Games and Economic Behavior 136 (2022) 428–453

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Serial Rules in a Multi-Unit Shapley-Scarf Market ☆

Péter Biró [a,b,1], Flip Klijn [c,d,*,2], Szilvia Pápai [e,3]

[a] Institute of Economics, Research Centre for Economic and Regional Studies, Hungarian Academy of Sciences, Hungary
[b] Department of Operations Research and Actuarial Sciences, Corvinus University of Budapest, Hungary
[c] Institute for Economic Analysis (CSIC), Spain
[d] Barcelona School of Economics, Spain
[e] Department of Economics, Concordia University, Canada

ARTICLE INFO

Article history:
Received 17 May 2021
Available online 18 October 2022

JEL classification:
C72
C78
D61

Keywords:
Indivisible goods
Circulation
Shapley-Scarf market
Serial dictatorship
Efficiency

ABSTRACT

We study generalized Shapley-Scarf exchange markets where each agent is endowed with multiple units of an indivisible and agent-specific good and monetary compensations are not possible. An outcome is given by a circulation which consists of a balanced exchange of goods. We focus on circulation rules that only require as input ordinal preference rankings of individual goods, and agents are assumed to have responsive preferences over bundles of goods. We study the properties of serial dictatorship rules which allow agents to choose either a single good or an entire bundle sequentially, according to a fixed ordering of the agents. We also introduce and explore extensions of these serial dictatorship rules that ensure individual rationality. The paper analyzes the normative and incentive properties of these four families of serial dictatorships and also shows that the individually rational extensions can be implemented with efficient graph algorithms.

1. Introduction

Background

Many different situations call for exchanging goods, services, or other items of interest in a centralized manner without using money or prices that facilitate the exchange. In a student exchange program, for instance, universities send exchange

☆ We thank the editor, an associate editor, and two anonymous reviewers for their comments that helped to improve the paper. We also thank seminar and workshop participants at Boston College, Keio University (Third International Workshop on Market Design Technologies for Sustainable Development), Université Paris Dauphine, Corvinus Game Theory Seminar, 10th Matching in Practice workshop (Toulouse), Workshop on Game Theory and Social Choice (Budapest), 21st CTN Workshop (Moscow), 2016 GTMD conference (St. Petersburg), 2016 GAMES Congress (Maastricht), MATCH-UP 2017 Conference (Boston), Lausanne Economic Theory Workshop 2017, 1st Catalan Economic Society Conference, Axiomatizations in Game Theory Workshop (Pécs), and in particular Haris Aziz, Jérôme Lang, Julien Lesca, Rosemarie Nagel, Hans Peters, Madhav Raghavan, Tayfun Sönmez, William Thomson, and Utku Ünver for their comments and suggestions.

* Corresponding author.
E-mail addresses: peter.biro@krtk.mta.hu (P. Biró), flip.klijn@iae.csic.es (F. Klijn), szilvia.papai@concordia.ca (S. Pápai).

[1] P. Biró gratefully acknowledges the financial support by the Hungarian Academy of Sciences, Momentum Grant No. LP2021-2, and by the Hungarian Scientific Research Fund, OTKA, Grant No. K143858.

[2] 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).

[3] 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" (SE-144698).

https://doi.org/10.1016/j.geb.2022.10.006
0899-8256/

students to other universities and receive exchange students from elsewhere in return. Centralized transplant exchanges at a national level or internationally are also important examples, since buying and selling organs is prohibited in most countries. Various websites support the bartering of goods such as clothing and books, or swapping services based on professional skills. Similarly, time banks serve the purpose of exchanging different services in a beneficial manner in a town or neighborhood. The reallocation of shifts in healthcare, swapping sabbatical homes, and timeshare exchanges which allow for trading holiday rights in specific resorts are further examples of exchange without making any payments. Even the financial sector has cases of cyclic liabilities that can be resolved using financial clearing, which means the releasing of financial liabilities in a cyclic manner without further compensation.

Some of the pertinent common features of these exchanges can be captured by the simple circulation model introduced by Biró et al. (2022). Each market in this circulation model is a generalized Shapley-Scarf market (Shapley and Scarf, 1974), where agents are endowed with multiple units of an indivisible and agent-specific good. It is often desirable, if not crucial, to obtain a balanced exchange in these exchange markets, i.e., for each agent the number of units of other goods received equals the number of units of her own good given to other agents. We therefore require that the outcome of a market be given by a circulation, which consists of a balanced exchange of goods. Furthermore, we study circulation rules which take as input preferences over individual goods only. These types of rules are attractive in practical situations where it may be difficult to elicit complex preferences over bundles of goods.

For classical Shapley-Scarf markets, where each agent is endowed with one unit of her good, one exchange rule stands out on the domain of strict preferences: David Gale's celebrated Top Trading Cycles (TTC) rule. Gale's TTC rule is the unique rule that satisfies three key properties, namely individual rationality,[4] Pareto-efficiency,[5] and strategy-proofness[6] in the classical Shapley-Scarf markets (Ma, 1994). Biró et al. (2022) showed that for generalized Shapley-Scarf markets there is no circulation rule that satisfies the combination of individual rationality, a weak version of Pareto-efficiency, and a weak version of strategy-proofness. Given this incompatibility, Biró et al. (2022) explored to what extent two natural generalizations of the TTC rule satisfy these and other properties. In particular, they found that these rules do not have good efficiency properties. One source of inefficiency are the multiple ways in which preferences over individual goods can be extended to preferences over bundles: TTC-based rules do not properly take into account that some agent may prefer bundles that contain her top good together with some rather inferior good while some other agent may prefer, to the contrary, "intermediate" bundles, see Biró et al. (2022, Proposition 3) for more details.

Our approach and motivation

In this paper, we take a different approach. Assuming that there is a natural order of the agents determined by e.g. priority or seniority, an intuitive assignment procedure is to let the agents pick their most preferred goods or bundles following this order, known in general as a serial dictatorship.

Priority rules or serial dictatorships are not only natural and common in many practical situations, but they can be more easily explained and "understood" than TTC-inspired rules; practitioners or participants in these markets may not have an adequate background to properly understand or appreciate the latter. Moreover, as it turns out, serial dictatorship rules tend to have better efficiency properties in the circulation model compared to TTC-based rules.[7]

Pareto efficiency and individual rationality of circulations are more basic requirements than – and incompatible with – strategy-proofness. However, in practice, agents may not be able to report more than an ordinal ranking of the individual goods. Therefore, achieving Pareto efficiency plus individual rationality based on such lean reports raises both an existence and computational complexity issue.[8] In this context we study serial dictatorship rules for the circulation model.

1.1. Illustrative examples

We first illustrate the working of the families of serial dictatorship rules that we study. Consider a market with three agents: $N = \{1, 2, 3\}$. Agent 1 has capacity $q_1 = 1$ and agents 2 and 3 have capacity $q_2 = q_3 = 2$. Initially each agent i is endowed with a "null bundle" that consists of q_i units of her own good. Since goods are agent-specific we will refer to the good of agent i as good i. Each agent is interested in obtaining a bundle of exactly q_i units of possibly different goods in total. Let the agents' preferences over individual goods in this market be given as follows:

$1: 2 \succ_1 1 \succ_1 3$
$2: 3 \succ_2 1 \succ_2 2$
$3: 2 \succ_3 3 \succ_3 1$

Here, for example, agent 3 prefers good 2 to her own good, but finds good 1 unacceptable (i.e., worse than her own good).

[4] A rule is individually rational if at each problem it assigns an acceptable bundle to each agent.
[5] A rule is Pareto-efficient if at each problem the assigned circulation is not Pareto-dominated by some other circulation, i.e., it is not possible to make all agents weakly better off and at least one agent strictly better off.
[6] A rule is strategy-proof if at each problem no agent can obtain a more preferred bundle by misrepresenting her preferences.
[7] A summary of the efficiency properties of the families of serial dictatorship rules studied in this paper is provided in Table 5 in Section 6.1.
[8] For instance, Pareto-efficiency of our so-called Multiple-Serial-IR rules follows by definition; the difficulty is in showing that the rules operate on the underlying profiles of ordinal preferences of individual goods (Theorem 15).

We assume that agents' preferences over bundles are "responsive" to preferences over individual goods: (1) an agent finds a bundle unacceptable if it contains any unacceptable good[9] and (2) preferences are monotonic in the sense that replacing one unit of a good in a bundle by one unit of a more preferred good yields a more preferred bundle. The agents' preferences over bundles are partially determined by responsiveness. For instance, in the case of agent 3, the bundle $(1,1,0)$ that consists of one unit of good 1 and one unit of good 2 is not acceptable (i.e., worse than agent 3's null bundle $(0,0,2)$) because it contains a unit of the unacceptable good 1. Also, agent 3 prefers bundle $(0,2,0)$ to bundle $(0,1,1)$ because the former is obtained from the latter by replacing one unit of a good by a more preferred good. Responsiveness typically does not pin down the complete preference relation over bundles: for example, in the case of agent 2 responsiveness does not tell which of the two acceptable bundles $(2,0,0)$ and $(0,1,1)$ is preferred to the other bundle.

Using this market, we demonstrate the four families of rules that we study in this paper. For each family of rules all goods are first collected and then agents pick single goods/bundles sequentially. More specifically, the common feature of the allocation processes that we explore is that we start with the "empty allocation" and then agents sequentially, following a fixed order, take one good or a complete bundle in turn until their capacities are reached. At the end of the sequential process we obtain a circulation that is not necessarily individually rational. However, if we require the intermediate allocations to be extendable to an individually rational circulation then our final circulation will also be individually rational. First we illustrate the so-called Single-Serial rules where agents select the goods one by one, together with their individually rational counterparts, the Single-Serial-IR rules. Then we demonstrate the so-called Multiple-Serial rules where agents choose complete bundles sequentially, as well as their individually rational counterparts, the Multiple-Serial-IR rules.

Single-Serial rules. This first family of rules lets agents pick a single good at a time following a fixed order. Each agent i appears q_i times in the fixed order. As an illustration, we consider the Single-Serial rule based on the order $\pi = (1,2,3,2,3)$. Following π, at each step an agent picks her most preferred available good, as depicted in Table 1. Thus, the resulting bundles are $x_1 = (0,1,0)$, $x_2 = (0,0,2)$, and $x_3 = (1,1,0)$, for agents 1, 2, and 3, respectively. Note that at step 5 agent 3 was obliged to pick the only remaining (unit of) good 1. Since good 1 is unacceptable to agent 3, her bundle x_3 is unacceptable. Therefore, the circulation $x \equiv (x_i)_{i\in N}$ obtained by the Single-Serial rule is not individually rational.

Table 1
Single-Serial rule.

at step	1	2	3	4	5
agent	1	2	3	2	3
picks good	2	3	2	3	1

Single-Serial-IR rules. The family of Single-Serial-IR rules is obtained by adapting the family of Single-Serial rules to ensure individual rationality. Specifically, at each step an agent picks her most preferred good among the available goods such that this choice is compatible with an individually rational final circulation. Using again the order $\pi = (1,2,3,2,3)$, the goods that are picked are shown in Table 2. Steps 1–3 are the same in Tables 1 and 2. However, at step 4, agent 2 is obliged to pick good 1 to ensure that at the last step agent 3 can pick an acceptable good. Thus, the resulting bundles are $y_1 = (0,1,0)$, $y_2 = (1,0,1)$, and $y_3 = (0,1,1)$, for agents 1, 2, and 3, respectively. By construction, the circulation $y \equiv (y_i)_{i\in N}$ obtained by the Single-Serial-IR rule is individually rational.

Table 2
Single-Serial-IR rule.

at step	1	2	3	4	5
agent	1	2	3	2	3
picks good	2	3	2	1	3

Multiple-Serial rules. The third family of rules lets agents sequentially pick a complete bundle following a fixed order. Hence, each agent appears once in the fixed order. As an illustration, we consider the Multiple-Serial rule based on the order $\bar{\pi} = (2,1,3)$. Following $\bar{\pi}$, at each step an agent picks her most preferred available bundle from the available goods, as depicted in Table 3. Note that at step 3 agent 3 was obliged to pick the remaining goods: one unit of good 1 and one unit of good 2. Since good 1 is unacceptable to agent 3, her bundle is unacceptable. Therefore, the circulation obtained by the Multiple-Serial rule is not individually rational.

Table 3
Multiple-Serial rule.

at step	1	2	3
agent	2	1	3
picks bundle	(0,0,2)	(0,1,0)	(1,1,0)

[9] This assumption and its possible relaxation are discussed in Section 6.3.

Multiple-Serial-IR rules. The last family of rules is obtained by adapting the family of Multiple-Serial rules to ensure individual rationality. Specifically, at each step an agent picks her most preferred bundle from the available goods such that this choice is compatible with an individually rational final circulation. Using again the order $\bar{\pi} = (2, 1, 3)$, the bundles that are picked are shown in Table 4. Step 1 is the same in Tables 3 and 4. However, at step 2 agent 1 is obliged to pick good 1 to ensure that at the last step agent 3 can pick an acceptable bundle. By construction, the circulation obtained by the Multiple-Serial-IR rule is individually rational.

Table 4
Multiple-Serial-IR rule.

at step	1	2	3
agent	2	1	3
picks bundle	(0,0,2)	(1,0,0)	(0,2,0)

1.2. Our main contributions

We first show that if a circulation is Pareto-efficient then it can be obtained by some Single-Serial rule (Proposition 2). Hence, Single-Serial rules are exhaustive in the sense that they yield all Pareto-efficient circulations. However, as the illustrative example shows, Single-Serial rules are not necessarily individually rational. Moreover, we show that when they do yield an individually rational circulation it is only guaranteed to be Pareto-efficient for lexicographic preferences (Lemma 1 and Example 2).

Single-Serial-IR rules are by definition individually rational. Since checking whether an intermediate allocation can still be extended to an individually rational circulation is non-trivial, it is important to show that the Single-Serial-IR rules can be implemented efficiently, i.e., in polynomial time. We prove this by establishing that extendability to an individually rational circulation is equivalent to the existence of a maximum flow in an associated maximum flow problem. More specifically, this equivalence allows us to give an alternative definition of Single-Serial-IR rules (Theorem 7) and then, using the Ford-Fulkerson theorem, show its efficient implementation (Corollary 10). Finally, we prove that Single-Serial-IR rules are Pareto-efficient for lexicographic preferences and any individually rational and Pareto-efficient circulation can be obtained by some Single-Serial-IR rule (Proposition 11).

