisely, for all $i \in N$, if there is a top trading cycle coalition S at P such that $i \in S$ and i is assigned k units of good j in this top trading cycle then $x_{ij} \geq k$, where $x = f(P)$.

Top unanimity means that the circulation assigned to each preference profile is consistent with the first round of the cTTC rule, which makes at least one of the trading agents filled. It is easy to see that not only the cTTC rule satisfies top unanimity, but also the SSTC rule, since the sequentially filled segments give rise to the same first-round trades in segments which contain a unit of good from each agent.

We show next that the cTTC rule and the SSTC rule are only identical for irreducible capacity configurations. This is not a coincidence, since this equivalence is closely related to the main impossibility result and characterization stated in Theorem 2, as we will see at the end of this section.

Proposition 15. *The cTTC rule and the SSTC rule are identical if and only if the capacity configuration is irreducible.*

Proof. Given the proof of Proposition 8, it suffices to show that if the capacity configuration is in the (rd1) category then the two rules are not identical. That is, if there exist distinct $a, b, c \in N$ with $q_a \leq q_b < q_c$, then there exists some preference profile to which the cTTC and the SSTC assign different circulations. Let $\succ$ satisfy the following: $c \succ_a a$, $c \succ_b b$, and $a \succ_c b \succ_c c$. Then $\{a, c\}$ is a top trading cycle in the first round and thus, by the top unanimity of both the cTTC and SSTC, a and c trade units corresponding to the maximum flow q_a. Thus, in the cTTC rule $\{b, c\}$ is a top trading cycle coalition in the next round of the procedure, and b and c then trade units corresponding to the maximum flow which is given by $\min\{q_b, q_c - q_a\}$. However, in the SSTC rule coalition $\{b, c\}$ trades $q_b - q_a$ units only. Since $q_a > 0$ and $q_b < q_c$, $q_b - q_a < \min\{q_b, q_c - q_a\}$. Therefore, the cTTC and the SSTC rules assign different circulations to $\succ$. □

Given Proposition 15, we now revisit and further refine the statement of Theorem 2 as follows.

Theorem 2 (revisited). *There is a circulation rule that is individually rational, Pareto-efficient, and strategy-proof if and only if the capacity configuration is irreducible. For irreducible capacity configurations, the SSTC rule, which is identical to the cTTC rule, is the* unique *circulation rule that satisfies all three properties.*

Proposition 16. *For any capacity configuration, if a circulation rule is top unanimous and strategy-proof then it is also individually rational.*

Proof. Suppose that a circulation rule f is top unanimous and strategy-proof, but there exists $P \in \mathcal{P}$ such that individual rationality is violated at P by f. Then there exists $i \in N$ such that for the null bundle 0_i, $0_i P_i f_i(P)$. Let $P'_i \in \mathcal{P}_i$ such that at $\succ^{P'_i}$ all goods $j \neq i$ are unacceptable. Then top unanimity implies that $f_i(P'_i, P_{-i}) = 0_i$. Thus, $f_i(P'_i, P_{-i}) P_i f_i(P)$, which contradicts strategy-proofness. Therefore, f is individually rational. □

Our final main result is a characterization of the SSTC rule by top unanimity and self-enforcing group-strategy-proofness. Note that Proposition 16 makes individual rationality a redundant property for this characterization, although the SSTC rule satisfies it.

Theorem 3. *For any capacity configuration, the SSTC rule is the unique circulation rule that is top unanimous and self-enforcing group-strategy-proof.*

The proof of Theorem 3 is in Appendix B. We know that the SSTC rule is top unanimous and have already shown that it is self-enforcing group-strategy-proof (Proposition 12), and thus the proof demonstrates the converse, namely that the only top unanimous and self-enforcing group-strategy-proof rule is the SSTC rule. The proof proceeds in several steps. We show first that each circulation can be decomposed into market segments and that each circulation assigned to a preference profile by a top unanimous and self-enforcing group-strategy-proof rule has the same market segment decomposition at each preference profile. Moreover, given top unanimity, this market segment decomposition is the same as the market segment decomposition of the SSTC rule. Then we show that strategy-proofness is satisfied within each market segment, and this combined with individual rationality and top unanimity implies that the classical TTC outcome is assigned within each market segment at each preference profile, completing the proof. The last step also provides a new characterization of the classical TTC rule using individual rationality, top unanimity and strategy-proofness (replacing Pareto-efficiency by top unanimity in the characterization of Ma [14]).

Remark 5. Self-enforcing group-strategy-proofness and top unanimity are logically independent properties. To see this note that, by definition, the cTTC rule is top unanimous. Moreover, it follows from Proposition 6 that for reducible capacity configurations the cTTC rule is not necessarily strategy-proof. Thus, by Proposition 11, the cTTC rule is not self-enforcing group-strategy-proof. Therefore, the cTTC rule is top unanimous but not self-enforcing group-strategy-proof. On the other hand, the no-trade rule, or any Segmented Trading Cycle rule other than the SSTC rule, is self-enforcing group-strategy-proof (Propositions 11 and 12) but does not satisfy top unanimity. ◇

Note that the SSTC rule is not ig-group-strategy-proof for all capacity configurations,[29] and thus this axiom cannot replace self-enforcing group-strategy-proofness in the characterization of the SSTC rule. However, since the Single-cTTC rule is top unanimous by definition and satisfies ig-group-strategy-proofness (see Remark 4), these two axioms are compatible. Indeed, these properties of the Single-cTTC rule also imply that we cannot substitute ig-group-strategy-proofness for self-enforcing group-strategy-proofness in Theorem 3.

As a corollary to Theorem 3 and Proposition 11, we obtain the following alternative characterization of the SSTC rule.

Corollary 3. *For any capacity configuration, the SSTC rule is the unique circulation rule that is top unanimous, strategy-proof, and nonbossy.*

Another corollary to Theorem 3, using Proposition 10 and the fact that group-strategy-proofness implies self-enforcing group-strategy-proofness, is the following.

Corollary 4. *There are capacity configurations for which there exists no circulation rule that is top unanimous and group-strategy-proof.*

We can now describe the tradeoffs regarding the two main proposed circulation rules, cTTC and SSTC, more precisely. While both rules satisfy individual rationality, top unanimity, and nonbossiness, cTTC does better in terms of efficiency since it is ig-Pareto-efficient, while SSTC satisfies the weak efficiency property of "efficiency-on-the-range."[30] On the other hand, SSTC satisfies strategy-proofness (in fact, it satisfies the more demanding self-enforcing group-strategy-proofness property), while cTTC is only ig-dropping-proof, a weaker incentive property than strategy-proofness.

[29] Proposition 10 demonstrates that the SSTC rule is not group-strategy-proof. A larger example is required than the one in the proof of this proposition to show that the SSTC rule does not satisfy ig-group-strategy-proofness either (assuming that the capacity configuration is not irreducible, since for irreducible capacity configurations it is ig-group-strategy-proof, by Proposition 14). Such an example is available from the authors upon request.

[30] Note that "efficiency-on-the-range" is not implied by ig-Pareto-efficiency (but, clearly, implied by Pareto-efficiency), and is in fact not satisfied by the cTTC rule.

To conclude, we use the two characterization theorems (Theorems 1 and 3) to give a short intuitive proof of the impossibility result stated in Proposition 2, showing that the three main axioms in the classical characterization result of the TTC by Ma (1994) cannot be reconciled for multiple-unit Shapley-Scarf markets when the capacity configuration is reducible.

Proof of Proposition 2. Let a circulation rule f be individually rational, Pareto-efficient and strategy-proof. Then f is ig-Pareto-efficient and ig-dropping-proof, and thus by Theorem 1 f must be the cTTC rule for any capacity configuration. This implies that f is also top unanimous and nonbossy, since the cTTC rule satisfies these properties. Then, given that f is also strategy-proof and thus self-enforcing group-strategy-proof for any capacity configuration, as shown by Proposition 11, Theorem 3 implies that f must be the SSTC rule for any capacity configuration. By Proposition 15, the cTTC and the SSTC rules are only identical for irreducible capacity configurations, and therefore we get a contradiction when the capacity configuration is reducible. This implies that if the capacity configuration is reducible then there is no circulation rule that is individually rational, Pareto-efficient and strategy-proof. □

6 Extensions and future research

Instead of (or besides) the agent-capacities we could have **link-capacities** on the number of units the agents can send to others through specific links (arcs in $D(N, A(P))$). This is a very typical setting for circulation problems in graph theory, and some practical applications do have this kind of requirement, e.g., in the Erasmus exchange program the number of students from university U that visit university V is bounded by the specifications in the bilateral contract between U and V. We opted for defining our model through node-capacities to easily relate it to existing models on the exchange of indivisible goods. However, any arc-capacitated market can always be transformed into a node-capacitated market under responsive preferences by introducing nodes for each arc, as follows. For each $(i, j) \in A(P)$, let $\overline{ij}$ be a new node with the same capacity as the original arc. Let j be replaced by $\overline{ij}$ in i's preference list and let node $\overline{ij}$ only find j acceptable. The feasible circulations are in one-to-one correspondence in the two markets. Moreover, the original agents evaluate any circulation in the same way in the two markets. Finally, since the new nodes do not have any strategic role in the extended market (they only have one partner to trade with), the manipulability of any circulation rule does not change from one setting to the other.