While Multiple-Serial rules are Pareto-efficient by definition, they are not necessarily individually rational, as demonstrated by the illustrative example. We establish that Multiple-Serial-IR rules do satisfy both properties. The (non-trivial) key issue here is to show that Multiple-Serial-IR rules satisfy our requirement that they only depend on the ordinal preferences over *individual* goods (Theorem 15). Moreover, since Multiple-Serial-IR rules are particular Single-Serial-IR rules, they can be implemented efficiently (Corollary 16).

1.3. Organization of the paper

In Section 2, we introduce the circulation model. In Sections 3 and 4, we present our Single-Serial-(IR) rules and Multiple-Serial-(IR) rules, respectively, and prove our main results on individual rationality, Pareto-efficiency, and computational complexity. In a separate section, Section 5, we show which of our rules also satisfy strategy-proofness or other weaker incentive properties. Section 6 provides a concise summary of the properties satisfied by our rules, discusses generalized serial rules, and describes how our positive results may be extended to more general models. Finally, Section 7 discusses the related literature and applications.

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 finite number of $q_i \in \mathbb{N}$ units of an indivisible, homogeneous, and agent-specific good. We call the non-negative integer q_i agent i's *capacity*. Let $q = (q_i)_{i \in N}$ be the capacity profile. Since goods are agent-specific, for each $i \in N$, we often refer to the good of agent i as good i.

An *assignment* for agent i is a vector $x_i = (x_{ij})_{j \in N} \in \mathbb{N}^N$ with $\sum_{j \in N} x_{ij} \leq q_i$, where x_{ij} denotes the amount (i.e., number of units) of good j that i receives. One particular assignment for agent i is the null assignment 0_i where agent i receives no good, i.e., for each j, $0_{ij} = 0$. An *allocation* is a vector of assignments $x = (x_i)_{i \in N}$ such that for each good $j \in N$, $\sum_{i \in N} x_{ij} \leq q_j$.

A *bundle* for agent i is a vector $x_i = (x_{ij})_{j \in N} \in \mathbb{N}^N$ with $\sum_{j \in N} x_{ij} = q_i$. Clearly, each bundle is an assignment. One particular bundle for agent i is the null bundle e_i where agent i receives no good different from her own, i.e., $e_{ij} = 0$ for all $j \neq i$, or equivalently, $e_{ii} = q_i$. Let X_i denote the set of possible bundles for agent i. A circulation is a vector of bundles such that each agent receives as many goods as she gives away from her initial endowment. Formally, a *circulation* is a vector of bundles $x = (x_i)_{i \in N} \in (X_i)_{i \in N}$ such that for each good $j \in N$, $\sum_{i \in N} x_{ij} = q_j$. Let X denote the set of circulations.

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 receiving (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 $j, l \in N$ 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 to agent i; otherwise it is unacceptable to i.

Each agent i has a linear order P_i on the set of possible bundles X_i. A bundle x_i is acceptable to i if $x_i P_i e_i$ or $x_i = e_i$; it is unacceptable to i otherwise. 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 conditions. Let $x_i, x_i' \in X_i$.

(r1) $e_i P_i x_i$ if there is $j \in N \backslash \{i\}$ with $i \succ_i j$ such that $x_{ij} > 0$; and
(r2) $x_i' P_i x_i$ if there are $j, l \in N$ with $j \succ_i l$ such that $x_{ij}' = x_{ij} + 1$, $x_{il}' = x_{il} - 1$, and $x_{ik}' = x_{ik}$ for all $k \in N \backslash \{j, l\}$.

Condition (r1) is a property of "absolute desirability": it states that agent i finds a bundle unacceptable if it contains some good that is unacceptable to her. Condition (r2) is a monotonicity property: it states that agent i prefers bundle x_i' to 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.

Remark 1. Note that if a bundle only contains acceptable goods to some agent then, by repeated application of (r2), the agent finds the bundle acceptable. Hence, it follows from (r1) and (r2) that a bundle is acceptable if and only if it only contains acceptable goods. $\diamond$

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 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 $P \in \mathcal{P}$, or simply 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 classes of additive and lexicographic preferences. An agent has additive preferences if there is a (cardinal) utility function on the set of acceptable goods such that for any pair of acceptable bundles, the agent prefers the bundle with highest sum of utilities (of the goods in the bundle). We can assume without loss of generality that the utility of her own good equals 0. Formally, agent i's responsive preferences P_i are *additive* if there exists a utility function $u_i : \{j \in N : j \succ_i i\} \to \mathbb{R}_{++}$ such that

$$\text{for all } x_i, x_i' \in X_i \text{ with } x_i, x_i' \, R_i \, e_i, \left[x_i' \, P_i \, x_i \text{ if and only if } \sum_{j: j \succ_i i} x_{ij}' u_i(j) > \sum_{j: j \succ_i i} x_{ij} u_i(j) \right]. \tag{1}$$

An agent has lexicographic preferences if whenever she compares any two acceptable bundles, she prefers the bundle with the largest number of units of her most preferred good; if the two bundles have the same number of units of her most preferred good, then she prefers the bundle with the largest number of units of her second most preferred good; etc. In other words, the agent first maximizes the number of units of her top good, then maximizes the number of units of her second most preferred good, and so on. Therefore, lexicographic preferences are a specific type of additive preferences, i.e., additive preferences that require a particular scheme of "strongly decreasing" utilities. Formally, agent i's responsive preferences P_i are *lexicographic* if there exists a utility function $u_i : \{j \in N : j \succ_i i\} \to \mathbb{R}_{++}$ where

$$\text{for all } k, l \succ_i i, \quad [k \succ_i l \text{ if and only if } u_i(k) > q_i u_i(l)] \tag{2}$$

such that condition (1) holds. Condition (2) says that receiving a unit of the top good is "more important" than receiving any number of other goods, receiving a unit of the second most preferred good is "more important" than receiving any number of the third most preferred or less preferred goods, etc. When preferences are lexicographic, the ordinal ranking over acceptable bundles is completely determined by the ordinal ranking over individual goods.[10] 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.

We require the exchange of the indivisible goods to be balanced. In other words, any outcome of a market should be a circulation. Our aim is to study rules that can be used by a centralized clearinghouse to obtain a circulation for each market. In practice such clearinghouses often only collect 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 a circulation rule.

Fix the set of agents N and the vector of capacities q. A *circulation rule* $f : \mathcal{P} \to X$ specifies a circulation for each preference profile. For each preference profile $P \in \mathcal{P}$, $f_i(P)$ denotes agent i's bundle at P. In view of the discussion above, *we require circulation rules to operate on the underlying profiles of ordinal preferences over individual goods.* In other words, 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,

[10] None of our results requires a similar assumption on the ordinal ranking over unacceptable bundles.

$$\text{for all } P, P' \in \mathcal{P} \text{ with } \succ^P = \succ^{P'}, \quad f(P) = f(P'). \tag{3}$$

In fact, the first two families of rules that we study will be defined directly on the domain of ordinal preferences over individual goods, and hence satisfy (3) by definition.

We first introduce the key desiderata. The first property that we consider indispensable is individual rationality, a standard property which requires that each agent receives a bundle that is acceptable to 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., $x_i R_i e_i$, or equivalently, for each j with $x_{ij} > 0$, $j \succ^{P_i} i$. 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. $\diamond$

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 property where applicable: "necessarily satisfied" and "possibly satisfied," indicating whether the property holds for every responsive extension of the underlying preferences over individual goods, or only for the lexicographic extension that can be inferred from the ordering of individual goods. Thus, "necessarily satisfied" corresponds to the property being satisfied by the entire domain of responsive preferences and is the standard version of the property for responsive preferences over bundles. The "possibly satisfied" version is weaker; namely, it corresponds to the property being satisfied by the lexicographic extension of any preferences over the individual goods. Henceforth, we will denote the weaker version of each property by adding the prefix "ig" (the acronym for "individual good") to the name of the standard version of the property. However, note that by Remark 1, the two versions of individual rationality are equivalent.

In view of the discussion above, we introduce two versions of the other key property, Pareto-efficiency.

Definition 2. A circulation x is Pareto-dominated by another circulation y at $P \in \mathcal{P}$ if for each agent $i \in N$, $y_i R_i x_i$ and for some agent $j \in N$, $y_j P_j x_j$. 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. $\diamond$

Remark 2. Checking whether a circulation is Pareto-efficient for additive preferences is NP-hard (Aziz et al., 2019), where the input is given as the cardinal utilities of agents over the individual goods. On the other hand, checking ig-Pareto-efficiency of a circulation from the agents' ordinal preferences is tractable in polynomial time (Aziz et al., 2015). This suggests that a centralized clearinghouse may find ig-Pareto-efficiency sufficient, especially since the agents may relatively easily detect if the circulation does not satisfy it. However, ensuring Pareto-efficiency (i.e., not "just" ig-Pareto-efficiency) is relevant beyond the detectability argument, as it is an important requirement from the point of view of social welfare. $\diamond$

Proposition 1 in Biró et al. (2022) shows that individual rationality and ig-Pareto-efficiency are not compatible with another important desideratum, ig-strategy-proofness.[11] Given this incompatibility, Biró et al. (2022) focused on two different generalizations of Gale's Top Trading Cycles rule (which does satisfy the three properties in the basic model where each agent has unit capacity). In this paper, we take a different approach by studying classes of serial dictatorships to achieve individual rationality and ig-Pareto-efficiency or Pareto-efficiency. In fact, it is not obvious that there exist rules that satisfy both individual rationality and Pareto-efficiency. The reason is that our requirement that circulation rules operate on the underlying profiles of ordinal preferences over individual goods, i.e., satisfy condition (3), creates tension with Pareto-efficiency on the domain of responsive preferences. We refer to Example 2 in the next section for an illustration of this tension. We postpone the statement (and proof) that individual rationality and Pareto-efficiency are compatible (Corollary 18), as it follows from the result that our class of Multiple-Serial-IR rules satisfies all requirements (Proposition 17).

3. Single-Serial rules

Given a capacity profile q, a q-priority order of agents is an ordered sequence in which each agent i appears exactly q_i times. Formally, let $Q = \sum_{i \in N} q_i$. A *q-priority order of agents* is a vector π in

$$\left\{(i_1, i_2, \ldots, i_Q) : \text{for all } k = 1, \ldots, Q,\ i_k \in N \text{ and } |\{l = 1, \ldots, Q : i_l = i_k\}| = q_{i_k}\right\}.$$

The **Single-Serial rule** associated with a q-priority order π is defined as follows. Fix a preference profile. Following the order π, each agent sequentially chooses her most preferred good among the remaining goods (i.e., goods that have not been exhausted yet). Next we provide a formal definition. For each allocation x and each $j \in N$, let the *remainder* $r_x(j)$ be the number of units of good j that have not been allocated at x, i.e., $r_x(j) = q_j - \sum_{i \in N} x_{ij}$.

[11] A circulation rule is strategy-proof if for each agent it is a weakly dominant strategy to reveal her true preferences. As explained in the discussion on "ig," the (weaker) property ig-strategy-proofness requires that the true preferences are a weakly dominant strategy only for the lexicographic extension of any preferences over the individual goods. We refer to Definition 3 in Section 5 for the formal definition.

INPUT: A q-priority order $\pi = (i_1, i_2, \dots, i_Q)$ and preferences over individual goods $\succ = \succ^P$.

STEP 0: For all $i \in N$, let $x_i^0 = 0_i$ be agent i's null assignment.

STEP $k = 1, \dots, Q$: Let $j^* \in N$ be the good with $r_{x^{k-1}}(j^*) > 0$ such that $j^* \succeq_{i_k} l$ for all $l \in N$ with $r_{x^{k-1}}(l) > 0$. Define x^k by $x^k_{i_k j^*} = x^{k-1}_{i_k j^*} + 1$ and $x^k_{ij} = x^{k-1}_{ij}$ for all $(i, j) \neq (i_k, j^*)$.

OUTPUT: The circulation of the Single-Serial rule associated with π evaluated at profile $\succ$ is x^Q.

Single-Serial rules are well-defined, since they operate on profiles of ordinal preferences over individual goods. However, as the following example illustrates, they need not be individually rational even if preferences are lexicographic.

Example 1. Consider the market (N, q, P) where $N = \{1, 2\}$, $q_1 = q_2 = 1$, and (lexicographic) preferences P given by:

$1: 2 \succ_1 1$
$2: 2 \succ_2 1$

Consider the q-priority order $(1, 2)$. The corresponding Single-Serial rule yields the individually irrational circulation where agent 1 receives her most preferred good and agent 2 receives an unacceptable good. $\diamond$

However, whenever a Single-Serial rule yields an individually rational circulation, it is also Pareto-efficient, provided that preferences are lexicographic. We say that an allocation x is *extendable* to a circulation y if $x \leq y$, i.e., for each agent i and each good j, $x_{ij} \leq y_{ij}$.

Lemma 1. *Let $P \in \mathcal{P}^L$ be a profile of lexicographic preferences. Let x be an individually rational circulation. If some Single-Serial rule yields x at $\succ^P$, then x is Pareto-efficient at P.*

Proof. Suppose $\pi = (i_1, i_2, \dots, i_Q)$ is a q-priority order such that its associated Single-Serial rule yields circulation x and x is not Pareto-efficient at P.

Let y be a circulation that Pareto-dominates x. Since $x \neq y$, there is a *smallest* $k = 1, \dots, Q$ such that x^k (i.e., the allocation at the end of step k of the assignment procedure) is not extendable to y. This implies that at step k, agent i_k chooses a good j^* such that $x^k_{i_k j^*} > y_{i_k j^*}$.

Let $l \in N$ be such that $l \succ_{i_k} j^*$. By the definition of step k, for each $i \in N$, $x^{k-1}_{il} \leq y_{il}$ and $r_{x^{k-1}}(l) = 0$, i.e., $\sum_{i \in N} x^{k-1}_{il} = q_l$. Since y is a circulation, $\sum_{i \in N} y_{il} = q_l$. Hence, $x^{k-1}_{i_k l} = y_{i_k l}$.

We conclude that $x_{i_k j^*} \geq x^k_{i_k j^*} > y_{i_k j^*}$ and for each good $l \in N$ with $l \succ_{i_k} j^*$ we have $x_{i_k l} \geq x^{k-1}_{i_k l} = y_{i_k l}$. But then, since x_{i_k} and y_{i_k} are acceptable bundles to i_k, condition (2) of the definition of lexicographic preferences implies that $x_{i_k} P_{i_k} y_{i_k}$, which contradicts the fact that y Pareto-dominates x. $\square$

The following example shows that the requirement of lexicographic preferences in Lemma 1 cannot be omitted.

Example 2. Consider the market (N, q, P) where $N = \{1, 2, 3, 4, 5, 6\}$, $q_1 = q_2 = q_3 = q_4 = 1$, $q_5 = q_6 = 2$, and responsive[12] preferences P with $(0, 1, 1, 0, 0, 0)$ P_5 $(1, 0, 0, 1, 0, 0)$, $(1, 0, 0, 1, 0, 0)$ P_6 $(0, 1, 1, 0, 0, 0)$ and such that the underlying preferences $\succ_i$ $(i \in N)$ over acceptable individual goods are as follows:

$1: 5 \succ_1 1$
$2: 5 \succ_2 2$
$3: 6 \succ_3 3$
$4: 6 \succ_4 4$
$5: 1 \succ_5 2 \succ_5 3 \succ_5 4 \succ_5 5$
$6: 1 \succ_6 2 \succ_6 3 \succ_6 4 \succ_6 6$

Consider the q-priority order $(5, 6, 6, 5, 1, 2, 3, 4)$. The corresponding Single-Serial rule gives the individually rational circulation x where each agent $i \in \{1, 2, 3, 4\}$ receives one unit of her most preferred good, agent 5 receives bundle $(1, 0, 0, 1, 0, 0)$, and agent 6 receives bundle $(0, 1, 1, 0, 0, 0)$. However, this circulation is not Pareto-efficient, as switching the bundles of agents 5 and 6 is a Pareto improvement. So, x is not Pareto-efficient at P.

The market above also allows us to illustrate why the requirement that circulation rules operate on the underlying profiles of ordinal preferences over individual goods, i.e., satisfy condition (3), creates tension with Pareto-efficiency on the domain of responsive preferences. Consider the market (N, q, P') that is the same as (N, q, P) except that now $(1, 0, 0, 1, 0, 0)$ P'_5 $(0, 1, 1, 0, 0, 0)$ and $(0, 1, 1, 0, 0, 0)$ P'_6 $(1, 0, 0, 1, 0, 0)$. One easily verifies that x is Pareto-efficient at P'.

[12] It is easy to see that there are additive preferences whose underlying preferences over individual goods are $\succ$.

Similarly, there are individually rational circulations that are Pareto-efficient at P, but not at P'. This raises the question whether for any profile of ordinal preferences over individual goods there exists an individually rational allocation that is Pareto-efficient at all possible responsive extensions. Corollary 18 provides an affirmative answer: there exist rules that satisfy both individual rationality and Pareto-efficiency. ◇

The next example shows that the requirement of individual rationality in Lemma 1 cannot be omitted either.

Example 3. Consider the market (N, q, P) where $N = \{1, 2, 3, 4, 5, 6\}$, $q_1 = q_2 = q_3 = q_4 = 1$, $q_5 = q_6 = 2$, and consider lexicographic preferences P such that the underlying preferences $\succ_i$ $(i \in N)$ over individual goods are as follows:[13]

$1: 5 \succ_1 1 \succ_1 \cdots$
$2: 5 \succ_2 2 \succ_2 \cdots$
$3: 6 \succ_3 3 \succ_3 \cdots$
$4: 6 \succ_4 4 \succ_4 \cdots$
$5: 1 \succ_5 2 \succ_5 3 \succ_5 5 \succ_5 4 \succ_5 6$
$6: 1 \succ_6 2 \succ_6 3 \succ_6 4 \succ_6 5 \succ_6 6$

Consider the q-priority order $(1, 2, 3, 4, 5, 6, 6, 5)$. The corresponding Single-Serial rule gives the circulation x where each agent $i \in \{1, 2, 3, 4\}$ receives one unit of her most preferred good, agent 5 receives (unacceptable) bundle $(1, 0, 0, 1, 0, 0)$, and agent 6 receives (acceptable) bundle $(0, 1, 1, 0, 0, 0)$. Clearly, x is not individually rational. Moreover, x is not Pareto-efficient, since switching the bundles of agents 5 and 6 is a Pareto improvement: agent 5 would receive an acceptable bundle and agent 6 would be better off as well by obtaining a unit of her most preferred good. ◇

Our next result shows that Single-Serial rules are "exhaustive" in terms of Pareto-efficiency. More precisely, for any profile of preferences, each Pareto-efficient circulation (whether individually rational or not) can be obtained by applying some Single-Serial rule to the preference profile. This result was proved by Cechlárová et al. (2014) in a similar setting. However, since their result does not imply ours, we present a (simpler) proof for our model.

Proposition 2. *Let $P \in \mathcal{P}$ be a profile of preferences. If a circulation x is Pareto-efficient at P, then x is obtained by some Single-Serial rule applied to $\succ^P$.*

Proof. Let $P \in \mathcal{P}$ be a profile of preferences. Let x be a circulation that is Pareto-efficient at P and suppose, by contradiction, that there is no q-priority order of agents for which the corresponding Single-Serial rule applied to $\succ^P$ yields x.

We first construct a partial q-priority order that results in an allocation as "close" to x as possible. Consider the following procedure. Let y be the empty allocation (where each agent receives her null assignment) and let σ be the empty order. Check whether there are any agent i and good j on the market such that 1) j is i's first choice among the goods that are on the market and 2) by assigning one additional unit of good j to i this extended allocation y would still be extendable to x. If there are an agent i and a good j that satisfy conditions 1) and 2), pick one such pair, say agent i^* and good j^*, and update $y_{i^*j^*} \equiv y_{i^*j^*} + 1$ and $\sigma \equiv (\sigma, i^*)$. If $\sum_{j\in N} y_{i^*j} = q_{i^*}$, then remove agent i^* from the market. Similarly, if $\sum_{i\in N} y_{ij^*} = q_{j^*}$, then remove good j^* from the market. We repeat this incremental procedure until we reach an allocation y that is not extendable with the first choice of any agent. (We reach such an allocation by the assumption that x cannot be obtained with any Single-Serial rule.)

Let k be a good that is on the market. Then $\sum_{i\in N} y_{ik} < q_k = \sum_{i\in N} x_{ik}$. Since for each agent $i \in N$, $y_{ik} \le x_{ik}$, there is some agent $i^* \in N$ with $y_{i^*k} < x_{i^*k}$. In fact, any such i^* is an agent that is still on the market.[14]

We now build a directed graph $D(y)$ on the remaining goods as follows. Let k and j be any two goods that are on the market. If for some agent i^* that is on the market we have 1) $y_{i^*k} < x_{i^*k}$ and 2) j is i^*'s most preferred good among the available goods (hence $y_{i^*j} = x_{i^*j}$), then there is a directed edge from k to j. From the above it follows that for each good k that is still on the market, there is a directed edge going out from k. Therefore, $D(y)$ contains at least one directed cycle, say $(b_1, b_2, \dots, b_r)$. For each directed edge from b_i to b_{i+1}, let a_i be an agent that is on the market such that 1) a_i has strictly more units of good b_i at x than at y and 2) good b_{i+1} is the most preferred good for a_i among all goods that remain on the market (modulo r). Since each agent has strict preferences, it follows that $\{a_1, \dots, a_r\}$ contains r different agents.

Construct the circulation x' from x by carrying out the trades in the cycle. That is, move one unit of good b_{i+1} from agent a_{i+1} to agent a_i (modulo r). Since preferences satisfy condition (r2) of responsiveness, each agent a_i strictly prefers x'_{a_i} to x_{a_i}. Hence, x' Pareto-dominates x, which obviously contradicts the fact that x is Pareto-efficient. □

13 $\cdots$ indicates that preferences can be completed in an arbitrary way.

14 Let p be a removed agent. For each good $r \in N$ (removed or not), $y_{pr} \le x_{pr}$. Since $\sum_{r\in N} y_{pr} = q_p$, it follows that for each good $r \in N$ we have in fact $y_{pr} = x_{pr}$.

3.1. Single-Serial-IR rules

Since Single-Serial rules do not necessarily yield individually rational circulations, we adjust them by demanding that for each (sequential) choice of a good the resulting allocation be IR-extendable. An allocation x is *IR-extendable* if there exists an individually rational circulation x' such that x is extendable to x'. The adjusted serial rules will henceforth be referred to as **Single-Serial-IR rules**.

Input: A q-priority order $\pi = (i_1, i_2, \ldots, i_Q)$ and preferences over individual goods $\succ$.

Step 0: For each $i \in N$, let $x_i^0 = 0_i$ be agent i's null assignment.

Step $k = 1, \ldots, Q$: Let $J \subseteq N$ consist of goods j such that

- $r_{x^{k-1}}(j) > 0$ and
- the allocation z defined by $z_{i_k j} = x_{i_k j}^{k-1} + 1$ and $z_{il} = x_{il}^{k-1}$ for all $(i, l) \neq (i_k, j)$ is IR-extendable.

Let $j^* \in J$ be such that for each $j \in J$, $j^* \succeq_{i_k} j$.
Define x^k by $x_{i_k j^*}^k = x_{i_k j^*}^{k-1} + 1$ and $x_{il}^k = x_{il}^{k-1}$ for all $(i, l) \neq (i_k, j^*)$.

Output: The circulation of the Single-Serial-IR rule associated with π evaluated at profile $\succ$ is x^Q.

The empty allocation (where each agent receives her null assignment) is IR-extendable, because we can give each agent her original endowments. Hence, the algorithm above is well-defined. Also, Single-Serial-IR rules are well-defined, as they operate on profiles of ordinal preferences over individual goods. Moreover, by definition, Single-Serial-IR rules yield individually rational circulations.

Remark 3. Note that if a Single-Serial rule associated with some q-priority order π yields an individually rational circulation x at a preference profile $\succ$, then the Single-Serial-IR rule associated with π also yields circulation x at $\succ$. $\diamond$

Remark 3 coupled with Example 2 shows that Single-Serial-IR rules need not be Pareto-efficient. However, we will establish that Single-Serial-IR rules are ig-Pareto-efficient (Proposition 11). Towards a proof of this result, we will first establish some technical results so that we can provide and use an alternative description (Theorem 7) of the Single-Serial-IR rules. The alternative description also enables us to show that Single-Serial-IR rules can be efficiently implemented from a computational point of view (Corollary 10).

We introduce some additional notation. Let (N, q, P) be a market. Let x be an allocation. Let $|x_i|$ be the number of units of (possibly different) goods that agent i receives at x, i.e., $|x_i| = \sum_{j \in N} |x_{ij}|$, and let $d_x(i) = q_i - |x_i|$ be the *demand* of agent i at allocation x. For $S \subseteq N$, let $d_x(S) = \sum_{i \in S} d_x(i)$ and $r_x(S) = \sum_{i \in S} r_x(i)$. Note that

$$\begin{aligned} r_x(N) &= \sum_{j \in N} r_x(j) = \sum_{j \in N} \left(q_j - \sum_{i \in N} x_{ij} \right) \\ &= \sum_{j \in N} q_j - \sum_{j \in N} \sum_{i \in N} x_{ij} \\ &= \sum_{i \in N} q_i - \sum_{i \in N} \sum_{j \in N} x_{ij} \\ &= \sum_{i \in N} q_i - \sum_{i \in N} |x_i| \\ &= \sum_{i \in N} (q_i - |x_i|) = \sum_{i \in N} d_x(i) = d_x(N). \end{aligned} \tag{4}$$

For $S \subseteq N$, let $I(S)$ denote the goods that are acceptable to some member of S, i.e., $I(S) = \{i \in N : i \succeq_j j \text{ for some } j \in S\}$. We say that S has *overdemand at* x if $d_x(S) > r_x(I(S))$. If some set of agents has overdemand at an allocation x then x is certainly not IR-extendable. Lemma 3 below shows that the reverse of this statement is also true, i.e., IR-extendability is characterized by absence of overdemand.

Lemma 3. *An allocation x is IR-extendable if and only if no set of agents has overdemand at x, i.e., for each $S \subseteq N$, $d_x(S) \leq r_x(I(S))$.*

Before we provide a proof of Lemma 3, we introduce a directed graph that turns out to be a useful tool for the proof of Lemma 3 and for the description of an efficient implementation of Single-Serial-IR rules.

Let (N, q, P) be a market. Let $N = \{1, \ldots, n\}$ and let x be an allocation. Consider the maximum flow problem associated with the directed graph $D_x = (V, A, c)$ with nodes V, arcs A, and arc-capacities c specified as follows. Let $V = \{s\} \cup U \cup W \cup \{t\}$, where $U = \{u_1, u_2, \ldots, u_n\}$ is the set corresponding to the agents and $W = \{w_1, w_2, \ldots, w_n\}$ is the set corresponding to their goods.[15] Regarding the arcs, let $u_i w_j \in A$ if j is an acceptable good to i, for each i let $su_i \in A$ and for each j let $w_j t \in A$. The capacities of the arcs depend on x: let $c(su_i) = d_x(i)$ and $c(w_j t) = r_x(j)$, whilst we have no capacities (or equivalently, capacity ∞) on the arcs between U and W. Let M denote the total remaining demand of the agents (and so the total number of remaining goods), i.e., let $M = d_x(N) = r_x(N)$, where the second equality follows from (4).

A flow ζ from the source s to the sink t in D_x has value of at most M and describes the allocation of the remaining goods as follows. We say that an allocation y (of the remaining goods) corresponds to ζ if $y_{ij} = \zeta(u_i w_j)$ for each pair $(i, j) \in N \times N$. We also say that an allocation x' is the extension of x by y if $x' = x + y$ (i.e., $x'_{ij} = x_{ij} + y_{ij}$ for each pair $(i, j) \in N \times N$). Note that $x' = x + y$ is an individually rational circulation (in which case x is IR-extendable) if and only if all the remaining goods are allocated in y, i.e., if and only if ζ is a maximum flow of value M in D_x.

Proof of Lemma 3. We already noted that the "only if" part is immediate: if some set of agents has overdemand at an allocation x then x is certainly not IR-extendable.

To prove the "if" part, suppose x is not IR-extendable, i.e., the maximum value of a flow in D_x is strictly less than M. We will show that there exists a set $S \subseteq N$ such that $d_x(S) > r_x(I(S))$. The maximum flow – minimum cut theorem (see, e.g., Schrijver, 2003) says that the value of a maximum flow in D_x is always equal to the capacity of a minimum cut in D_x, where a cut is a partition (V_1, V_2) of V such that $s \in V_1$ and $t \in V_2$, and its capacity is the total capacity of the arcs going from V_1 to V_2. Let (V_1, V_2) be a minimum cut of capacity strictly less than M. We complete the proof by showing that $S \equiv V_1 \cap U$ has overdemand.