If agents have **heterogeneous goods** instead of homogeneous goods then some aspects of the circulation problem can still be studied in our model, although the strategic issues will be different. We can reduce the model with heterogeneous goods to our circulation model as follows. Let each good of the heterogeneous case be an artificial agent with unit capacity in our circulation model, and let her only find the original owner acceptable. Let all original agents preferences with respect to the new artificial agents be consistent with their original preferences. A generalization of the TTC rule for the heterogeneous case, which was introduced and studied in Fujita et al. [8], is equivalent to our cTTC rule for the reduced circulation market with homogeneous goods. The new artificial agents cannot manipulate the cTTC rule, as they have unit capacity. However, since these artificial agents are "owned" by the original agents, the strategic manipulations are different. For instance, a manipulation in which an agent in the heterogeneous goods market hides some of her goods corresponds to a group manipulation in the reduced market. The precise connections between the two markets and the properties of the circulation rules could be pursued in future research.

One could also consider different **input** for circulation rules. In this paper we only elicited the ordinal preferences of the agents over the goods, but circulation rules could also be based on agents' cardinal utilities of goods (see, e.g., Aziz et al. [2]), linear preferences over bundles, or some choice functions on the set of bundles. Different input for circulation rules would also make it possible to investigate more general preference domains, e.g., substitutable choice functions.

Another future line of research could focus on a relaxation of our assumption on **acceptable bundles**. We have assumed that a bundle is acceptable if and only if it does not contain any unacceptable good. A relaxation of this assumption could be of interest for some real-life applications and the results might turn to out be quite different from ours.

Finally, in our main impossibility result we characterize the **capacity configurations** for which individual rationality, Pareto-efficiency, and strategy-proofness are incompatible. It could be interesting to carry out a similar exercise for other combinations of desiderata.

Appendix

A Proof of Proposition 2

Proposition 2. *If the capacity configuration is reducible, then there is no circulation rule that is individually rational, Pareto-efficient, and strategy-proof.*

Proof. We first note that all preferences considered in the proof are assumed to satisfy (resp-3). In particular, it can be easily verified that all preferences that are constructed can indeed satisfy (resp-3). Throughout the proof we only consider individually rational circulations and circulation rules. However, to not further increase the length of the proof *we very often do not explicitly refer to individual rationality.* We prove the result for the two classes of reducible capacity configurations separately.

Case 1: there are different $a, b, c \in N$ with $q_b > q_c \geq q_a$ (rd1). Without loss of generality we may assume that $N = \{a, b, c\}$. (Any additional agents can be assumed to prefer the empty bundle to any other bundle.) We show that any Pareto-efficient circulation rule is not strategy-proof. We first provide the set of Pareto-efficient circulations for 4 different types of preference profiles (types I, II, III, and IV).

Type I: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $b \succ_a a$,[31] $c \succ_b a \succ_b b$, and $a \succ_c b \succ_c c$. Figure 2 depicts the underlying preferences and the resulting graph $D(N, A(P))$.[32]

Claim 1. *For any $P \in \mathcal{P}$ of Type I, a circulation x is Pareto-efficient if and only if*
(i). $x_{ca} + x_{ba} = q_a$;
(ii). $x_{ca} + x_{cb} = q_c$;
(iii). $q_c + q_a - q_b \leq x_{ca}$.

[31] In particular, $a \succ_a c$. Because of individual rationality we always omit unacceptable goods from the description of the preferences.

[32] Edges correspond to acceptable goods. Continuous edges denote most preferred goods and discontinuous edges denote second most preferred goods.

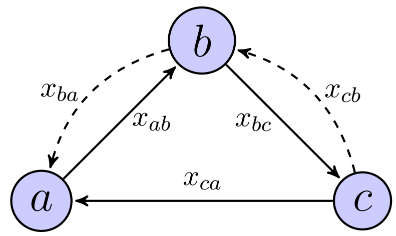

Figure 2: $\succ^P$ and $D(N, A(P))$ of Type I

Proof. Let x be a Pareto-efficient circulation. We show that it satisfies (i), (ii), and (iii).

Suppose (i) does not hold. Then, by balancedness and feasibility, $x_{a\cdot} = x_{ab} + x_{ac} = x_{ca} + x_{ba} < q_a$, i.e., a is unfilled at x. Suppose $x_{cb} = 0$. Then, $x_{c\cdot} = x_{ca} \leq x_{ca} + x_{ba} < q_a \leq q_c$ and $x_{b\cdot} = x_{\cdot b} = x_{ab} \leq q_a < q_b$. Hence, b and c are unfilled at x. But then x is Pareto-dominated by only increasing the flow in cycle (b, c) until b or c becomes filled.[33] Suppose now $x_{cb} > 0$. Let x' be the circulation obtained from x by decreasing the flow in cycle (b, c) by $\delta = \min\{q_a - x_{a\cdot}, x_{cb}\} > 0$ and increasing the flow in cycle (a, b, c) by δ. Then, agents a and c will be strictly better off at x', while agent b gets the same bundle at x and x'. Hence, x' Pareto-dominates x. So, (i) holds (and a is filled at x).

Suppose (ii) does not hold. Then, by balancedness and feasibility, $x_{c\cdot} < q_c$, i.e., c is unfilled at x. Suppose $x_{ba} = 0$. Then, $x_{b\cdot} = x_{bc} \leq x_{\cdot c} = x_{c\cdot} < q_c < q_b$, i.e., b is unfilled at x. Hence, x is Pareto-dominated by increasing the flow in cycle (b, c). Suppose now $x_{ba} > 0$. Let x' be the circulation obtained from x by decreasing the flow in cycle (a, b) by $\delta = \min\{q_c - x_{c\cdot}, x_{ba}\} > 0$ and increasing the flow in cycle (a, b, c) by δ. Then, agents b and c will be strictly better off at x', while agent a gets the same bundle at x and x'. Hence, x' Pareto-dominates x. So, (ii) holds (and c is filled at x).

By individual rationality, $x_{ac} = 0$. Hence, by feasibility at b and balancedness at a, $q_b \geq x_{\cdot b} = x_{ab} + x_{cb} = (x_{ab} + x_{ac}) + x_{cb} = (x_{ca} + x_{ba}) + x_{cb}$. By (i), $x_{ca} = q_a - x_{ba}$, and by (ii), $x_{cb} = q_c - x_{ca}$. Hence, $q_b \geq (q_a - x_{ba}) + x_{ba} + (q_c - x_{ca}) = q_a + q_c - x_{ca}$, which proves (iii).

Let x be a circulation that satisfies (i), (ii), and (iii). Suppose x is not Pareto-efficient. Then, since there is a finite number of goods, x is Pareto-dominated by some Pareto-efficient circulation y. In particular, y satisfies (i), (ii), and (iii).

Suppose $y_{ca} = x_{ca}$. Then, from (i) and (ii) for y and balancedness of y it follows that $x = y$, which contradicts that y Pareto-dominates x.

Suppose $y_{ca} < x_{ca}$. At x, agent c's bundle consists of $x_{cb} = q_c - x_{ca}$ units of b and x_{ca} units of a (her most preferred good). At y, agent c's bundle consists of $y_{cb} = q_c - y_{ca}$ units of b and y_{ca} units of a. Hence, agent c strictly prefers x_c to y_c, which contradicts that y Pareto-dominates x.

Suppose $y_{ca} > x_{ca}$. At x, agent b's bundle consists of $x_{ba} = q_a - x_{ca}$ units of a and $x_{bc} = x_{ca} + x_{cb} = x_{ca} + (q_c - x_{ca}) = q_c$ units of c. Similarly, at y, agent b's bundle consists of $y_{ba} = q_a - y_{ca}$ units of a and $y_{bc} = q_c$ units of c. Hence, agent b strictly prefers x_b to y_b, which contradicts that y Pareto-dominates x.

Hence, x is Pareto-efficient. □

[33] More precisely, "increasing the flow in a cycle" means that all agents in the cycle consume an additional and same amount of good of the next agent in the cycle. We will use similarly the expression "decreasing the flow in a cycle."

Type II: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $b \succ_a a$, $a \succ_b c \succ_b b$, and $a \succ_c b \succ_c c$. Figure 3 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

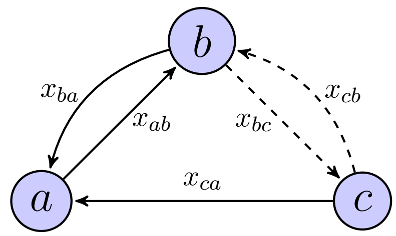

Figure 3: $\succ^P$ and $D(N, A(P))$ of Type II

Claim 2. *For any $P \in \mathcal{P}$ of Type II, a circulation x is Pareto-efficient if and only if*
(i). $x_{ca} + x_{ba} = q_a$;
(ii). $x_{ba} + x_{bc} = q_b$ or $x_{ca} + x_{cb} = q_c$;
(iii). $x_{cb} = \min\{q_b - (x_{ca} + x_{ba}), q_c - x_{ca}\} = \min\{q_b - q_a, q_c - q_a + x_{ba}\}$.