First note that $S \neq \emptyset$ (otherwise $U \subseteq V_2$ and (V_1, V_2) would have capacity $\geq M$). But then $I(S) \neq \emptyset$ and $I(S) \subseteq V_1$ (otherwise some $j \in I(S)$ would be in V_2 and (V_1, V_2) would have infinite capacity). Since, trivially, $I(S) \subseteq W$, $I(S) \subseteq V_1 \cap W$.

The set of arcs going from V_1 to V_2 are precisely those from s to $U \setminus S$, denoted by A_1, together with those from $V_1 \cap W$ to t, denoted by A_2. This holds because it cannot contain an arc from U to W, otherwise (V_1, V_2) would have infinite capacity. Since $I(S) \subseteq V_1 \cap W$, A_2 contains all the arcs going from $I(S)$ to t, denoted by A_3. Hence, for the capacity $c(A_1 \cup A_2)$ of the cut (V_1, V_2), we have

$$M > c(A_1 \cup A_2) = c(A_1) + c(A_2) \geq c(A_1) + c(A_3) = \sum_{i \in N \setminus S} c(su_i) + \sum_{j \in I(S)} c(w_j t).$$

Using

$$M = d_x(N) = \sum_{i \in S} c(su_i) + \sum_{i \in N \setminus S} c(su_i),$$

we then obtain

$$d_x(S) = \sum_{i \in S} c(su_i) = M - \sum_{i \in N \setminus S} c(su_i) > \sum_{j \in I(S)} c(w_j t) = r_x(I(S)),$$

which completes the proof. □

Remark 4. Lemma 3 is a generalization of Hall's marriage theorem (Hall, 1935). Using additional notation and arguments, the lemma could be obtained from the original Hall's theorem with a graph reduction using as many copies of nodes in the corresponding bipartite graph as the capacities of agents.[16] We provided a straightforward proof of Lemma 3 instead which does not rely on Hall's theorem. Note also that parts of this proof are used in the proofs of Claims 1-5 below. ◇

Let x be an IR-extendable allocation and let $S \subseteq N$. We say that S is constrained at x if the total remaining capacity of the agents in S is equal to the number of units of remaining goods that any member of S finds acceptable. Formally, S is *constrained (at IR-extendable allocation x)* if $d_x(S) = r_x(I(S))$.

Remark 5. Note that $\emptyset$ and N are constrained at each IR-extendable allocation x. ◇

Let x be an IR-extendable allocation. We say that agent i is *unsatisfied at* x if $d_x(i) > 0$, i.e., $\sum_{k \in N} x_{ik} < q_i$. Similarly, we say that good j is *unallocated at* x if $r_x(j) > 0$, i.e., $\sum_{k \in N} x_{kj} < q_j$. Finally, we say that unallocated good j is *feasible to receive* for unsatisfied agent i (at IR-extendable allocation x) if after giving one additional unit of good j to agent i, the resulting allocation x' is still IR-extendable (here x' is the allocation defined by $x'_{ij} = x_{ij} + 1$ and $x'_{kl} = x_{kl}$ for all pairs $(k, l) \neq (i, j)$).

[15] More precisely, u_i corresponds to agent i and w_j corresponds to good j.

[16] We are grateful to an anonymous reviewer for pointing out the connection with Hall's marriage theorem.

Lemma 4. *Let x be an IR-extendable allocation. Let i be an unsatisfied agent and j be an unallocated good at x. Then good j is feasible to receive for agent i at x if and only if there exists no constrained set S with $i \notin S$ and $j \in I(S)$.*

Proof. Let x' be the allocation obtained from x by allocating one more unit of good j to agent i. Let $S \subseteq N$. One easily verifies that if $i \notin S$ and $j \in I(S)$ then

$$r_{x'}(I(S)) = r_x(I(S)) - 1 \text{ and } d_{x'}(S) = d_x(S). \tag{5}$$

Similarly, if $i \in S$ or $j \notin I(S)$ then

$$r_{x'}(I(S)) - d_{x'}(S) \geq r_x(I(S)) - d_x(S) \geq 0, \tag{6}$$

where the last inequality follows from Lemma 3.

By definition, good j is feasible to receive for agent i at x if and only if x' is IR-extendable. Given Lemma 3, this is the case if and only if $d_{x'}(S) \leq r_{x'}(I(S))$ for each $S \subseteq N$. The proof is completed by observing that it follows from (5) and (6) that $d_{x'}(S) \leq r_{x'}(I(S))$ for each $S \subseteq N$ if and only if there is no constrained set S at x (i.e., $d_x(S) = r_x(I(S))$) with $i \notin S$ and $j \in I(S)$. □

Lemma 5. *Let x be an IR-extendable allocation. If $S, T \subseteq N$ are constrained sets at x, then $S \cup T$ and $S \cap T$ are also constrained sets at x.*

Proof. Let x be an IR-extendable allocation. Suppose $S, T \subseteq N$ are constrained sets at x. Then

$$\begin{aligned} d_x(S \cap T) + d_x(S \cup T) &= d_x(S) + d_x(T) \\ &= r_x(I(S)) + r_x(I(T)) \\ &= r_x(I(S) \cap I(T)) + r_x(I(S) \cup I(T)) \\ &\geq r_x(I(S \cap T)) + r_x(I(S \cup T)), \end{aligned} \tag{7}$$

where the first equality follows from the definition of d_x, the second equality from S and T being constrained sets at x, the third equality from the definition of r_x, and the inequality is due to $I(S \cup T) = I(S) \cup I(T)$ and $I(S \cap T) \subseteq I(S) \cap I(T)$. Since x is IR-extendable, it follows from Lemma 3 that no set has overdemand at x. In particular, $d_x(S \cap T) \leq r_x(I(S \cap T))$ and $d_x(S \cup T) \leq r_x(I(S \cup T))$. But then from (7) we have $d_x(S \cap T) = r_x(I(S \cap T))$ and $d_x(S \cup T) = r_x(I(S \cup T))$, i.e., both $S \cap T$ and $S \cup T$ are constrained at x. □

Let x be an IR-extendable allocation. We say that a constrained set S (at x) is *minimal* if it is non-empty and it has no proper non-empty subset $T \subsetneq S$ that is also constrained (at x). Note that, by Lemma 5, if S is a minimal constrained set and $S \cap T \neq \emptyset$ for another constrained set $T \neq S$ then $S \subsetneq T$.

Lemma 6. *Let x be an IR-extendable allocation. If i is an unsatisfied agent in a minimal constrained set, then any unallocated good j that is acceptable to i is feasible for i to receive at x.*

Proof. Let x be an IR-extendable allocation. Suppose for a contradiction that for some unsatisfied agent i in a minimal constrained set S there is an acceptable and unallocated good j that is not feasible for i to receive at x (so, in particular, $i \in S$ and $j \in I(S)$). By Lemma 4 there is a constrained set T such that $i \notin T$ and $j \in I(T)$. Suppose $S \cap T = \emptyset$. Then

$$\begin{aligned} d_x(S \cup T) &= d_x(S) + d_x(T) \\ &= r_x(I(S)) + r_x(I(T)) \\ &> r_x(I(S \cup T)), \end{aligned}$$

where the first equality follows from $S \cap T = \emptyset$, the second equality from the fact that both S and T are constrained, and the inequality from $j \in I(S) \cap I(T)$. Hence, $S \cup T$ has overdemand at x, which contradicts Lemma 3.

Now suppose $S \cap T \neq \emptyset$. Since both S and T are constrained, it follows from Lemma 5 that $S \cap T$ is also constrained at x. Since $i \notin S \cap T$, it follows that $S \cap T$ is a proper non-empty subset of S, which contradicts the minimality of S. □

Lemma 6 allows us to provide an alternative definition of the Single-Serial-IR rules which does not require checking IR-extendability. The circulation obtained by applying the Single-Serial-IR rule associated with a q-priority order to a preference profile can be computed as follows. Initially, each agent's assignment is empty. At each step, find the first entry in the q-priority order, say ℓ, of a member in some minimal constrained subset, say i. Add one unit of agent i's most preferred available good to her assignment. Update the q-priority order by removing entry ℓ, and move to the next step. Next we provide a formal description of the alternative definition.

Alternative definition of the **Single-Serial-IR rules**.

INPUT: A q-priority order π and preferences over individual goods $\succ$.

STEP 0: For each $i \in N$, let $x_i^0 = 0_i$ be agent i's null assignment. Let $\tilde{N} = N$ denote the agents present in the market. Let $\tilde{\pi} = \pi$.

STEP $k = 1, \ldots, Q$: Let i^k be the first agent in $\tilde{\pi}$ that belongs to a minimal constrained subset of $\tilde{N}$ at allocation x^k.[17]

Let $j^* \in N$ be the good with $r_{x^k}(j^*) > 0$ such that $j^* \succeq_{i^k} l$ for all $l \in N$ with $r_{x^k}(l) > 0$. Define x^k by $x^k_{i^k j^*} = x^{k-1}_{i^k j^*} + 1$ and $x^k_{il} = x^{k-1}_{il}$ for all $(i, l) \neq (i^k, j^*)$.
Update $\tilde{\pi}$ by removing[18] the first instance of i^k in $\tilde{\pi}$. If $d_{x^k}(i^k) = 0$, then update $\tilde{N} = \tilde{N} \backslash \{i_k\}$.

OUTPUT: The circulation of the Single-Serial-IR rule associated with π evaluated at $\succ$ is x^Q.

Remark 6. Other alternative definitions are possible. The reason is that it is not necessary to pick the very first agent that belongs to some minimal constrained subset. According to the arguments in the proof of Theorem 7, one could pick some other minimal constrained subset, say S, and take the first agent (according to the q-priority order) that belongs to S. $\diamond$

Theorem 7. *The alternative definition of Single-Serial-IR rules is equivalent to the original definition.*

Proof. To see that the alternative definition of a Single-Serial-IR rule is equivalent to its original definition, it suffices to make the following observations. If at some step of the alternative definition some $S \subsetneq \tilde{N}$ is a minimal constrained set, then agents in $\tilde{N} \backslash S$ can no longer receive any good from $I(S)$, since otherwise the allocation would not remain IR-extendable. Moreover, the agents in S will obviously only choose from $I(S)$, i.e., the goods they find acceptable. So we can treat S independently from $\tilde{N} \backslash S$.

Finally, by rearranging the agents as described, the agents in turn can choose their most preferred goods from the ones that are on the market, since these goods are feasible to receive for them. If the minimal constrained set is $\tilde{N}$ then this is obvious. Moreover, if the minimal constrained set S is a strict subset of $\tilde{N}$, then any agent in S can in turn choose freely her most preferred good among the ones that are still available in $I(S)$, as stated in Lemma 6. $\square$

Corollary 8. *Let $P \in \mathcal{P}$. Let x be a circulation that is obtained by some Single-Serial-IR rule applied to $\succ^P$. Then there is a Single-Serial rule that also yields x at $\succ^P$.*

Proof. Let $(i^1, i^2, \ldots, i^Q)$ be the q-priority order generated by the algorithm of the alternative definition of the Single-Serial-IR rule applied to $\succ^P$. Then the Single-Serial rule associated with q-priority order $(i^1, i^2, \ldots, i^Q)$ applied to $\succ^P$ gives the same circulation. $\square$

Next we show that given an IR-extendable allocation we can decide efficiently whether any given agent is involved in a minimal constrained set. The following lemma shows that it is sufficient to show that for each agent we can determine in polynomial time the *smallest* constrained set that contains the agent.[19]

Lemma 9. *Let x be an IR-extendable allocation. For each $j \in N$, let S_j be the smallest constrained set at x that contains agent j. Then $i \in N$ is in a minimal constrained set at x if and only if for each $j \in N$, $S_i \subseteq S_j$ or $S_i \cap S_j = \emptyset$. Moreover, if i is in a minimal constrained set T at x, then $T = S_i$.*

Proof. Suppose $i \in N$ is in a minimal constrained set, say T, and that for some $j \in N$, neither $S_i \subseteq S_j$ nor $S_i \cap S_j = \emptyset$. Then $S_i \cap S_j$ is a non-empty, strict subset of S_i. By Lemma 5, $S_i \cap S_j$ is constrained. Since S_i is the smallest constrained set that contains i, $S_i \subseteq T$. Hence, $\emptyset \neq S_i \cap S_j \subsetneq S_i \subseteq T$, in contradiction to the assumption that T is a minimal constrained set.

Suppose that for each $j \in N$, $S_i \subseteq S_j$ or $S_i \cap S_j = \emptyset$. We prove that S_i is a minimal constrained set. Suppose S_i is not a minimal constrained set. Then there is a constrained set $T \subseteq N$ with $\emptyset \neq T \subsetneq S_i$. Let $j \in T$. From the first part of the proof $S_j \subseteq S_i$ or $S_j \cap S_i = \emptyset$. Since $j \in S_i$, $S_j \subseteq S_i$. By assumption, $S_i \subseteq S_j$. Thus, $S_i = S_j$. Since S_j is the smallest constrained

[17] Recall that (by our definition) minimality requires non-emptiness. So, $\emptyset$ is not a minimal constrained set. Since N and $N \backslash \tilde{N}$ are constrained at x, $\tilde{N}$ is also constrained at allocation x. Hence, there exists a minimal constrained subset of $\tilde{N}$.

[18] For instance, if at step 1 we have $\tilde{\pi} = (3, 2, 1, 2, 4)$ and $i_1 = 2$, then the updated order is $\tilde{\pi} = (3, 1, 2, 4)$.

[19] Let x be an IR-extendable allocation and $i \in N$. Let $\mathcal{S}$ be the collection of constrained sets at x that contain i. From Remark 5, $N \in \mathcal{S}$. Hence, $\mathcal{S} \neq \emptyset$. From Lemma 5 it follows that $\cap_{S \in \mathcal{S}} S$ is the smallest constrained set at x that contains agent i.

set that contains i, $S_j \subseteq T$. Hence, $S_i = S_j \subseteq T$, which contradicts $T \subsetneq S_i$. Therefore, S_i is a minimal constrained set (which contains i).

The second statement follows immediately from the second part of the proof. $\square$

Finding the smallest constrained set that contains a given agent.

INPUT: A preference profile $\succ$, an IR-extendable allocation x, and an agent $i \in N$.

STEP 1: From the directed graph $D_x = (V, A, c)$ construct the directed graph $D'_x = (V, A, c')$ by only increasing the capacity of arc su_i by one, i.e., $c'(su_i) = c(su_i) + 1$ and $c'(a) = c(a)$ for all arcs $a \neq su_i$.

STEP 2: Run the Ford-Fulkerson algorithm (see, e.g., Schrijver, 2003) to find the smallest minimum cut (V_1^*, V_2^*) of D'_x. In other words, for each minimum cut (V_1, V_2) of D'_x, $V_1^* \subseteq V_1$.

OUTPUT: Let $S = \{j \in N : u_j \in V_1\}$. Set S is the smallest constrained set at x that contains i.

Claims 1–5 below establish that the algorithm above is well-defined and yields the asserted output.

Claim 1. *Each minimum cut (V_1, V_2) of D'_x has capacity $d_x(N)$, i.e., the capacity of each minimum cut of D_x. Moreover, $u_i \in V_1$.*

Proof. Let (V_1, V_2) be a minimum cut of D'_x. Let v' be the capacity of (V_1, V_2). Since x is IR-extendable, it follows from the discussion that precedes the proof of Lemma 3 that $d_x(N)$ is the capacity of each minimum cut of D_x. Denote $v = d_x(N)$.

On the one hand, since $c' \geq c$ we have $v' \geq v$. On the other hand, $v' \leq \sum_{j \in N} c'(w_j t) = \sum_{j \in N} c(w_j t) = r_x(N) = v$. Here the inequality follows from the maximum flow – minimum cut theorem and the fact that the maximum flow in D'_x is at most $\sum_{j \in N} c'(w_j t)$. The last equality follows from (4). Hence, $v' = v = d_x(N)$.

Suppose $u_i \notin V_1$. Suppose (V_1, V_2) is not a minimum cut of D_x. Since $s \in V_1$ and $u_i \notin V_1$, arc (s, u_i) is in (V_1, V_2). Now note that the only difference between D'_x and D_x is the capacity of arc (s, u_i): $c'(su_i) > c(su_i)$. Thus, (V_1, V_2) is not a minimum cut of D'_x either. This contradiction proves that (V_1, V_2) is a min cut of D_x. Since $su_i \in A \cap (V_1 \times V_2)$, $v = \sum_{a \in A \cap (V_1 \times V_2)} c(a) = \sum_{a \in A \cap (V_1 \times V_2)} c'(a) - 1 = v' - 1$, which contradicts $v' = v$. Hence, $u_i \in V_1$. $\square$

Claim 2. *There is a (unique) smallest minimum cut (V_1^*, V_2^*) of D'_x. The Ford-Fulkerson algorithm can be used to find this smallest minimum cut.*

Proof. Let (Y_1, Y_2) and (Z_1, Z_2) (or Y_1 and Z_1 for short) be minimum cuts. Then, for the capacities of the cuts $Y_1 \cap Z_1$ and $Y_1 \cup Z_1$, one easily establishes that $c(Y_1 \cap Z_1) + c(Y_1 \cup Z_1) \leq c(Y_1) + c(Z_1)$. From the minimality of the cuts Y_1 and Z_1, $c(Y_1 \cap Z_1), c(Y_1 \cup Z_1) \geq c(Y_1) = c(Z_1)$. Hence, $c(Y_1 \cap Z_1) = c(Y_1 \cup Z_1) = c(Y_1) = c(Z_1)$. So, $Y_1 \cap Z_1$ is a minimum cut.

Let $\mathcal{C}$ be the collection of minimum cuts Y. From the above it follows that $\cap_{Y \in \mathcal{C}} Y$ is the smallest minimum cut. It is well-known that the Ford-Fulkerson algorithm can be used to find this smallest minimum cut.[20] $\square$

Claim 3. *The set $S \equiv \{j \in N : u_j \in V_1^*\}$ is constrained at x and contains i.*

Proof. From Claim 1, $i \in S$. Note that

$$A \cap \left[(V_1^* \cap U) \times (V_2^* \cap W) \right] = \emptyset; \tag{8}$$

otherwise (V_1^*, V_2^*) would have infinite capacity, contradicting Claim 1. Therefore, the capacity of the cut (V_1^*, V_2^*) of D'_x equals

$$\begin{aligned} c'(V_1^*, V_2^*) &\equiv \sum_{v \in V_2^* \cap U} c'(sv) + \sum_{a \in A \cap [(V_1^* \cap U) \times (V_2^* \cap W)]} c'(a) + \sum_{v \in V_1^* \cap W} c'(vt) \\ &= \sum_{j \in N \setminus S} c'(su_j) + \sum_{w_l \in V_1^* \cap W} c'(w_l t). \end{aligned} \tag{9}$$

Next, define $W(I(S)) \equiv \{w_k : k \in I(S)\}$. We prove that

$$W(I(S)) \subseteq V_1^* \cap W \text{ and, for each } w_k \in (V_1^* \cap W) \setminus W(I(S)),\ r_x(k) = 0. \tag{10}$$

Suppose that for some $k \in I(S)$, $w_k \in V_2^*$. Since $k \in I(S)$, there is $j \in S$ for which good k is acceptable, i.e., $u_j w_k \in A$. Since $j \in S$, $u_j \in V_1^*$. This yields a contradiction to (8). Therefore, the first part of (10) holds.

[20] A proof is available from the authors upon request.

To prove the second part of (10), suppose that for some $w_k \in (V_1^* \cap W)\backslash W(I(S))$ we have $r_x(k) > 0$. Since $w_k \notin W(I(S))$, $k \notin I(S)$. Hence, there is no $j \in N$ with $u_j \in V_1^*$ that finds k acceptable. So,

$$\text{there is no arc in } A \text{ that goes from some node in } V_1^* \text{ to } w_k. \tag{11}$$

Then the capacity of the cut $(V_1^*\backslash\{w_k\}, V_2^* \cup \{w_k\})$ of D'_x equals

$$\sum_{v\in(V_2^*\cup\{w_k\})\cap U} c'(sv) + \sum_{a\in A\cap[(V_1^*\backslash\{w_k\})\cap U]\times[(V_2^*\cup\{w_k\})\cap W]} c'(a) \; + \sum_{v\in(V_1^*\backslash\{w_k\})\cap W} c'(vt) =$$
$$\sum_{j\in N\backslash S} c'(su_j) + \sum_{w_l\in(V_1^*\backslash\{w_k\})\cap W} c'(w_l t) \;=\; c'(V_1^*, V_2^*) - c'(w_k t) \;<\; c'(V_1^*, V_2^*), \tag{12}$$

where the first equality follows from (8) and (11), the second equality from (9), and the inequality from $c(w_k t) = r_x(k) > 0$. Inequality (12) contradicts that (V_1^*, V_2^*) is a minimum cut of D'_x (Claim 2). Hence, the second part of (10) holds. This completes the proof of (10).

Finally, we show that S is constrained at x. The capacity of the cut (V_1^*, V_2^*) of D'_x equals

$$\begin{aligned} d_x(N) &= \sum_{j\in N\backslash S} c'(su_j) + \sum_{w_l\in V_1^*\cap W} c'(w_l t) \\ &= \sum_{j\in N\backslash S} d_x(j) + \sum_{w_l\in V_1^*\cap W} r_x(l) \\ &= \sum_{j\in N\backslash S} d_x(j) + \sum_{w_l\in W(I(S))} r_x(l) \\ &= \sum_{j\in N\backslash S} d_x(j) + \sum_{l\in I(S)} r_x(l) \\ &= [d_x(N) - d_x(S)] + r_x(I(S)), \end{aligned}$$

where the first equality follows from Claim 1 and (9), the second equality from $i \in S$ and the definition of the capacities c' of the arcs, and the third equality from (10). Hence, $d_x(S) = r_x(I(S))$, i.e., S is constrained at x. □

Claim 4. *Let $T \subseteq N$ be a constrained set at x that contains i. Let*

$$\bar{V}_1 \equiv \{s\} \cup \{u_j : j \in T\} \cup \{w_l : l \in I(T)\}$$

and $\bar{V}_2 \equiv V\backslash\bar{V}_1$. Then $(\bar{V}_1, \bar{V}_2)$ is a minimum cut of D'_x.

Proof. First note that by the definition of $(\bar{V}_1, \bar{V}_2)$ and A,

$$A \cap [(\bar{V}_1 \cap U) \times (\bar{V}_2 \cap W)] = A \cap [\{u_j : j \in T\} \times \{w_l : l \notin I(T)\}] = \emptyset. \tag{13}$$

The capacity of the cut $(\bar{V}_1, \bar{V}_2)$ of D'_x equals

$$\begin{aligned} &\sum_{v\in\bar{V}_2\cap U} c'(sv) + \sum_{a\in A\cap[(\bar{V}_1\cap U)\times(\bar{V}_2\cap W)]} c'(a) \; + \sum_{v\in\bar{V}_1\cap W} c'(vt) = \\ &\sum_{j\in N\backslash T} c'(su_j) + \sum_{l\in I(T)} c'(w_l t) = \\ &d_x(N) - d_x(T) + r_x(I(T)) = d_x(N) \end{aligned}$$

where the first equality follows from (13), the second equality from $i \in T$ and the definition of the capacities c' of the arcs, and the third equality from the fact that T is constrained at x. Hence, by Claim 1, $(\bar{V}_1, \bar{V}_2)$ is a minimum cut of D'_x. □

Claim 5. *Set S is the smallest constrained set at x that contains i.*

Proof. From Claim 3, S is a constrained set at x that contains i. Let $T \subseteq N$ be a constrained set at x that contains i. Let $(\bar{V}_1, \bar{V}_2)$ be the minimum cut of D'_x induced by T, as stated in Claim 4. From Claim 2, $V_1^* \subseteq \bar{V}_1$. Then $V_1^* \cap U \subseteq \bar{V}_1 \cap U$. Thus, $S = \{j \in N : u_j \in V_1^*\} \subseteq \{j \in N : u_j \in \bar{V}_1\} = T$. □

Corollary 10 below shows that, from a computational point of view, the alternative definition of Single-Serial-IR rules provides an efficient implementation.[21]

Corollary 10. *Single-Serial-IR rules can be efficiently implemented through their alternative definition. The runtime is bounded by* $O(n^6 Q)$.

Proof. Using the alternative definition (Theorem 7), at each step $k = 1, \ldots, Q$ we have to compute at most n smallest constrained sets to find the first agent i in the remaining q-priority order $\tilde{\pi}$ such that i is in a minimal constrained set. Throughout, the number of edges in the graph is $|E| \leq n^2 + 2n$ and the number of vertices in the graph is $|V| \leq 2n + 2$. Then, since the runtime of Ford-Fulkerson algorithm is bounded by $O(|V| \cdot |E|^2) = O(n^5)$ (see e.g. Schrijver, 2003), the total runtime (of applying at most nQ times the algorithm) is bounded by $O(nQ \cdot n^5) = O(n^6 Q)$. □

Proposition 11. *(i) Let $P \in \mathcal{P}$ be a profile of preferences. Let x be a circulation that is obtained by some Single-Serial-IR rule applied to $\succ^P$. Then x is individually rational at P and if preferences are lexicographic it is also Pareto-efficient at P.*
(ii) Let $P \in \mathcal{P}$ be a profile of preferences. Let x be a circulation that is individually rational and Pareto-efficient at P. Then x is obtained by some Single-Serial-IR rule applied to $\succ^P$.

Proof. We first prove (i). Let f be a Single-Serial-IR rule. By the definition of f, $x = f(\succ^P)$ is individually rational at P. Suppose preferences P are lexicographic. We show that x is also Pareto-efficient at P. By Corollary 8, there is a q-priority order π such that x is obtained by letting the agents sequentially choose their most preferred (and available) goods, following π. In other words, the Single-Serial rule based on π and applied to $\succ^P$ yields x. Then, by Lemma 1, x is Pareto-efficient at P.

Next we prove (ii). Let P be a profile of preferences. Let x be a circulation that is individually rational and Pareto-efficient at P. By Proposition 2, there is a Single-Serial rule associated with some q-priority order π that yields x at $\succ^P$. By Remark 3, the Single-Serial-IR rule associated with π also yields x at $\succ^P$. □

We note that Remark 3 and Example 2 show that requiring lexicographic preferences cannot be omitted from the last part of (i) in Proposition 11.

Corollary 12. *Single-Serial-IR rules are individually rational and ig-Pareto-efficient.*

Remark 7. Biró et al. (2022) showed that the cTTC rule τ (a generalization of the top trading cycles rule) is individually rational and ig-Pareto-efficient. Therefore, by Proposition 11, for each profile of lexicographic preferences P, the circulation $\tau(P)$ can be obtained by some Single-Serial-IR rule. More specifically, the Single-Serial-IR rule that is based on the q-priority order generated by the top trading cycles during the execution of the generalized top trading cycles algorithm, where the order among the agents in the same top trading cycle can be arbitrary, yields the circulation $\tau(P)$. ◇

4. Multiple-Serial rules

The **Multiple-Serial rule** associated with an order of the agents is defined as follows. Fix a preference profile. Following the order, each agent sequentially chooses her most preferred bundle among the remaining goods (i.e., goods that have not been exhausted yet). Next we provide a formal definition.

Input: An order $\pi = (i_1, \ldots, i_n)$ of the agents and a preference profile $P \in \mathcal{P}$.

Step 0: For each $i \in N$, let $x_i^0 = 0_i$ be agent i's null assignment.

Step $k = 1, \ldots, n$: Let Y_{i_k} denote the collection of available bundles for agent i_k, i.e.,

$$Y_{i_k} = \left\{ x_{i_k} \in X_{i_k} : \text{ for each good } j \in N,\ x_{i_k j} \leq r_{x^{k-1}}(j) \right\}.$$

Let $y^*_{i_k} \in Y_{i_k}$ be the bundle such that for each $y_{i_k} \in Y_{i_k}$, $y^*_{i_k} R_{i_k} y_{i_k}$. Define x^k by setting $x^k_{i_k} = y^*_{i_k}$ and $x^k_i = x^{k-1}_i$ for all $i \neq i_k$.

Output: The circulation of the Multiple-Serial rule associated with π evaluated at profile P is x^n.

An important observation is that Multiple-Serial rules operate on the underlying profiles of ordinal preferences over individual goods, i.e., (3) is satisfied. This observation follows from responsiveness: an agent's most preferred bundle from the

[21] We note that deciding whether an allocation x is IR-extendable is also possible with a more general strongly polynomial algorithm (i.e., the running time of the algorithm is polynomial in the number of agents and does not depend on the number of goods or the characteristic of the allocation concerned). This follows from the solvability of the classical circulation problem with lower arc-capacities (see, e.g., Schrijver, 2003). Using our terminology, the latter problem focuses on the question whether there exists an individually rational circulation given the acceptability graph, the capacities of the nodes, and the lower arc-capacities (which equal the flow through the arcs as stipulated by the allocation x).

remaining goods coincides with the bundle obtained from a greedy procedure where the agent picks the most preferred (and available) goods one by one. Therefore, when analyzing the computational complexity of Multiple-Serial rules, we can assume that the input of the algorithm is $\succ^P$, rather than P.