Proof. Let x be a Pareto-efficient circulation. We show that it satisfies (i), (ii), and (iii).

Suppose (i) does not hold. Then, by feasibility, $x_{ca} + x_{ba} < q_a$, i.e., a is unfilled at x. Suppose $x_{cb} = 0$. Then, $x_{c\bullet} = x_{ca} \leq x_{ca} + x_{ba} < q_a \leq q_c$ and $x_{b\bullet} = x_{\bullet b} = x_{ab} \leq q_a < q_b$. Hence, b and c are unfilled at x. But then x is Pareto-dominated by increasing the flow in cycle (b, c). Suppose now $x_{cb} > 0$. Let x' be the circulation obtained from x by decreasing the flow in cycle (b, c) by $\delta = \min\{q_a - x_{a\bullet}, x_{cb}\} > 0$ and increasing the flow in cycle (a, b, c) by δ. Then, agents a and c will be strictly better off at x', while agent b gets the same bundle at x and x'. Hence, x' Pareto-dominates x. So, (i) holds (and a is filled at x).

Suppose (ii) does not hold. Then, both b and c are unfilled at x. But then x is Pareto-dominated by increasing the flow in cycle (b, c). So, (ii) holds.

We now prove (iii). Note that $x_{cb} + x_{ab} \leq q_b$ is equivalent to $x_{cb} \leq q_b - x_{ab} = q_b - (x_{ca} + x_{ba})$.[34] Similarly, $x_{cb} + x_{ca} \leq q_c$ is equivalent to $x_{cb} \leq q_c - x_{ca}$. In particular, b is filled if and only if $x_{cb} = q_b - (x_{ca} + x_{ba})$ and c is filled if and only if $x_{cb} = q_c - x_{ca}$. Since (ii) holds, b or c is filled at x, and hence $x_{cb} = \min\{q_b - (x_{ca} + x_{ba}), q_c - x_{ca}\}$. By (i), $x_{ca} = q_a - x_{ba}$. Hence, $x_{cb} = \min\{q_b - (q_a - x_{ba} + x_{ba}), q_c - (q_a - x_{ba})\} = \min\{q_b - q_a, q_c - q_a + x_{ba}\}$. This proves (iii).

Let x be a circulation that satisfies (i), (ii), and (iii). Suppose x is not Pareto-efficient. Then, x is Pareto-dominated by some Pareto-efficient circulation y. Note y satisfies (i), (ii), and (iii).

Suppose $y_{ba} = x_{ba}$. Then, from (i), $y_{ca} = x_{ca}$. By balancedness at a, $y_{ab} = y_{ba} + y_{ca} = x_{ba} + x_{ca} = x_{ab}$. Moreover, by (iii), $x_{cb} = y_{cb}$. By balancedness at c, $y_{bc} = y_{ca} + y_{cb} = x_{ca} + x_{cb} = x_{bc}$. Hence, $y = x$, which contradicts that y Pareto-dominates x.

Suppose $y_{ba} < x_{ba}$. At x, agent b's bundle consists of x_{ba} units of a (her most preferred good) and $x_{bc} = x_{ca} + x_{cb} = x_{ca} + \min\{q_b - (x_{ca} + x_{ba}), q_c - x_{ca}\} = \min\{q_b - x_{ba}, q_c\}$ units of c. Similarly, at y, agent b's bundle consists of y_{ba} units of a and $y_{bc} = \min\{q_b - y_{ba}, q_c\}$ units of c. Hence, agent b strictly prefers x_b to y_b, which contradicts that y Pareto-dominates x.

Suppose $y_{ba} > x_{ba}$. At x, agent c's bundle consists of $x_{ca} = q_a - x_{ba}$ units of a (her most preferred good) and $x_{cb} = \min\{q_b - q_a, q_c - q_a + x_{ba}\}$ units of b. Similarly, at y, agent b's bundle

[34] Here we use that x is an individually rational circulation and that preferences satisfy (resp-3).

consists of $y_{ca} = q_a - y_{ba}$ units of a and $y_{cb} = \min\{q_b - q_a, q_c - q_a + y_{ba}\}$ units of b. Hence, agent c strictly prefers x_c to y_c, which contradicts that y Pareto-dominates x.

Hence, x is Pareto-efficient. □

Type III: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $b \succ_a a$, $a \succ_b b$, and $a \succ_c b \succ_c c$.
Figure 4 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

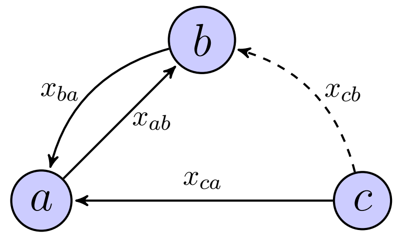

Figure 4: $\succ^P$ and $D(N, A(P))$ of Type III

Claim 3. *For any $P \in \mathcal{P}$ of Type III, a circulation x is Pareto-efficient if and only if $x_{ba} = q_a$.*

Proof. The only cycle in $D(N, A(P))$ is (a, b). Hence, a circulation x is Pareto-efficient if and only if (a, b) has maximum flow, i.e., q_a. □

Type IV: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $b \succ_a a$, $c \succ_b a \succ_b b$, and $a \succ_c c$.
Figure 5 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

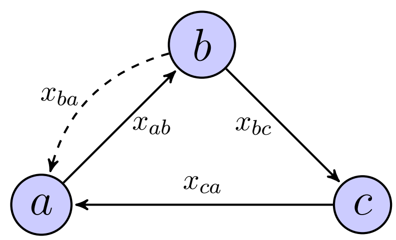

Figure 5: $\succ^P$ and $D(N, A(P))$ of Type IV

Claim 4. *For any $P \in \mathcal{P}$ of Type IV, a circulation x is Pareto-efficient if and only if $x_{ba} = 0$ and $x_{ca} = q_a$.*

Proof. Let x be a circulation with $x_{ba} > 0$. Since $x_{ba} + x_{ca} \leq q_a \leq q_c$, it follows that $\delta = \min\{q_c - x_{ca}, x_{ba}\} > 0$. Let x' be the circulation defined by $x'_{ba} = x_{ba} - \delta$, $x'_{bc} = x_{bc} + \delta$, $x'_{ca} = x_{ca} + \delta$, and $x'_{ab} = x_{ab}$. Then, x' Pareto-dominates x.

Let x be a circulation with $x_{ba} = 0$, i.e., there is only flow through cycle (a, b, c). Then, x is Pareto-efficient if and only if there is maximum flow through (a, b, c), i.e., $x_{ca} = q_a$. □

We can now prove the result for Case 1. Suppose that f is a circulation rule that is both Pareto-efficient and strategy-proof.

Let $P \in \mathcal{P}$ be of Type I such that c's preferences are lexicographic. Let $x = f(P)$. Suppose $x_{ca} < q_a$. Then, agent c can submit preferences P'_c such that $P' = (P_{-c}, P'_c)$ is of Type IV. Let $x' = f(P')$. Since c's preferences P_c are lexicographic, it follows from Claim 4 that agent c strictly prefers x'_c to x_c, which contradicts strategy-proofness of f. Therefore $x_{ca} = q_a$. Then, by Claim 1(i,ii),

$$x_{ba} = 0 \text{ and } x_{bc} = x_{ca} + x_{cb} = q_c.$$

Let $P \in \mathcal{P}$ be of Type I. Let $x = f(P)$. Since circulation rules are individual-good-preference based, i.e., only take into account the underlying ordinal preferences over individual goods (see (2) in Section 2), it follows from the above that

$$x_{ba} = 0 \text{ and } x_{bc} = x_{ca} + x_{cb} = q_c. \tag{3}$$

Let $P \in \mathcal{P}$ be of Type II such that b's preferences are lexicographic. Let $y = f(P)$. Suppose $y_{ba} < q_a$. Then, agent b can submit preferences P'_b such that $P' = (P_{-b}, P'_b)$ is of Type III. Let $y' = f(P')$. Since b's preferences P_b are lexicographic, it follows from Claim 3 that agent b strictly prefers y'_b to y_b, which contradicts strategy-proofness of f. Therefore $y_{ba} = q_a$.