Obviously, Multiple-Serial rules need not be individually rational: the Single-Serial rule that yields an individually irrational circulation in Example 1 is a Multiple-Serial rule, since all agents' capacities are equal to one. The following result is immediate.

Proposition 13. *Multiple-Serial rules are Pareto-efficient.*

Next we introduce and study Multiple-Serial-IR rules which are the rules obtained by adjusting the Multiple-Serial rules to guarantee individual rationality while maintaining Pareto-efficiency.

4.1. Multiple-Serial-IR rules

Since Multiple-Serial rules do not necessarily yield individually rational circulations, we adjust them by demanding that for each (sequential) choice of a bundle the resulting allocation be IR-extendable. The adjusted serial rules will henceforth be referred to as **Multiple-Serial-IR rules**.

INPUT: An order $\pi = (i_1, \dots, i_n)$ of the agents and a preference profile $P \in \mathcal{P}$.

STEP 0: For each $i \in N$, let $x_i^0 = 0_i$ be agent i's null assignment.

STEP $k = 1, \dots, n$: Let Y_{i_k} denote the collection of bundles $y_{i_k} \in X_{i_k}$ for agent i_k such that

- for each good $j \in N$, $y_{i_k j} \leq r_{x^{k-1}}(j)$ and
- the allocation z defined by $z_{i_k} = y_{i_k}$ and $z_i = x_i^{k-1}$ for all $i \neq i_k$ is IR-extendable.

Let $y_{i_k}^* \in Y_{i_k}$ be the bundle such that for each $y_{i_k} \in Y_{i_k}$, $y_{i_k}^* R_{i_k} y_{i_k}$. Define x^k by setting $x_{i_k}^k = y_{i_k}^*$ and $x_i^k = x_i^{k-1}$ for all $i \neq i_k$.

OUTPUT: The circulation of the Multiple-Serial-IR rule associated with π evaluated at profile P is x^n.

Remark 8. If a Multiple-Serial rule associated with some order π yields an individually rational circulation x at a preference profile P, then the Multiple-Serial-IR rule associated with π also yields circulation x at P. $\diamond$

Next we show that Multiple-Serial-IR rules operate on the underlying profiles of ordinal preferences over individual goods, i.e., (3) is satisfied. This will allow us to assume that the input of the algorithm above is $\succ^P$, rather than P, and show that Multiple-Serial-IR rules can be efficiently implemented for such concise inputs.

We first prove a technical lemma. Given an IR-extendable allocation x, we say that an assignment y_i is *feasible to receive* for agent i (at x) if after giving the goods in y_i to agent i (on top of those in x_i), the resulting allocation x' is IR-extendable (here x' is the allocation defined by $x'_{ij} = x_{ij} + y_{ij}$ for all j and $x'_l = x_l$ for all $l \neq i$).

Lemma 14. *Let $i \in N$. Let x be an IR-extendable allocation and let y_i and z_i be two assignments that are both feasible to receive for agent i at x and such that $|y_i| < |z_i|$. Then there is a good j in z_i such that after adding j to y_i the extended assignment is also feasible to receive for agent i at x.*

Proof. Let $i \in N$. First we prove the lemma for $|y_i| = 1$ and $|z_i| = 2$. Suppose that the statement is not true. Let y_i consist of good j and let z_i consist of two goods, k and l, with the possibility that $k = l$. Let x' denote the extension of x by y_i and let x'' denote the extension of x by z_i. Note that x' and x'' are both IR-extendable but, by our assumption, at x' neither k nor l is feasible to receive for i. Therefore, by Lemma 4, there is a constrained set S at x' such that $i \notin S$ and $k \in I(S)$, and similarly, there is a constrained set T (which possibly coincides with S) at x' such that $i \notin T$ and $l \in I(T)$. In particular, $k, l \in I(S) \cup I(T) = I(S \cup T)$.

Note that x' only differs from x by adding a unit of good j to x_i. Since $i \notin S$, it follows that for each $s \in S$, $d_x(s) = d_{x'}(s)$. Hence, $d_x(S) = d_{x'}(S)$. Suppose $j \notin I(S)$. Then, for each $s \in I(S)$, $r_x(s) = r_{x'}(s)$. Hence, $r_x(I(S)) = r_{x'}(I(S))$. Since k is feasible to receive for i at x, it follows from Lemma 4 that S is not constrained at x. Thus, $d_x(S) \neq r_x(I(S))$. But then it follows from the above that $d_{x'}(S) \neq r_{x'}(I(S))$ as well, which contradicts that S is constrained at x'. Hence, $j \in I(S)$. Similarly, $j \in I(T)$. Hence, $j \in I(S) \cup I(T) = I(S \cup T)$.

We now show that $S \cup T$ has overdemand at x''. First, since $i \notin S \cup T$, it follows that for all $p \in S \cup T$ we have $x'_p = x_p = x''_p$. Thus, $d_{x'}(S \cup T) = d_{x''}(S \cup T)$. Second, by the definitions of x' and x'' and the fact that $j, k, l \in I(S \cup T)$, it follows that $r_x(I(S \cup T)) = r_{x'}(I(S \cup T)) + 1$ and $r_{x''}(I(S \cup T)) = r_x(I(S \cup T)) - 2$. Hence, $r_{x''}(I(S \cup T)) = r_{x'}(I(S \cup T)) - 1$. By Lemma 5, $S \cup T$ is a constrained set at x'. Therefore,

$$d_{x''}(S \cup T) = d_{x'}(S \cup T) = r_{x'}(I(S \cup T)) = r_{x''}(I(S \cup T)) + 1,$$

which shows that $S \cup T$ has overdemand at x''. Using Lemma 3 we obtain a contradiction with the fact that x'' is IR-extendable. Hence, the lemma holds when $|y_i| = 1$ and $|z_i| = 2$.

Now we complete the proof of the lemma by extending the previous argument and by using the above subcase. Suppose the lemma is not true. Among all triples (x, y_i, z_i) that violate the statement pick one for which $|y_i|$ is minimal. Note $|y_i| > 0$.

Next, note that y_i and z_i do not have any good in common. Otherwise, this good could be added to x_i and omitted from both y_i and z_i, resulting in another triple, say (x', y_i', z_i'), which violates the statement while $|y_i'| < |y_i|$, contradicting the minimality of $|y_i|$.

Let k be some good in y_i. Let y_i' be the assignment that results from removing good k from y_i. (Then obviously y_i' is feasible to receive for agent i at x.) If there is still no good from z_i that can be added to y_i' while keeping the thus extended assignment feasible to receive for i at x, then the triple (x, y_i', z_i) violates the statement while $|y_i'| < |y_i|$, which contradicts the minimality of $|y_i|$. Hence, there is a good in z_i, say j, that can be added to y_i' such that the thus extended assignment (say y_i'') would still be feasible to receive for i at x.

Extend x by assigning j to i and let x' be the resulting allocation, i.e., the only difference between x and x' is that $x'_{ij} = x_{ij} + 1$. Let z_i' be the assignment obtained from z_i by removing good j. We will show that the triple (x', y_i', z_i') also violates the statement and, since $|y_i'| < |y_i|$, we obtain a contradiction with the minimality of $|y_i|$. First, by definition of x' and good j, assignments y_i' and z_i' are both feasible to receive for agent i at x'. Moreover, since $|y_i| < |z_i|$, we also have $|y_i'| < |z_i'|$. Finally, there is no good l in z_i' such that after adding l to y_i' the extended assignment is also feasible to receive for agent i at x'. To show the last claim, suppose this is not the case, i.e., ($\star$) there *does* exist a good l in z_i' such that after adding l to y_i' the extended assignment is also feasible to receive for agent i at x'. Then consider the triple (x^*, y_i^*, z_i^*) where x^* is obtained from x by adding y_i' to x_i and where y_i^* is the assignment that consists of good k and z_i^* is the assignment that consists of goods j and l.

We verify that (x^*, y_i^*, z_i^*) violates the statement of the lemma. First, since y_i' is feasible to receive for agent i at x, allocation x^* is IR-extendable. Second, $|y_i^*| = 1 < 2 = |z_i^*|$. Third, since y_i is feasible to receive for agent i at x, it follows that y_i^* is feasible to receive for agent i at x^*. Fourth, by ($\star$), z_i^* is feasible to receive for agent i at x^*. Finally, since (x, y_i, z_i) violates the statement, it follows that we cannot add either good j or good l to y_i^* such that the extended assignment is feasible to receive for agent i at x^*. However, given that $|y_i^*| = 1$ and $|z_i^*| = 2$, it follows from the first part of the proof that (x^*, y_i^*, z_i^*) does not violate the statement of the lemma. This contradiction completes the proof. □

Theorem 15. *Each Multiple-Serial-IR rule operates on the underlying profiles of ordinal preferences over individual goods, i.e., (3) is satisfied. More precisely, the bundle of each agent can also be obtained in a greedy way by selecting (when it is her turn) one by one the most preferred goods from the goods that are feasible to receive for the agent.*

Proof. First note that checking IR-extendability only requires the ordinal preferences over individual goods. Hence, to determine whether a good is feasible to receive also only requires the ordinal preferences over individual goods.

Let $i \in N$. Let the greedy method yield bundle g_i for agent i. Suppose that the Multiple-Serial-IR rule yields a different bundle, say f_i.

Let us order the goods in both g_i and f_i according to agent i's ordinal preferences over individual goods. More specifically, for each $k \in \{1, \ldots, q_i\}$, let $f_i(k)$ and $g_i(k)$ be the k-th most preferred good in f_i and g_i, respectively. (Note that some good may appear multiple times in f_i and/or g_i. Therefore it is possible that for some $k \in \{1, \ldots, q_i\}$ we have $f_i(k) = f_i(k+1)$ and/or $g_i(k) = g_i(k+1)$.) Suppose i weakly prefers $g_i(k)$ to $f_i(k)$ for each $k \in \{1, \ldots, q_i\}$. Then, by responsiveness, g_i is weakly preferred to f_i. Since $g_i \neq f_i$, it follows that g_i is strictly preferred to f_i, which contradicts the optimality of agent i's choice in the Multiple-Serial-IR rule.

Therefore, there is some index $k \in \{1, \ldots, q_i\}$ such that i strictly prefers $f_i(k)$ to $g_i(k)$, and thus i also strictly prefers each of the goods $f_i(l)$ with $1 \le l \le k$ to $g_i(k)$. Let z_i be the assignment that consists of the k most preferred goods in f_i and let y_i be the assignment that consists of the $k-1$ most preferred goods in g_i (here multiple units of the same good are also counted). By Lemma 14, there is a good j in z_i such that after adding j to y_i the extended assignment is also feasible to receive for agent i, which contradicts the selection of the greedy method. □

Remark 9. Theorem 15 can be proved alternatively using matroids as follows.[22] Let $i \in N$. Let x be an IR-extendable allocation. The collection $\mathcal{M}$ of sets of up to k goods ($k \le r_x(i)$) that are feasible to receive for i at x subject to IR-extendability is a matroid. To see this, note that the exchangeability property of matroids is precisely the contents of Lemma 14 (the other matroid properties are satisfied trivially). Thus, by applying Theorem 1 in Gourvès (2019) it follows that for any responsive preferences P_i, the most preferred bundle of k goods in $\mathcal{M}$ can be obtained by choosing k goods greedily according to $\succ^{P_i}$, which shows Theorem 15. ◇

[22] We are grateful to an anonymous reviewer for pointing out the alternative approach with matroids.

Corollary 16. *Each Multiple-Serial-IR rule is a Single-Serial-IR rule. In particular, it can be efficiently implemented.*

Proof. Consider any Multiple-Serial-IR rule. Let $\pi = (i_1, \dots, i_n)$ be the associated order of the agents. Let $\bar{\pi}$ be the q-priority order in which the first q_1 entries are agent i_1, the next q_2 entries are agent i_2, etc. According to Theorem 15, the Multiple-Serial-IR rule (associated with π) coincides with the Single-Serial-IR rule associated with $\bar{\pi}$. Then efficient implementation follows from Corollary 10. □

The following proposition follows easily from the definition of the Multiple-Serial-IR rules.

Proposition 17. *Multiple-Serial-IR rules are individually rational and Pareto-efficient.*

Proof. Individual rationality is immediate. Suppose a Multiple-Serial-IR rule associated with order $\pi = (i_1, \dots, i_n)$ is not Pareto-efficient. Then there is a preference profile $P \in \mathcal{P}$ such that the rule applied to P gives a circulation x that is Pareto-dominated by some circulation x'. Following the order π, consider the first agent i_k such that $x_{i_k} \neq x'_{i_k}$. Since $x'_{i_k} P_{i_k} x_{i_k}$, it follows from Step k of the definition of the Multiple-Serial-IR rule that x' cannot be individually rational. So, there is an agent $i \in N$ for which $e_i P_i x'_i$. Since x is individually rational, $x_i R_i e_i$. Hence, $x_i P_i x'_i$ which contradicts the fact that x' Pareto-dominates x. □

Proposition 17 also shows that individual rationality and Pareto-efficiency together are compatible with our requirement that circulation rules operate on the underlying profiles of ordinal preferences over individual goods.

Corollary 18. *There are individually rational and Pareto-efficient rules.*

A converse to Proposition 17 does not hold. More precisely, there are markets where some individually rational and Pareto-efficient circulation cannot be obtained with any Multiple-Serial-IR rule (and hence, by Remark 8, also not with any Multiple-Serial rule; thus a converse statement to Proposition 13 does not hold either). We demonstrate this with the following example.

Example 4. Consider the market (N, q, P) where $N = \{1, 2, 3\}$, $q_1 = q_2 = q_3 = 2$, and consider preferences P such that the underlying preferences $\succ_i$ $(i \in N)$ over acceptable individual goods are as follows:

$1: 3 \succ_1 1$
$2: 3 \succ_2 2$
$3: 1 \succ_3 2 \succ_3 3$

The three circulations

$x: x_{13} = x_{22} = x_{31} = 2,$
$x': x'_{11} = x'_{23} = x'_{32} = 2,$ and
$x'': x''_{11} = x''_{13} = x''_{22} = x''_{23} = x''_{31} = x''_{32} = 1$

are individually rational and Pareto-efficient independently of the particular responsive preferences P_3 of agent 3 over bundles (in particular, we can assume that all preferences are lexicographic).[23] However, the only (individually rational and Pareto-efficient) circulations that can be obtained by Multiple-Serial-IR rules are x and x'. Specifically, orders (1,2,3), (1,3,2), (3,1,2), and (3,2,1) yield x, while orders (2,1,3) and (2,3,1) yield x'. Thus, the individually rational and Pareto-efficient circulation x'' cannot be obtained by any Multiple-Serial-IR rule.

Given Remark 8, it is clear that the class of Multiple-Serial rules can only yield a subset of the circulations obtained by the Multiple-Serial-IR rules, in addition to possibly some individually *ir*rational circulations. Specifically in this example, Multiple-Serial rules lead to the following: orders (1,2,3), (1,3,2), and (3,1,2) yield x, order (2,1,3) yields x', and orders (2,3,1) and (3,2,1) yield the individually irrational circulation y given by $y_{12} = y_{23} = y_{31} = 2$. Therefore, the individually rational and Pareto-efficient circulation x'' cannot be obtained by any Multiple-Serial rule either. ◇

Example 4 together with Proposition 11 demonstrate an interesting difference between Multiple-Serial-IR and Single-Serial-IR rules: given any profile of lexicographic preferences P and any circulation x that is individually rational and Pareto-efficient at P, x can be obtained by some Single-Serial-IR rule, but possibly not by any Multiple-Serial-IR rule.

When we compare the Single-Serial rules with the Multiple-Serial rules (with or without individual rationality), the Multiple-Serial rules achieve Pareto-efficiency even for responsive preferences, but the price we pay is that not all Pareto-efficient circulations can be obtained as illustrated by Example 4. In particular, the circulations obtained by the Multiple-Serial rules tend to be rather unfair, since the agents who choose first can obtain the goods that possibly all agents prefer

[23] Given $\succ_1$ and $\succ_2$, the responsive preferences P_1 and P_2 of agents 1 and 2 are uniquely determined. For agent 3, $\succ_3$ together with responsiveness does not specify whether receiving two units of good 2 is preferred to one unit of good 1 together with one unit of good 3.

most. More "equitable" (but still Pareto-efficient) circulations in which several agents receive (possibly commonly) most preferred goods are typically not obtained through Multiple-Serial rules. This indicates that there is a trade-off between the rule being Pareto-efficient for responsive preferences and being "equitable."

5. Manipulability

In this section, we determine which of our rules satisfy incentive properties. As a starting point, we note that Proposition 1 in Biró et al. (2022) shows that individual rationality and ig-Pareto-efficiency are not compatible with another important desideratum, ig-strategy-proofness. For any $P \in \mathcal{P}$ and any $i \in N$, denote $P_{-i} = (P_j)_{j \neq i}$.

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 $P \in \mathcal{P}^L$.[24] $\diamond$

The following example illustrates the incompatibility of individual rationality, ig-Pareto-efficiency, and ig-strategy-proofness. Specifically, Serial-IR rules satisfy the first two properties, and are hence vulnerable to manipulations.

Example 5. Consider the market (N, q, P) where $N = \{1, 2, 3\}$, $q_1 = q_2 = q_3 = 1$, and lexicographic preferences P such that the preferences over acceptable goods are given by

$1: 3 \succ_1 1$
$2: 3 \succ_2 1 \succ_2 2$
$3: 2 \succ_3 3$

Consider the Single-Serial-IR rule associated with the order $(1, 2, 3)$. This rule yields circulation x where $x_{13} = x_{21} = x_{32} = 1$. However, if agent 2 removes good 1 from her list of acceptable goods then the rule yields circulation x' where $x'_{11} = x'_{23} = x'_{32} = 1$. Obviously, agent 2 prefers x'_2 to x_2. $\diamond$

The literature considered several weaker incentive properties, which we explore next. For each agent $i \in N$, let $\mathcal{L}_i$ denote the set of strict (ordinal) preferences over individual goods for agent i. A truncation of a preference list over individual goods is a preference list obtained by making some of the lowest-ranked acceptable goods unacceptable. Formally, a preference $\succ'_i \in \mathcal{L}_i$ is a *truncation* of $\succ_i \in \mathcal{L}_i$ if for all $k, l \in N$ we have [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 $l \succ'_i i$]. The first condition says that if two goods are listed as acceptable under the "manipulation" $\succ'_i$, then they are ordered in the same way as in the true preferences $\succ_i$. The second condition says that if a good is listed as acceptable under the "manipulation" $\succ'_i$ and there is some other good that is more preferred in the true preferences $\succ_i$, then the latter good is also acceptable under the "manipulation" $\succ'_i$.

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.[25] A circulation rule f is *ig-truncation-proof* if no agent can manipulate f by means of truncation at any profile $P \in \mathcal{P}^L$ of lexicographic preferences.[26] $\diamond$

A preference $\succ'_i \in \mathcal{L}_i$ is a *dropping* of $\succ_i \in \mathcal{L}_i$ if for all $k, l \in N$, [if $k \succeq'_i l \succeq'_i i$, then $k \succeq_i l \succeq_i i$]. Obviously, since the requirement in the definition of dropping is exactly the first condition in the definition of truncation, it follows that each truncation is a dropping.

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.[27] $\diamond$

[24] Since circulation rules operate on profiles of ordinal preferences over individual goods, equivalent definitions of strategy-proofness and ig-strategy-proofness are obtained by demanding that the deviation P'_i is lexicographic.
[25] Kojima (2013) similarly defined "non-manipulability via truncation" in the context of resource allocation with multi-unit demand.
[26] Since circulation rules operate on profiles of ordinal preferences over individual goods, equivalent definitions of truncation-proofness and ig-truncation-proofness are obtained by requiring that the deviation P'_i is lexicographic.
[27] Since circulation rules operate on profiles of ordinal preferences over individual goods, equivalent definitions of dropping-proofness and ig-dropping-proofness are obtained by requiring that the deviation P'_i is lexicographic.

Note that strategy-proofness implies dropping-proofness, which in turn implies truncation-proofness. Similarly, ig-strategy-proofness implies ig-dropping-proofness, which in turn implies ig-truncation-proofness.

Since preferences are lexicographic and agent 2's manipulation is a truncation, Example 5 provides an instance of a Single-Serial-IR rule that is not ig-truncation-proof. Hence, there are Single-Serial-IR rules that do not satisfy any of the six incentive properties! Note that since in Example 5 each agent's capacity equals one, the Single-Serial-IR rule is a Multiple-Serial-IR rule. Thus, there are Multiple-Serial-IR rules that do not satisfy any of the six incentive properties.

However, as a positive result we show that Multiple-Serial-IR rules are safe against so-called swapping manipulations. A preference $\succ_i' \in \mathcal{L}_i$ is a *swapping* of $\succ_i \in \mathcal{L}_i$ if for all $k \in N$, $k \succeq_i i \iff k \succeq_i' i$. Hence, a swapping can swap (the order of) goods, but what is (un)acceptable remains (un)acceptable.

Definition 6. Agent $i \in N$ can manipulate circulation rule f at $P \in \mathcal{P}$ by means of swapping if there exists a deviation $P_i' \in \mathcal{P}_i$ such that $\succ^{P_i'}$ is a swapping of $\succ^{P_i}$ and $f_i(P_i', P_{-i}) P_i f_i(P)$. A circulation rule f is (necessarily) *swapping-proof* if no agent can manipulate f at any $P \in \mathcal{P}$ by means of swapping. A circulation rule f is *ig-swapping-proof* if no agent can manipulate f by means of swapping at any profile $P \in \mathcal{P}^L$ of lexicographic preferences.[28] ◇

Proposition 19. *Multiple-Serial-IR rules are swapping-proof.*

Proof. Consider the kth agent, say i_k, in the order of a Multiple-Serial-IR rule. This agent cannot change the choices of the first $k-1$ agents by replacing her true preferences by some swapping. The reason is that restrictions on choices are determined by IR-extendability, which does not vary between agent i_k's true preferences and any swapping (because the set of acceptable goods is the same). Furthermore, at step k, agent i_k weakly prefers choosing her most preferred (feasible) bundle with respect to her true preferences to choosing her most preferred (feasible) bundle with respect to any swapping. □

An agent $i \in N$ is said to be of unit-capacity if $q_i = 1$.

Corollary 20. *Single-Serial-IR rules are swapping-proof for unit-capacity agents.*

Remark 10. We note that Single-Serial rules are in general not ig-swapping-proof. This is a well-known weakness of serial rules (see e.g. Hatfield, 2009) that is experienced for instance in sports drafts when teams sequentially choose one player at a time: sometimes it can be beneficial to choose a popular player rather than a personal favorite among the remaining players, since the latter may still be available in subsequent rounds, while the popular player will surely be taken. Moreover, we also conclude by the same token that Single-Serial-IR rules are not ig-swapping-proof (except for unit-capacity agents, as described in Corollary 20). ◇

Finally, we consider a different kind of manipulation, namely the possibility of hiding endowments.[29] Let $i \in N$ and let q be a capacity profile. We denote $q_{-i} = (q_j)_{j \neq i}$. In the next definition, we express the circulation outcome explicitly as a function of the capacity profile, in addition to the preference profile.

Definition 7. A circulation rule f is *hiding-proof* if for all $i \in N$, $P \in \mathcal{P}$, and $q_i' < q_i$, $f_i(P, q)\ R_i\ f_i(P, (q_{-i}, q_i')) + \frac{q_i - q_i'}{q_i} e_i$.[30] A circulation rule f is *ig-hiding-proof* if for all $i \in N$, $P \in \mathcal{P}^L$, and $q_i' < q_i$, $f_i(P, q)\ R_i\ f_i(P, (q_{-i}, q_i')) + \frac{q_i - q_i'}{q_i} e_i$. ◇

Remark 11. Hiding-proofness implies individual rationality. For instance, in the market exhibited in Example 1 agent 2 can profit by hiding her resources. Hence, there are Single-Serial and Multiple-Serial rules that are not ig-hiding-proof (and hence not hiding-proof). It is easy to check that Single-Serial-IR and Multiple-Serial-IR rules are hiding-proof (and hence ig-hiding-proof). ◇

6. Concluding remarks

6.1. Summary of properties

Table 5 summarizes our findings regarding the properties of the families of circulation rules that we have studied in this paper. In the table ✓ indicates that a property (row) is satisfied by any rule in the family (column) and ✗ indicates that it is not. For a comparison, the table also includes the properties of the most important rules studied in Biró et

[28] Since circulation rules operate on profiles of ordinal preferences over individual goods, equivalent definitions of swapping-proofness and ig-swapping-proofness are obtained by requiring that the deviation P_i' is lexicographic.

[29] In the context of classical exchange economies, Postlewaite (1979) was the first to introduce and study "non-manipulability by withholding."

[30] Note that e_i is the bundle that consists of q_i units of good i. Hence, $\frac{q_i - q_i'}{q_i} e_i$ consists of (the hidden) $q_i - q_i'$ units of good i.

Table 5
Properties of rules.

	Serial rules					
	Single	Single-IR	Multiple	Multiple-IR	cTTC	STC
individually rational	✗	✓	✗	✓	✓	✓
Pareto-efficient	✗	✗	✓	✓	✗	✗
ig-Pareto-efficient	✗	✓	✓	✓	✓	✗
strategy-proof	✗	✗	✓	✗	✗	✓
ig-strategy-proof	✗	✗	✓	✗	✗	✓
dropping-proof	✓	✗	✓	✗	✗	✓
ig-dropping-proof	✓	✗	✓	✗	✓	✓
truncation-proof	✓	✗	✓	✗	✓	✓
ig-truncation-proof	✓	✗	✓	✗	✓	✓
swapping-proof	✗	✗	✓	✓	✗	✓
ig-swapping-proof	✗	✗	✓	✓	✗	✓
hiding-proof	✗	✓	✗	✓	✓	✓
ig-hiding-proof	✗	✓	✗	✓	✓	✓

Table 6
Basis for the properties.

	Serial rules			
	Single	Single-IR	Multiple	Multiple-IR
individually rational	Example 1	By def.	Example 1	By def.
Pareto-efficient	Example 2	Remark 3 + Example 2	By def.	Proposition 17
ig-Pareto-efficient	Example 3	Corollary 12	By def.	Proposition 17
strategy-proof	Remark 10	Example 5	Trivial	Example 5
ig-strategy-proof	Remark 10	Example 5	Trivial	Example 5
dropping-proof	Trivial	Example 5	Trivial	Example 5
ig-dropping-proof	Trivial	Example 5	Trivial	Example 5
truncation-proof	Trivial	Example 5	Trivial	Example 5
ig-truncation-proof	Trivial	Example 5	Trivial	Example 5
swapping-proof	Remark 10	Remark 10	Trivial	Proposition 19
ig-swapping-proof	Remark 10	Remark 10	Trivial	Proposition 19
hiding-proof	Remark 11	Remark 11	Remark 11	Remark 11
ig-hiding-proof	Remark 11	Remark 11	Remark 11	Remark 11

al. (2022): the circulation Top Trading Cycles (cTTC) rule and the family of Segmented Trading Cycle (STC) rules. As we noted earlier, Proposition 1 in Biró et al. (2022) shows that there is no rule that satisfies individual rationality, ig-Pareto-efficiency, and ig-strategy-proofness. As is clear from Table 5, there are rules that satisfy any two of the three properties: (1) Multiple-Serial rules satisfy ig-Pareto-efficiency and ig-strategy-proofness, (2) STC rules satisfy individual rationality and ig-strategy-proofness, and (3) Single/Multiple-Serial-IR rules (and the cTTC rule) satisfy individual rationality and ig-Pareto-efficiency. To accompany Table 5, we display in Table 6 where the proof comes from for each entry regarding the serial rules. For the proofs of the entries on cTTC and the STC rules we refer to Biró et al. (2022).

6.2. Generalized serial rules

For each Single-Serial/Single-Serial-IR rule we have assumed that there is a fixed q-priority order of the agents, i.e., independently of the preferences. Similarly, for each Multiple-Serial/Multiple-Serial-IR rule we have assumed that there is a fixed order of the agents. However, as we have focused our study on individual rationality and (ig)-Pareto-efficiency, to establish our results and examples in Sections 3 and 4 we have not compared outcomes across different preference profiles. Hence, our analysis also holds for "generalized serial rules" where we allow the order to depend on the preference profile. In particular, we obtain the following result as a corollary to Propositions 13 and 2.

Corollary 21. *Each generalized Multiple-Serial rule is Pareto-efficient. Each Pareto-efficient rule is a generalized Single-Serial rule.*

Note that not every Pareto-efficient rule is a generalized Multiple-Serial rule, see, e.g., the last paragraph in Example 4. Similarly, not every generalized Single-Serial rule is Pareto-efficient, see, e.g., Example 3. Note that in both examples preferences are lexicographic. Fig. 1 depicts our findings on Pareto-efficient circulations and Single-Serial and Multiple-Serial rules in a Venn diagram.