Let $P \in \mathcal{P}$ be of Type II. Let $y = f(P)$. From the above and the fact that f is individual-good-preference based, it follows that

$$y_{ba} = q_a. \tag{4}$$

Moreover, $y_{bc} = y_{cb} + y_{ab} - y_{ba} = \min\{q_b - (y_{ca} + y_{ba}), q_c - y_{ca}\} + y_{ab} - y_{ba}$ by Claim 2(iii). By Claim 2(i) and (4), we have $y_{ca} = 0$. Then, $y_{ab} = y_{ba} + y_{ca} = q_a + 0 = q_a$. Hence,

$$y_{bc} = \min\{q_b - q_a, q_c\} + q_a - q_a = \min\{q_b - q_a, q_c\}. \tag{5}$$

Let $P \in \mathcal{P}$ be of Type I such that b's preferences are additive and

$$1 < \frac{u_b(c)}{u_b(a)} < \frac{\min\{q_b, q_a + q_c\}}{q_c}. \tag{6}$$

By stating the truth, it follows from (3) that agent b's total sum of utility is $q_c u_b(c)$. Agent b can submit preferences P'_b such that $P' = (P_{-b}, P'_b)$ is of Type II. From (4) and (5) it follows that in that case she obtains a bundle of at least $q_a + \min\{q_b - q_a, q_c\} = \min\{q_b, q_c + q_a\}$ units of goods which gives her a total sum of utility of at least $\min\{q_b, q_c + q_a\}u_b(a)$. From (6) it follows that agent b can manipulate f at P, which contradicts the strategy-proofness of f. Therefore, there is no circulation rule that is both Pareto-efficient and strategy-proof.

Case 2: there are different $a, b, c, d \in N$ with $q_b = q_c = q_a > q_d$ (rd2). With slight abuse of notation denote $q = q_a$. Without loss of generality we may assume again that $N = \{a, b, c, d\}$. Let f be a circulation rule that satisfies Pareto-efficiency. We will show that it is not strategy-proof. Suppose, to the contrary, that f is strategy-proof. We first study 4 different types of preference profiles (types 1–9).

Type 8: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a a$, $a \succ_b c \succ_b b$, $a \succ_c c$, and $c \succ_d d$. Figure 6 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

Claim 5. *For any $P \in \mathcal{P}$ of Type 8, a circulation x is Pareto-efficient if and only if $x_{ca} = q_d$.*

Proof. The only cycle in $D(N, A(P))$ is (c, a, d). Hence, a circulation x is Pareto-efficient if and only if (c, a, d) has maximum flow, i.e., q_d. □

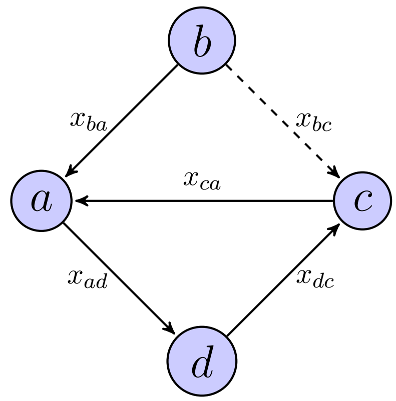

Figure 6: $\succ^P$ and $D(N, A(P))$ of Type 8

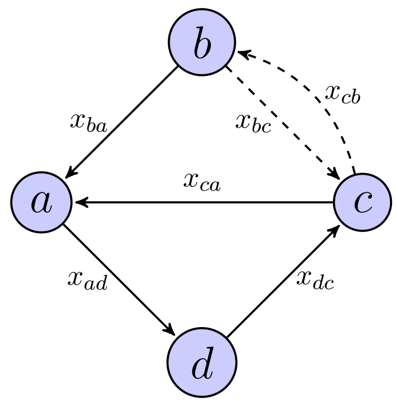

Figure 7: $\succ^P$ and $D(N, A(P))$ of Type 7

Type 7: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a a$, $a \succ_b c \succ_b b$, $a \succ_c b \succ_c c$, and $c \succ_d d$. Figure 7 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

Claim 6. *Let $P \in \mathcal{P}$ be of Type 7. Let $x = f(P)$. Then, $x_{ca} = q_d$.*

Proof. Since f is individual-good-preference based, we may assume that agent c has lexicographic preferences P_c. Then, if $x_{ca} < q_d$, agent c can manipulate f at P: from Claim 5 it follows that she can submit preferences P'_c such that she gets q_d units of a at $f(P')$ where $P' = (P_{-c}, P'_c)$. So, since f is strategy-proof, $x_{ca} = q_d$. □

Type 6: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a a$, $c \succ_b a \succ_b b$, $a \succ_c c$, and $c \succ_d d$. Figure 8 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

Claim 7. *For any $P \in \mathcal{P}$ of Type 6, a circulation x is Pareto-efficient if and only if $x_{ca} = q_d$.*

Proof. The only cycle in $D(N, A(P))$ is (c, a, d). Hence, a circulation x is Pareto-efficient if and only if (c, a, d) has maximum flow, i.e., q_d. □

Type 5: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a a$, $c \succ_b a \succ_b b$, $a \succ_c b \succ_c c$, and $c \succ_d d$. Figure 9 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

Claim 8. *Let $P \in \mathcal{P}$ be of Type 5. Let $x = f(P)$. Then, $x_{ca} = q_d$.*

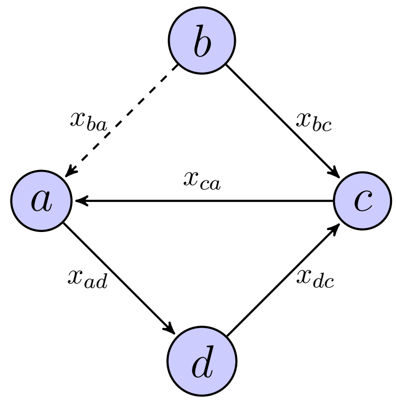

Figure 8: $\succ^P$ and $D(N, A(P))$ of Type 6

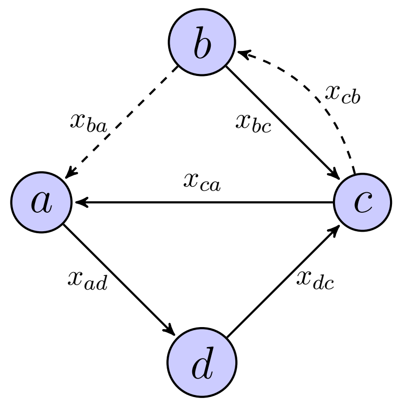

Figure 9: $\succ^P$ and $D(N, A(P))$ of Type 5

Proof. Since f is individual-good-preference based, we may assume that agent c has lexicographic preferences P_c. Then, if $x_{ca} < q_d$, agent c can manipulate f at P: from Claim 7 it follows that she can submit preferences P'_c such that she gets q_d units of a at $f(P')$ where $P' = (P_{-c}, P'_c)$. So, since f is strategy-proof, $x_{ca} = q_d$. □

Type 4: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a b \succ_a a$, $c \succ_b a \succ_b b$, $a \succ_c c$, and $c \succ_d d$. Figure 10 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

Claim 9. *Let $P \in \mathcal{P}$ be of Type 4. If x is a Pareto-efficient circulation, then $x_{ca} = q$.*

Proof. Suppose $x_{ca} \neq q$. Then, by feasibility, $x_{ca} < q$, i.e., c is unfilled.

Suppose $x_{ba} = 0$. Then, $x_{b\cdot} = x_{bc} \leq x_{ca} < q$. So, b is unfilled. Moreover, $x_{\cdot a} = x_{ca} < q$. So, also a is unfilled. Then, x is Pareto-dominated by increasing the flow in cycle (a, b, c), which contradicts Pareto-efficiency of x.

Now suppose $x_{ba} > 0$. Let $\delta = \min\{x_{ba}, q - x_{ca}\} > 0$. Let x' be the circulation obtained from x by setting $x'_{ba} = x_{ba} - \delta$, $x'_{bc} = x_{bc} + \delta$, and $x'_{ca} = x_{ca} + \delta$ (while maintaining the other flows). Then, each of agents b and c strictly prefers her bundle at x' to her bundle at x, while agents a and d are indifferent. This contradicts Pareto-efficiency of x. Hence, $x_{ca} = q$. □

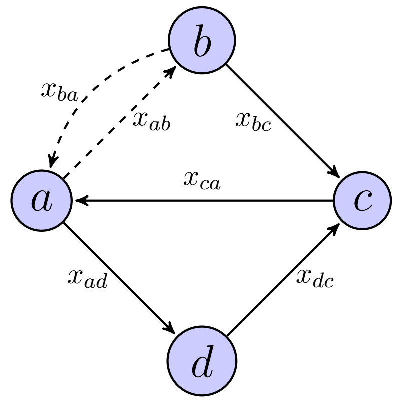

Figure 10: $\succ^P$ and $D(N, A(P))$ of Type 4

Type 1: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a b \succ_a a$, $c \succ_b a \succ_b b$, $a \succ_c b \succ_c c$, and $c \succ_d d$.
Figure 11 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

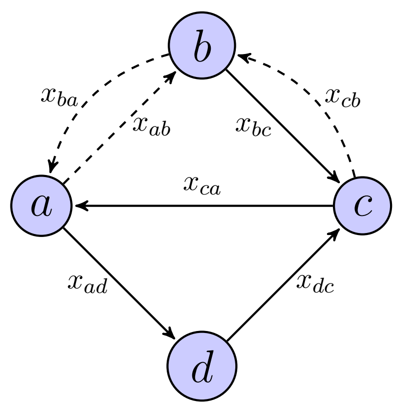

Figure 11: $\succ^P$ and $D(N, A(P))$ of Type 1

Claim 10. *Let $P \in \mathcal{P}$ be of Type 1. Let $x = f(P)$. Then,*
(i). d is filled, i.e., $x_{ad} = x_{dc} = q_d$;
(ii). $x_{ca} = q$;
(iii). $x_{ba} = x_{cb} = 0$ and $x_{bc} = x_{ab} = q - q_d$.

Proof. Since f is individual-good-preference based, we may assume that agents have lexicographic preferences. Suppose (i) does not hold. Then, by feasibility and balancedness at d, $x_{ad} = x_{dc} < q_d$. Then, agent a can submit preferences P'_a such that $P' = (P_{-a}, P'_a)$ is of Type 5. Let $x' = f(P')$. By Claim 8, $x'_{ca} = q_d$. Then, from balancedness of x' at a it follows that $x'_{ad} = x'_{ba} + x'_{ca} \geq q_d$. By feasibility of x' at d, $x'_{ad} \leq q_d$. Hence, $x'_{ad} = q_d$. Since agent a has lexicographic preferences, she prefers x'_a to x_a, which contradicts the strategy-proofness of f. Hence, (i) holds.

Suppose (ii) does not hold. Then, by feasibility, $x_{ca} < q$. Then, agent c can submit preferences P'_c such that $P' = (P_{-c}, P'_c)$ is of Type 4. Let $x' = f(P')$. By Claim 9, $x'_{ca} = q$. Since agent c has lexicographic preferences, she prefers x'_c to x_c, which contradicts the strategy-proofness of f. Hence, (ii) holds.

Finally, we show (iii). Since $x_{ca} = q = q_a = q_c$, it follows from feasibility and balancedness that $x_{cb} = 0 = x_{ba}$. Then, $x_{bc} = x_{ca}+x_{cb}-x_{dc} = q+0-q_d$. Similarly, $x_{ab} = x_{ca}+x_{ba}-x_{ad} = q+0-q_d$. $\square$

Type 9: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a b \succ_a a$, $a \succ_b c \succ_b b$, $a \succ_c c$, and $c \succ_d d$. Figure 12 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

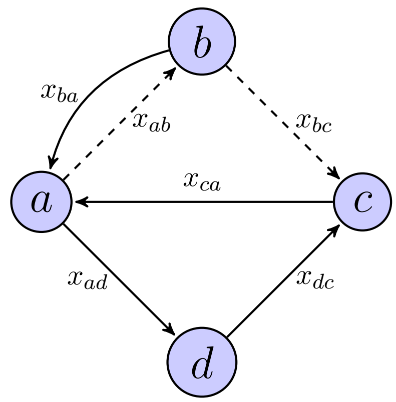

Figure 12: $\succ^P$ and $D(N, A(P))$ of Type 9

Claim 11. *Let $P \in \mathcal{P}$ be of Type 9. Let $x = f(P)$. Then, $x_{ca} \geq q_d$.*

Proof. Since f is individual-good-preference based, we may assume that agent a has lexicographic preferences P_a. Suppose $x_{ca} < q_d$. Then, $x_{ad} = x_{dc} = x_{ca} - x_{bc} \leq x_{ca} < q_d$. Agent a can submit preferences P'_a such that $P' = (P_{-a}, P'_a)$ is of Type 8. Let $x' = f(P')$. By Claim 5, $x'_{ca} = q_d$. Then, together with $x'_{ad} = x'_{dc} \leq q_d$ it follows that $x'_{ad} = x'_{ca} + x'_{ba} = q_d$. Therefore, agent a can manipulate f at P, which contradicts the strategy-proofness of f. Hence, $x_{ca} \geq q_d$. $\square$

Type 3: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a b \succ_a a$, $a \succ_b b$, $a \succ_c b \succ_c c$, and $c \succ_d d$. Figure 13 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

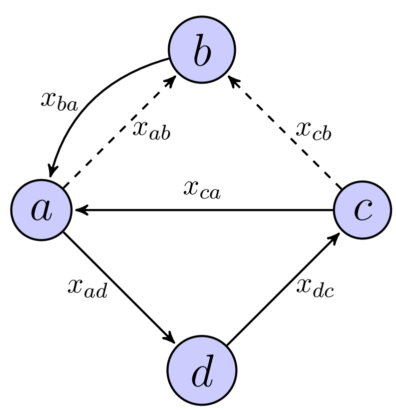

Figure 13: $\succ^P$ and $D(N, A(P))$ of Type 3

Claim 12. *Let $P \in \mathcal{P}$ be of Type 3. If x is a Pareto-efficient circulation, then $x_{ba} \geq q - q_d$.*

Proof. Note $x_{ca} \le x_{dc} = x_{ad} \le q_d$. Suppose $x_{ba} < q - q_d$. Then, $x_{ba} + x_{ca} < q - q_d + q_d = q$, i.e., a is unfilled. Since $x_{ba} < q$, b is also unfilled. Then x can be Pareto improved by increasing the flow in cycle (a, b). □

Type 2: $P \in \mathcal{P}$ is such that for $\succ \equiv \succ^P$ we have $d \succ_a b \succ_a a$, $a \succ_b c \succ_b b$, $a \succ_c b \succ_c c$, and $c \succ_d d$.
Figure 14 depicts the underlying preferences and the resulting graph $D(N, A(P))$.

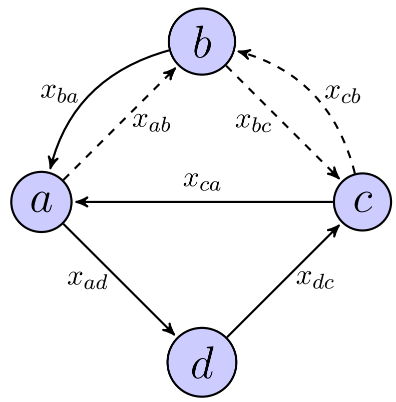

Figure 14: $\succ^P$ and $D(N, A(P))$ of Type 2

Claim 13. *Let $P \in \mathcal{P}$ be of Type 2. Let $x = f(P)$. Then,*
(i). d is filled, i.e., $x_{ad} = x_{dc} = q_d$;
(ii). $x_{ba} = q - q_d$ and $x_{ca} = q_d$;
(iii). $x_{ab} = q - q_d$;
(iv). $x_{bc} = x_{cb} = \min\{q - q_d, q_d\}$.

Proof. Since f is individual-good-preference based, we may assume that the preferences in P are lexicographic. Suppose (i) does not hold. From feasibility and balancedness at d, $x_{ad} = x_{dc} < q_d$. Then, agent a can submit preferences P'_a such that $P' = (P_{-a}, P'_a)$ is of Type 7. Let $x' = f(P')$. By Claim 6, $x'_{ca} = q_d$. Then, from balancedness of x' at a (see also Figure 7) it follows that $x'_{ad} = x'_{ba} + x'_{ca} \ge q_d$. Hence, agent a prefers x'_a to x_a, which contradicts the strategy-proofness of f. Hence, (i) holds.

To prove (ii), we first show that $x_{ca} \ge q_d$. Suppose $x_{ca} < q_d$. Agent c can submit preferences P'_c such that $P' = (P_{-c}, P'_c)$ is of Type 9. Let $x' = f(P')$. By Claim 11, $x'_{ca} \ge q_d$. Hence, c can manipulate f at P, which contradicts the strategy-proofness of f. Hence, $x_{ca} \ge q_d$.

We next show that $x_{ba} \ge q - q_d$. Suppose $x_{ba} < q - q_d$. Agent b can submit preferences P'_b such that $P' = (P_{-b}, P'_b)$ is of Type 3. Let $x' = f(P')$. By Claim 12, $x'_{ba} \ge q - q_d$. Hence, b can manipulate f at P, which contradicts the strategy-proofness of f. Hence, $x_{ba} \ge q - q_d$.

Since, $x_{ba} \ge q - q_d$ and $x_{ca} \ge q_d$, $x_{ba} + x_{ca} \ge q$. But since $x_{ba} + x_{ca} = x_{\cdot a} \le q$, it follows that in fact $x_{ba} = q - q_d$ and $x_{ca} = q_d$ (which proves (ii)).

Then, $x_{ab} = x_{ba} + x_{ca} - x_{ad} = (q - q_d) + q_d - q_d = q - q_d$, which shows (iii). By balancedness, $x_{bc} + x_{dc} = x_{ca} + x_{cb}$. Since $x_{ca} = q_d = x_{dc}$, $x_{bc} = x_{cb}$. Since x is Pareto-efficient, $x_{bc} = x_{cb}$ is maximal, i.e., b or c is filled. This shows that (iv) holds. □

We now complete the proof of Case 2. Let $P \in \mathcal{P}$ be of Type 1 such that b's preferences are additive and

$$1 < \frac{u_b(c)}{u_b(a)} < \frac{\min\{2(q-q_d), q\}}{q-q_d}. \tag{7}$$

It follows from Claim 10 that by stating the truth, agent b's total sum of utility is $(q-q_d)u_b(c)$. Let agent b submit preferences P'_b such that $P' = (P_{-b}, P'_b)$ is of Type 2. In that case she obtains a bundle that consists of $(q-q_d) + \min\{q-q_d, q_d\} = \min\{2(q-q_d), q\}$ units of goods which gives her a total sum of utility of at least $\min\{2(q-q_d), q\}u_b(a)$. From (7) it follows that agent b can manipulate f at P, which contradicts the strategy-proofness of f. Therefore, there is no circulation rule that is both Pareto-efficient and strategy-proof. □

B Proof of Theorem 3

Theorem 3. *For any capacity configuration, the SSTC rule is the unique circulation rule that is top unanimous and self-enforcing group-strategy-proof.*

The SSTC rule is clearly top unanimous, and Proposition 12 shows that it is self-enforcing group-strategy-proof. In order to prove the converse statement, we first establish two lemmas.

Lemma 1. *Let f be strategy-proof. Let $P, P' \in \mathcal{P}$, $i \in N$, and $j, j' \in N\backslash\{i\}$ such that $j, j' \succ_i i$ and $j, j' \succ'_i i$, where $\succ_i \equiv \succ_i^{P_i}$ and $\succ'_i \equiv \succ_i^{P'_i}$. Let $x \equiv f(P)$ and $x' \equiv f(P')$. Let $H = \{h \in N : h \succeq_i j\}$ and $H' = \{h \in N : h \succeq'_i j'\}$ such that $H = H'$. Then $\sum_{h \in H} x_{ih} = m$ implies that $\sum_{h \in H'} x'_{ih} = m$.*

Proof. Suppose that $\sum_{h \in H'} x'_{ih} \neq m$. Assume without loss of generality that $\sum_{h \in H'} x'_{ih} = m' > m$. Let P_i be additive such that for all $h \in H$, $u_i(h) = 1 + \frac{\epsilon}{k_h m}$, where $0 < \epsilon < 1$ and h is k_hth-ranked by $\succ_i$. Let P'_i be additive such that for all $h \in H$, $u'_i(h) = 1 + \frac{\epsilon}{k'_h m}$, where $0 < \epsilon < 1$ and h is k'_hth-ranked by $\succ'_i$. Then $\sum_{h \in H} x_{ih} u_i(h) \leq m\left(1 + \frac{\epsilon}{m}\right) = m + \epsilon$, and $\sum_{h \in H} x'_{ih} u'_i(h) > m' > m + \epsilon$, given that $m' > m$ and $\epsilon < 1$. Then i can manipulate at P via deviation P'_i, which contradicts the strategy-proofness of f. □

Lemma 2. *Let f be individually rational, strategy-proof and nonbossy. Fix $\bar{P} \in \mathcal{P}$ and let $x \equiv f(\bar{P})$. Fix $i \in N$ such that $\succ_i$ ranks j first, where $\succ_i \equiv \succ_i^{\bar{P}_i}$. Assume that $x_{ij} = 0$. Let $\hat{\succ}_i$ be the same as $\succ_i$, except that $\hat{\succ}_i$ does not rank j first, but otherwise the two preference orderings are the same. Let $\hat{P}_i$ be an extension of $\hat{\succ}_i$. Then $f(\bar{P}) = f(\hat{P}_i, P_{-i})$.*

Proof. Let $i \in N, \bar{P} \in \mathcal{P}$, $x \equiv f(\bar{P})$, and let $\succ_i$, $\hat{\succ}_i$, $\hat{P}$ be defined as above. For all $j \in N$, let P_i be a lexicographic extension of $\succ_i$, and assume without loss of generality that $\hat{P}_i$ is a lexicographic extension of $\hat{\succ}_i$. Let $\hat{x} \equiv f(\hat{P})$. Since f is individual-good-preference based, $f(P) = x$. Suppose $x_{ij} > \hat{x}_{ij}$. Then $x_{ij} > 0$ and i can manipulate at $\hat{P}$ via deviation P_i. Now suppose $x_{ij} < \hat{x}_{ij}$. Then $\hat{x}_{ij} > 0$ and by individual rationality $x_{ij} \neq 0$. Hence, $x_{ij} > 0$ and thus i can manipulate at P via deviation $\hat{P}_i$. Therefore, both cases contradict the strategy-proofness of f. Thus, $x_{ij} = \hat{x}_{ij} = 0$. Then Lemma 1 implies that $x_i = \hat{x}_i$ or, equivalently, $f_i(P) = f_i(\hat{P}_i, P_{-i})$. Then, by nonbossiness, $f(P) = f(\hat{P}_i, P_{-i})$. Finally, since $f(P) = f(\bar{P})$, $f(\bar{P}) = f(\hat{P}_i, P_{-i})$. □

Indexed-cTTC rule

In the proof of the converse statement we use an equivalent description of the SSTC rule in addition to its main definition, which is not only useful in the proof but also illuminates the relationship between the SSTC and cTTC rules, our two main competing rules. We refer to the equivalent description of the SSTC rule as the *Indexed-cTTC rule*, because it uses indices for each agent's units which are traded in a similar fashion to trading in the cTTC rule, but with the main difference that only units that have the same index may be traded. More precisely, each agent i's units are indexed from 1 to q_i, corresponding to the respective market segment that each unit is traded in, and the SSTC outcome is obtained by a cTTC-like procedure as follows. In each round, each agent i points to agent j who has the lowest-indexed good among the agents with whom agent i is "willing to" exchange *all* of her remaining common indices, in the sense that for each index that is still held by both i and j, agent j is the highest-ranked by i among all agents who hold this index. Note that there may be multiple agents with whom agent i is "willing to" exchange *all* of the remaining common indices in this sense, all of whom must therefore hold a disjoint set of indices (i.e., no two such agents hold the same index), and hence we specify that i points to the agent who holds the lowest index among all these agents. Note that it is feasible for i to exchange all of her units with a common index with all these agents, and i is going to point to all of them in subsequent rounds.

This procedure is well-defined, since there is always an agent to whom i can point (as long as i still holds any indices), because there must be a favorite agent for i among those agents with whom i has at least one common index left (but note that this agent may be i herself, if all agents with remaining common indices are unacceptable to i in some round). Observe that agent j may not necessarily be the agent who holds the lowest index that i still has for trade. For example, if in some round agent i holds units with indices 1, 2, 3, j holds 2 and 3, k holds 1 and 3, and $j \succ_i k \succ_i i$, then i points to agent j, since i is not willing to trade her index 3 unit with agent k, given that j is preferred by i to k and j still has his index 3 unit. On the other hand, if k holds index 1 only and everything else is unchanged, then i points to k in this round. All trading cycles are carried out simultaneously, just like in the cTTC, but each exchange involves only units of the same index. Observe that in the first round each agent points to her unrestricted top-ranked agent since each agent holds initially a unit indexed 1, which means that each trading cycle formed in the first round will be carried out until some agent's capacity is reached, hence satisfying top unanimity. Finally, note that the equivalence of the SSTC rule and the Indexed-cTTC rule is straightforward to verify, since agents only trade units with the same index, and an agent only points to another agent if this agent is top-ranked by her among all agents who hold a unit with the same index as the one to be traded.

Converse statement: *The only top unanimous and self-enforcing group-strategy-proof rule is the SSTC rule.*

Proof. We prove the statement in five steps. Let g be a top unanimous and self-enforcing group-strategy-proof circulation rule. Then g is strategy-proof and nonbossy by Proposition 11, and individually rational by Proposition 16.

Step 1: *We identify a decomposition of the circulation $g(P)$ into market segments at each preference profile $P \in \mathcal{P}$.*

Fix $P \in \mathcal{P}$. Let $x \equiv g(P)$. Since x is a circulation and thus it is balanced, we can partition the set of goods into market segments where market segments are defined to satisfy the following basic properties: a) each agent has at most one unit of her good in each market segment and b) all one-for-one exchanges of the goods that are required to reach circulation x from the initial endowments take place among goods within a market segment. Note that this definition of market segments is satisfied for the market segments of STC rules. However, unlike for STC rules, in this construction the market segments depend on the circulation, as seen from b).

For each circulation there may be multiple partitions of the set of goods into market segments with properties a) and b), and thus we require that the market segments $\overline{W}_1(P), \dots, \overline{W}_z(P)$ have the following properties:

i) Let the number of the market segments be minimized, that is, combine different goods from different exchanges, as long as no more than one unit of a good is included in a market segment, taking into account that there may be multiple ways of decomposing a circulations into one-for-one exchanges. For example, the exchange of one unit of good involving the same three agents in two different ways can be carried out in two market segments, while the equivalent pairwise trades for the three different pairs of two agents can only be carried out in three market segments. Nonetheless, since these trades may be combined with other trades within a markets segment, exchanges with smaller cycles may lead to fewer market segments overall.

ii) If there are two market segments $\overline{W}_t(P)$ and $\overline{W}_{t'}(P)$ such that $|\overline{W}_t(P)| > |\overline{W}_{t'}(P)|$ then $t < t'$.

iii) There is no unit of any good that could be moved from $\overline{W}_{t'}(P)$ to $\overline{W}_t(P)$, where $t < t'$, subject to properties a) and b).

Step 2: *We show that the market segment decomposition of $g(P)$ is the same at each preference profile $P \in \mathcal{P}$.*

Suppose that the market segment decomposition of g is different for two preference profiles $P, (P'_j, P_{-j}) \in \mathcal{P}$ for some $j \in N$, in the sense that the same market segment decomposition cannot lead to both $g(P)$ and $g(P'_j, P_{-j})$. Let $P' \equiv (P'_j, P_{-j})$. Let $x \equiv g(P)$ and $x' \equiv g(P'_j, P_{-j})$. Then there exist agents $i, l, \hat{l} \in N$ such that

a) i gets a unit of l's good in one market segment, say $\overline{W}_t$, at P, which does not have a unit of $\hat{l}$'s good;

b) i gets a unit of $\hat{l}$'s good in another market segment, say at P, which does not have a unit of l's good;

c) at P' one of these market segments, say $\overline{W}_t$ has both l's and $\hat{l}$'s unit, while $\overline{W}_{\tilde{t}}$ has neither;

d) and (including trade in other market segments) we have $x_{il} + x_{i\hat{l}} > x'_{il} + x'_{i\hat{l}}$.

We can assume that $\succ_i^P$ ranks l first and $\hat{l}$ second (since the setup is symmetric in l and $\hat{l}$) and we can also assume that $\succ_l^P$ and $\succ_{\hat{l}}^P$ rank i first, given that if the market segment decomposition changes at either P or P', due to these preference changes, such that $\overline{W}_t$ and $\overline{W}_{\tilde{t}}$ do not have features a) to d) above, then strategy-proofness is violated for i, l or $\hat{l}$, respectively. Note that $i \neq j$ since i could manipulate otherwise at P' via deviation $P_i = P_j$.

Given that g is top unanimous, agents i and l trade $\min(q_i, q_l)$ of their units between themselves. Moreover, i and l trade in market segment $\overline{W}_t$ and i and $\hat{l}$ trade in market segment $\overline{W}_{\tilde{t}}$ at preference profile P, while only i and l trade in $\overline{W}_t$ and i does not get $\hat{l}$ in $\overline{W}_{\tilde{t}}$ at P'. Since i and $\hat{l}$ trade in $\overline{W}_{\tilde{t}}$ at P, it follows that $q_i > q_l$ and $x_{il} = x'_{il} = q_l$. Furthermore, $x_{i\hat{l}} > x'_{i\hat{l}}$. Let P_i be lexicographic. This implies that $x_i P_i x'_i$.

Let $\bar{\succ}_{\hat{l}}$ rank l first and let $P'_{\hat{l}}$ be an extension of $\bar{\succ}_{\hat{l}}$. Let $\bar{P} \equiv (\bar{P}_{\hat{l}}, P_{-\hat{l}})$ and let $\bar{x} \equiv g(\bar{P})$. Note that the top unanimity of g implies that $x_{il} = q_l$ and thus $\bar{x}_{\hat{l}l} = 0$ by feasibility. Then $g(\bar{P}) = g(P)$, by Lemma 2. Now let $\succ'_i$ rank $\hat{l}$ first and l second, and let P'_i be an extension of $\succ'_i$. Let $\bar{P}'' \equiv (P'_i, \bar{P}_{\hat{l}}, P_{N\setminus\{i,\hat{l}\}})$ and $\bar{x}'' \equiv g(\bar{P}'')$. Then $\bar{x}_{il} + \bar{x}_{i\hat{l}} < x_{il} + x_{i\hat{l}}$.

Observe that $\bar{P}''$ and $\bar{P}$ only differ in agent i's preferences, and since both $\succ_i$ and $\succ'_i$ rank l and $\hat{l}$ as the top two most-preferred goods, Lemma 1 is violated. This is a contradiction, and therefore the market segment decomposition of g is the same at P and at (P'_j, P_{-j}). Since a similar argument holds for any two adjacent profiles which only differ in one agent's preferences, the market segment decomposition of g is the same at each preference profile.

Step 3: *We show that the top unanimity of rule g implies that the fixed market segment decomposition of g is identical to the market segment decomposition of the SSTC rule.*

Order all agents in N from i_1 to i_n in decreasing order of their capacities (i.e., if $k < k'$ then $q_{i_k} \geq q_{i_{k'}}$). For all $t = 1, \ldots, q_{\max}$, let n_t denote the number of agents who have a unit of good in market segment W_t in the SSTC rule. Then $n_1 = n$ and, for all $t = 1, \ldots, q_{\max} - 1$, $n_t \geq n_{t+1}$. For all $t = 1, \ldots, q_{\max}$, let $P^t \in \mathcal{P}$ be a preference profile such that for all $k = 1, \ldots, n_t$ (modulo n_t), $\succ^{P^t}_{i_k}$ ranks i_{k+1} first. Then, given that the market segment decomposition of $g(P)$ is the same at each preference profile $P \in \mathcal{P}$ by Step 2, due to the construction of the market segment decomposition in Step 1, the top unanimity of g implies that for all $t = 1, \ldots, q_{\max}$, the market segment $\overline{W}_t$ consists of one unit of good of each of agents i_1 to i_{n_t}. This means that there are $q_{\max}$ market segments in the fixed market segment decomposition of g and the market segments of g are identical to the market segments $W_1, \ldots, W_{q_{\max}}$ of the SSTC rule.

Step 4: *We prove that, given that rule g is top unanimous and self-enforcing group-strategy-proof and that its fixed market segment decomposition is identical to the market segment decomposition of the SSTC rule, g satisfies strategy-proofness within each market segment.*[35]

Given that at each preference profile each agent receives an item in each market segment $W_1, \ldots, W_{q_{\max}}$ in which the agent has a unit of her good, as shown in Step 3, we can decompose the circulation assigned to each preference profile $P \in \mathcal{P}$ into $q_{\max}$ rules, $g^1, \ldots, g^{q_{\max}}$, such that, for all $P \in \mathcal{P}$, $g^t(P)$ is the assignment of goods within market segment $t = 1, \ldots, q_{\max}$. Given Step 3, what remains to be shown is that for all $t = 1, \ldots, q_{\max}$, g^t is the classical TTC algorithm.

We prove first in this step that for all $t = 1, \ldots, q_{\max}$, g^t is strategy-proof. That is, there is no agent $i \in N$ who can manipulate the outcome of market segment W_t at some profile $P \in \mathcal{P}$. Suppose, to the contrary, that there exist $t \in \{1, \ldots, q_{\max}\}$, $P_{N_t} \in \times_{i \in N_t} \mathcal{P}_i, i \in N_t$ and $P'_i \in \mathcal{P}_i$, where N_t is the set of agents who have a unit of their good in market segment W_t, such that $g^t_i(P'_i, P_{N_t\setminus\{i\}}) \succ^{P_i}_i g^t_i(P_{N-t})$. Note that Step 3 implies that in each round all the exchanges that are carried out take place within market segments, and we can determine which trading cycle

[35]Step 4 is somewhat reminiscent of the decomposition theorem of Le Breton and Sen [12], however their theorem does not apply here because the underlying preferences in our model are identical for each component of the outcome, among other differences.

(or trading cycles) is the first one in the order of trading in the Indexed-cTTC rule that can be manipulated by an agent in the sense that the trading cycle is carried out at P but not at (P'_i, P_{-i}), and we can also determine the market segment within which this trading cycle takes place at P. Assume without loss of generality that the manipulated trading cycle is the first one in the order of trading in the Indexed-cTTC rule, that is, all trades in each previous round of the Indexed-cTTC rule at profile P are such that no agent can manipulate them. Moreover, if multiple trading cycles can be manipulated at P in the first round of the Indexed-cTTC rule in which manipulation is possible then choose one arbitrarily, and if multiple indices are traded then fix the lowest traded index t. Let this trading cycle coalition in market segment W_t be $S = \{j_1, \ldots, j_k\} \subseteq N$ such that for all $v \in \{1, \ldots, k\}$, modulo k, j_v obtains a unit of j_{v+1} in circulation $g^t(P)$.

Let $P' \equiv (P'_i, P_{-i})$. For all $v \in \{1, \ldots, k\}$, modulo k, let $\tilde{P}_{j_v} \in \mathcal{P}_{j_v}$ such that $\succ_j^{\tilde{P}_{j_v}}$ is the same as $\succ_j^{P_{j_v}}$, except that $\succ_j^{\tilde{P}_{j_v}}$ ranks j_{v+1} first. Let $\tilde{P} \equiv (\tilde{P}_S, P'_{-S})$. Given that by assumption all trades at P in the Indexed-cTTC rule prior to the exchange by the trading cycle coalition S are also carried out at P', there exists at least one agent $j_{\hat{v}} \in S$ such that $g_{j_{\hat{v}}} P_{j_{\hat{v}}} g_{j_{\hat{v}}}(P')$ and for all $v \in \{1, \ldots, k\}$, $g_{j_v}(P) R_{j_v} g_{j_v}(P')$.

Now suppose that there exists $j_{\bar{v}} \in S$ such that $f_{j_{\bar{v}}}(P'_{j_{\bar{v}}}, \tilde{P}_{S\setminus\{j_{\bar{v}}\}}, P'_{-S}) P'_{j_{\bar{v}}} f_{j_{\bar{v}}}(\tilde{P}_S, P'_{-S})$. Assuming that $P'_{j_{\bar{v}}}$ is a lexicographic extension of $\succ_{j_{\bar{v}}}^{P^{j_{\bar{v}}}}$, this is not possible, since all trades at P in the Indexed-cTTC rule prior to the exchange by the trading cycle coalition S are also carried out at P' and S is a minimal manipulating coalition at P' via $\tilde{P}_S$. Therefore, coalition S can manipulate circulation rule g at P' in a self-enforcing manner, which is a contradiction since g is self-enforcing group-strategy-proof. Therefore, there is no agent $i \in N$ who can manipulate the outcome of market segment W_t at P, and thus for all $t = 1, \ldots, q_{\max}$, g^t is strategy-proof.

Step 5: *We show that since g satisfies individual rationality, top unanimity and strategy-proofness within each market segment, it selects the TTC outcome in each market segment, and therefore g is the SSTC rule.*

By Step 4, for all $t = 1, \ldots, q_{\max}$, g^t is strategy-proof. Since g satisfies top unanimity, it also follows that for $t = 1, \ldots, q_{\max}$, g^t satisfies top unanimity. We now show that for all $t = 1, \ldots, q_{\max}$, g^t is the classical TTC rule. Fix $t \in \{1, \ldots, q_{\max}\}$ and let $N_t \subseteq N$ denote the set of agents who have a unit of their good in market segment W_t. Suppose, by way of contradiction, that there exists a preference profile $P_{N_t} \in \times_{i \in N_t} \mathcal{P}_i$ such that $g^t(P_{N_t})$ is not the classical TTC outcome at preference profile P_{N_t} . For all $i \in N_t$, let $\succ'_i$ rank i's TTC assignment first and i second, and for all $i \in N_t$, let P'_i be an extension of $\succ'_i$. Let $S \subseteq N_t$ be a trading cycle coalition in the TTC algorithm atP_{N_t} which has at least one member, say $j \in S$, who does not get her TTC assignment at this profile. Assume without loss of generality that S is in the first round of the TTC algorithm at P_{N_t} such that some member of S doesn't get her TTC assignment. That is, agents in all previous rounds of the TTC get their TTC assignments. Then top unanimity implies that S is not a trading cycle coalition in the first round of the TTC algorithm. Suppose S trades in round k of the TTC algorithm, where $k \geq 2$. Since g is strategy-proof and individually rational, and j prefers his TTC assignment to his assignment in market segment W_t at P_{N_t}, given that all agents who trade in previous rounds of the TTC receive their TTC assignment, $g_j(P'_j, P_{N_t\setminus\{j\}}) = j$. Let $l \in N_t$ be the agent who is assigned $j's$ unit in W_t at P_{N_t}. Then l is not assigned j's unit at $g_j(P'_j, P_{N_t\setminus\{j\}})$ and thus, since g is strategy-proof and individually rational, $g_l(P'_j, P'_l, P_{N_t\setminus\{j,l\}}) = l$. Continuing the same argument iteratively, we can show that for all $i \in S$, $g_i(P'_S, P_{N_t\setminus S}) = i$. This contradicts the fact that g satisfies top unanimity. Therefore, $g^t(P_{N_t})$ is the classical TTC outcome at preference profile P_{N_t}. Since the same holds for all $t = 1, \ldots, q_{\max}$ and for all preference profiles $P_{N_t} \in \times_{i \in N_t} \mathcal{P}_i$, this completes the proof that g is the SSTC rule. $\square$

References

[1] T. Andersson, Á. Cseh, L. Ehlers, and A. Erlanson (2021): "Organizing Time Banks: Lessons from Matching Markets," *American Economic Journal: Microeconomics*, 13(1):338–373.

[2] H. Aziz, P. Biró, J. Lang, J. Lesca, and J. Monnot (2015): "Optimal Reallocation under Additive and Ordinal Preferences," *In AAMAS 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore*, 907–913.

[3] S. Barberà, D. Berga, and B. Moreno (2017): "Immunity to Credible Deviations from the Truth," *Mathematical Social Sciences*, 90: 129–140.

[4] P. Biró, F. Klijn, and S. Pápai (2021): "Serial Rules in a Multi-Unit Shapley-Scarf Market," Barcelona School of Economics Working Paper 1255.

[5] K. Cechlárová, 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.

[6] U.M. Dur and M.U. Ünver (2019): "Two-Sided Matching via Balanced Exchange," *Journal of Political Economy*, 127: 1156-1177.

[7] L. Ehlers and B. Klaus (2003): "Coalitional Strategy-Proof and Resource-Monotonic Solutions for Multiple Assignment Problems," *Social Choice and Welfare*, 21(2): 265–280.

[8] E. Fujita, J. Lesca, A. Sonoda, T. Todo, and M. Yokoo (2015): "A Complexity Approach for Core-Selecting Exchange with Multiple Indivisible Goods under Lexicographic Preferences," *In Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 907–913.

[9] J.W. Hatfield (2009): "Strategy-Proof, Efficient, and Nonbossy Quota Allocations," *Social Choice and Welfare*, 33(3): 505–515.

[10] F. Kojima (2013): "Efficient Resource Allocation under Multi-Unit Demand," *Games and Economic Behavior*, 82: 1–14.

[11] H. Konishi, 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): 1–15.

[12] M. Le Breton, A. Sen (1999): "Separable Preferences, Strategyproofness, and Decomposability," *Econometrica*, 67(3): 605–628.

[13] J. Lesca (2015): private communication.

[14] J. Ma (1994): "Strategy-Proofness and the Strict Core in a Market with Indivisibilities," *International Journal of Game Theory*, 23(1): 75–83.

[15] V. Manjunath and A. Westkamp (2021): "Strategy-Proof Exchange under Trichotomous Preferences," *Journal of Economic Theory*, 193, 105197.

[16] E. Miyagawa (2002): "Strategy-Proofness and the Core in House Allocation Problems," *Games and Economic Behavior*, 38(2): 347–361.

[17] S. Pápai (2000): "Strategyproof Assignment by Hierarchical Exchange," *Econometrica*, 68(6): 1403–1433.

[18] S. Pápai (2001): "Strategy-Proof and Nonbossy Multiple Assignments," *Journal of Public Economy Theory*, 3(3): 257–271.

[19] S. Pápai (2003): "Strategyproof Exchange of Indivisible Goods," *Journal of Mathematical Economics*, 39(8): 931–959.

[20] S. Pápai (2007): "Exchange in a General Market with Indivisible Goods," *Journal of Economic Theory*, 132(1): 208–235.

[21] A. Postlewaite (1979): "Manipulation via Endowments," *Review of Economic Studies*, 46(2): 255–262.

[22] M. Pycia (2016): "Swaps on Networks," UCLA, Working paper, February 2016.

[23] A.E. Roth (1982): "Incentive Compatibility in a Market with Indivisible Goods," *Economics Letters*, 9(1): 127–132.

[24] A.E. Roth and A. Postlewaite (1977): "Weak versus Strong Domination in a Market with Indivisible Goods," *Journal of Mathematical Economics*, 4(2): 131–137.

[25] A.E. Roth, T. Sönmez and M.U. Ünver (2004): "Kidney Exchange," *Quarterly Journal of Economics*, 119(2): 457–488.

[26] M. Satterthwaite and H. Sonnenschein (1981): "Strategy-proof Allocation Mechanisms at Differentiable Points," *Review of Economic Studies*, 48(4): 587–597.

[27] A. Schrijver (2003): Combinatorial Optimization: Polyhedra and Efficiency. Algorithms and Combinatorics, Vol. 24. Springer-Verlag, Berlin.

[28] S. Serizawa (2006): "Pairwise Strategy-Proofness and Self-Enforcing Manipulation," *Social Choice and Welfare*, 26(2): 305–331.

[29] L.S. Shapley and H.E. Scarf (1974): "On Cores and Indivisibility," *Journal of Mathematical Economics*, 1(1): 23–37.

[30] T. Sönmez (1999): "Strategy-Proofness and Essentially Single-Valued Cores," *Econometrica*, 67(3): 677–689.

[31] A. Sonoda, E. Fujita, T. Todo, and M. Yokoo (2014): "Two Case Studies for Trading Multiple Indivisible Goods with Indifferences," *In Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 791–797.

[32] L.G. Svensson (1999): "Strategy-proof Allocation of Indivisible Goods," *Social Choice and Welfare*, 16(4): 557–567.

[33] W. Thomson (2016): "Non-bossiness," *Social Choice and Welfare*, 47(3): 665–696.

[34] T. Todo, H. Sun and M. Yokoo (2014): "Strategyproof Exchange with Multiple Private Endowments," *In Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 805–811.

[35] Y. Wang and A. Krishna (2006): "Timeshare Exchange Mechanisms," *Management Science*, 52(8): 1223–1237.