A special case of a generalized Single-Serial-IR rule is the cTTC rule studied in Biró et al. (2022) which is individually rational and ig-Pareto-efficient (see Remark 7). This also follows from the next result, which is a corollary to Proposition 11.

Corollary 22. *A rule is individually rational and ig-Pareto-efficient if and only if it is a generalized Single-Serial-IR rule.*

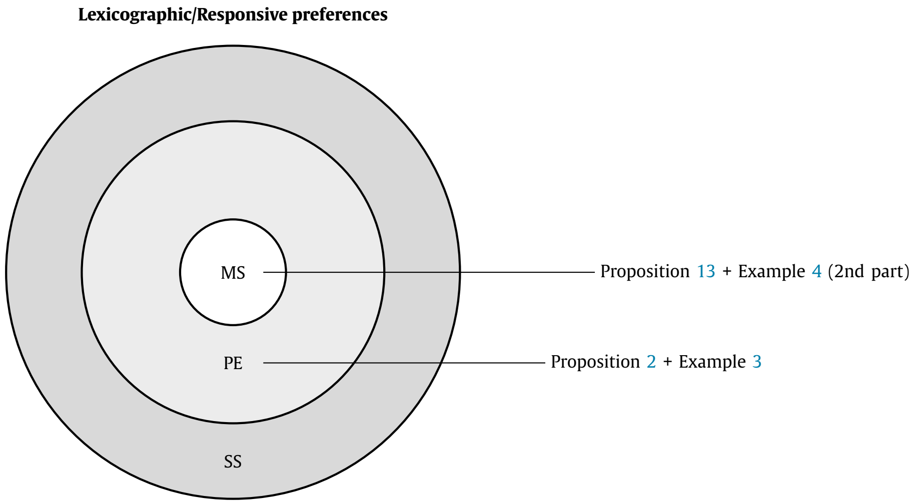

Fig. 1. Venn diagram. Fix preferences P. Let **PE** denote the set of Pareto-efficient circulations at P. Let SS (MS) denote the set of circulations obtained by applying Single-Serial (Multiple-Serial) rules to $\succ^P$. The examples show that the set inclusion can be strict.

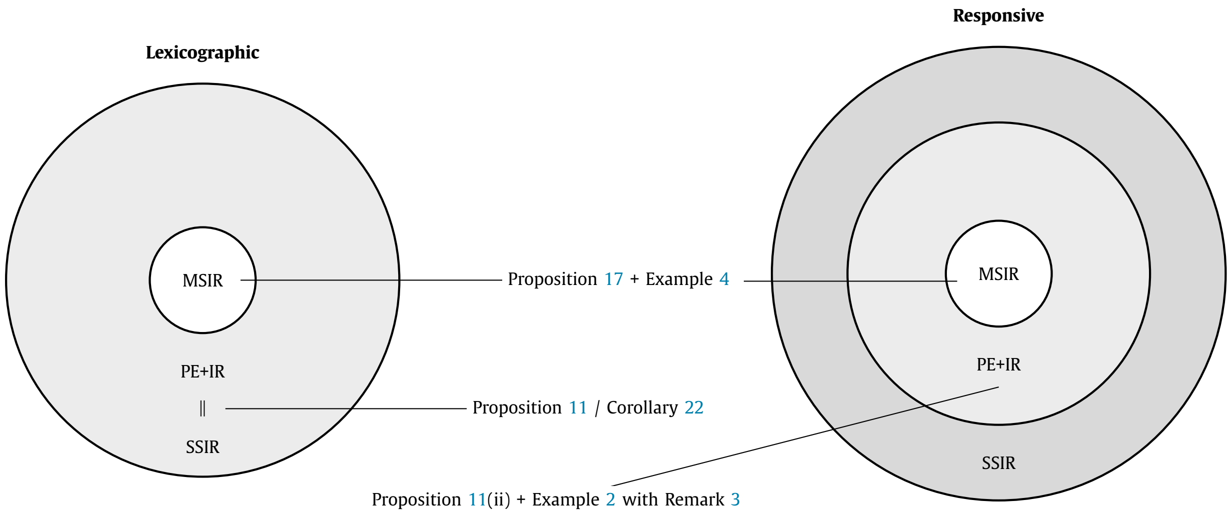

Fig. 2. Venn diagrams. Fix preferences P. Let **PE+IR** denote the set of Pareto-efficient and individually rational circulations at P. Let SSIR (MSIR) denote the set of circulations obtained by applying Single-Serial-IR (Multiple-Serial-IR) rules to $\succ^P$. The examples show that the set inclusion can be strict.

The following result is obtained as a corollary to Propositions 17 and 11(ii).

Corollary 23. *Each generalized Multiple-Serial-IR rule is Pareto-efficient and individually rational. Each Pareto-efficient and individually rational rule is a generalized Single-Serial-IR rule.*

Note that not every Pareto-efficient and individually rational rule is a generalized Multiple-Serial-IR rule, see, e.g., Example 4 (where preferences are lexicographic). When preferences are not lexicographic, not every generalized Single-Serial-IR rule is Pareto-efficient and individually rational, see, e.g., Example 2 coupled with Remark 3. Fig. 2 depicts our findings on Pareto-efficient and individually rational circulations and Single-Serial-IR and Multiple-Serial-IR rules in a Venn diagram.

Figs. 1 and 2 show that the natural interest in Pareto-efficient circulations and Pareto-efficient and individually rational circulations should motivate a further study of serial rules.

6.3. Extensions

Obviously, our negative results still hold in extended models. We describe below how our positive results may be extended to models with link-capacities, heterogeneous goods, or more complex preferences.

Link-capacities

Instead of (or besides) the agent-capacities we could have link-capacities, i.e., a cap on the number of goods an agent can send to, or receive, from other agents. 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 agent-capacities to easily relate it to existing models on the exchange of indivisible goods. However, any link-capacitated market can always be transformed into an agent-capacitated market under responsive preferences by introducing artificial agents. For instance, if agent j cannot receive more than q_{ij} units of good i, then we introduce an artificial agent/good $\overline{ij}$ with capacity q_{ij}. In agent i's preferences we replace good j by good $\overline{ij}$, and agent $\overline{ij}$ only finds good j acceptable. Thus, 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 unit capacity),[31] the manipulability of any circulation rule does not change from one setting to the other.

Heterogeneous goods

We can reduce the model with heterogeneous goods to our circulation model as follows. For the sake of exposition, assume that all (units of) goods are distinct. Let each unit of the heterogeneous case be an artificial agent with unit capacity in our circulation model. Each artificial agent only finds acceptable its original owner. Any original agent's preferences over artificial agents are induced by her original preferences over heterogeneous goods. As an illustration, the generalized TTC for the heterogeneous case, which was introduced and studied in Fujita et al. (2015), is equivalent to the cTTC rule for the reduced circulation market with homogeneous goods in Biró et al. (2022). The artificial agents cannot manipulate the cTTC rule because of their unit capacity. Yet the strategic possibilities of the original agents 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.

More complex preferences

First we discuss the relevance of the assumption that an unacceptable good makes a bundle unacceptable. In our definition of responsiveness of agents' preferences over bundles we assume that acceptable bundles can only contain acceptable goods (r1). Our results on Single-Serial and Multiple-Serial rules still hold when (r1) is dropped. The reason is that Single-Serial and Multiple-Serial rules do not satisfy individual rationality. However, the assumption is important for Single-Serial-IR and Multiple-Serial-IR rules. Since these rules were constructed to guarantee individual rationality, it is crucial to have enough structure on the set of acceptable bundles. For instance, to obtain the alternative definition of Single-Serial-IR rules (that does not require checks of IR-extendability) we use (r1), see, e.g., the maximum flow problem employed in the proof of Lemma 3.

Assumption (r1) is reasonable in many real-life applications, such as Erasmus exchanges (where a student cannot be sent to a university she never applied to) or organ exchanges (where only transplantable organs can be accepted by a country). However, there are also many applications where "negative utility goods" (a.k.a. *bads*) can be accepted by agents if they are compensated with "positive utility goods." For instance, consider the allocation of courses to university professors. A professor may have a usual set of acceptable courses, but she may be willing to teach a course she finds much less interesting than any of her usual courses, as long as she is compensated with a new special topics course of her choosing. The non-trivial question of extending our results on individually rational serial rules to cover situations where an acceptable bundle may contain unacceptable goods is left for future research.

One could also consider a different input for the rules. In this paper we only use the ordinal preferences of the agents over the individual goods and assume responsive and lexicographic preference extensions. But circulation rules could also be based on the agents' cardinal utilities of individual goods (see, e.g., Aziz et al., 2019), again with responsive and lexicographic preference extensions. This would extend the class of circulation rules and the set of possible strategic manipulations, for example. More generally, one could study the case where agents submit linear preferences over the whole set of bundles or even choice functions. It could be interesting to focus on particular preference domains, e.g. substitutable choice functions. Such general models are used in some recent studies on stable networks, e.g. Hatfield et al. (2013). However, the main challenge of allowing the agents to submit their full preferences over the possible bundles is that such input would be exponentially large in the number of agents/goods. This is a well-known issue in applications such as course allocation (Budish et al., 2017) or combinatorial auctions (Milgrom, 2000).

7. Related literature and applications

Our paper is in the intersection of two strands of literature, namely the literature that studies serial dictatorships for allocation problems and the literature on exchange with multiple indivisible goods.

[31] We refer to Biró et al. (2022) for further details.

In the first strand, serial dictatorships were shown to satisfy desirable properties such as Pareto-efficiency and strategy-proofness in various allocation problems. Svensson (1999) characterized serial dictatorships by Pareto-efficiency, non-bossiness Satterthwaite and Sonnenschein (1981), and neutrality in the classical house allocation problem where the houses are public endowments. For multiple object allocation, Pápai (2001), Ehlers and Klaus (2003), and Hatfield (2009) obtained the same characterization result on increasingly smaller preference domains. Namely, Pareto-efficiency, non-bossiness, and strategy-proofness characterize *sequential dictatorships* (a variation on serial dictatorships where only the first agent is fixed in the ordering and subsequent agents in the ordering are determined by previous assignments). On the domain where agents always desire a fixed quota of heterogeneous objects and preferences are responsive, Hatfield (2009) also proved that these three axioms together with neutrality characterize the subfamily of serial dictatorships. In a more general setting with agent-specific quotas, Hosseini and Larson (2019) proved that when preferences are lexicographic an allocation rule is strategy-proof, non-bossy, neutral, and satisfies a mild Pareto-efficiency requirement if and only if it is a serial dictatorship. Pápai (2000) studied multiple assignment problems with monotonic and quantity-monotonic preferences and established further similar characterizations of serial dictatorships. Consistency and solidarity axioms were considered in the same model by Klaus and Miyagawa (2001) who also derived serial dictatorship results.

All these papers study serial dictatorship rules that allow each agent to pick a good or a set of goods only once, which accounts for the positive result on incentives, as indicated by the use of strategy-proofness in many characterizations. When agents are allowed to choose only one good at a time and have multiple turns which are not necessarily consecutive, for example as in our Single-Serial and Single-Serial-IR rules, serial dictatorships possess an intricate strategic structure, which was investigated by Manea (2007). He considered a model in which all bundles are acceptable and preferences are represented by additive utility functions and proved that subgame perfect equilibrium circulations are not necessarily Pareto-efficient and generally not every Pareto-efficient circulation is sustained at some subgame perfect equilibrium in the perfect information game induced by serial rules. We discussed the incentive properties of our rules in Section 5. As we saw, the difficulty with incentives stems from two different sources: one is the above-mentioned multiple non-consecutive turns of agents, which applies to the Single-Serial and Single-Serial-IR rules. The other one is that requiring individual rationality interferes with the nice incentive properties of serial dictatorships and creates room for manipulation by truncation, which applies to the Single-Serial-IR and Multiple-Serial-IR rules.

In contrast to our set-up, all of the above papers explored allocation problems without initial private endowments. The second relevant strand of the literature focuses on the exchange of multiple indivisible goods, which presupposes that agents initially own the goods. The first generalization of the Shapley-Scarf market was due to Konishi et al. (2001) who studied the core in a model with multiple types of goods, where each agent initially owns one good of each type and only goods of the same type can be traded for each other. They showed that in this model there is no individually rational, Pareto-efficient, and strategy-proof rule. Klaus (2008) proved that the type-wise top trading cycle rule in this model is not Pareto-dominated by any other strategy-proof rule, while Pápai (2003) obtained an axiomatic characterization of a similar top trading cycles rule in a model with heterogeneous goods and responsive preferences. Pápai (2007) is a further axiomatic study of exchange in a model with general preferences over heterogeneous goods. With the exception of this last paper, all of the above papers on exchange either require or end up with a balanced exchange, depending on the approach they take.

Recently, artificial intelligence and computer science papers also considered related exchange problems. Todo et al. (2014) studied a model with multiple private endowments and showed that individual rationality, Pareto-efficiency, and strategy-proofness are not compatible for lexicographic preferences. Fujita et al. (2015) studied a model with lexicographic preferences and showed that their augmented TTC rule always yields an assignment in the core. Hence, their rule is individual rational and Pareto-efficient, but not strategy-proof. However, they proved that it is NP-hard to find a beneficial preference misreport. Lesca and Todo (2018) considered the so-called service exchange problem where each agent is willing to provide her service in order to receive in exchange the service of someone else. Assuming that each agent cares about the service that she receives and the person who receives her service, they showed that finding an individually rational and Pareto-efficient circulation is NP-hard, unless all preferences are "set-restricted."

Apart from Biró et al. (2022), there are four recent, closely related papers that studied the balanced exchange of multiple indivisible goods. Two of these papers are on tuition and student exchanges from a two-sided (Dur and Ünver, 2019) and one-sided (Dur et al., 2019) perspective, respectively. The other two papers are motivated by time banks (Andersson et al., 2021) and shift reallocation (Manjunath and Westkamp, 2021). We discuss below the main differences among the models as well as the main findings of these four papers.

Dur and Ünver (2019) studied a model where the agents on the two sides of the market are students and universities. Students want to exchange their seats and universities are interested in exchanging their enrolled students.[32] In the largest students exchange program of this kind, the European Erasmus program, students pay their tuition fee to their "home university" during the exchange period. So, to ensure the longevity of the program, it is essential that exchanges be balanced, i.e., for each university, the number of incoming students equals the number of outgoing students. Each university has a priority order over its outgoing students and responsive preferences over incoming students. The latter assumption on the universities' preferences fits many markets, especially labor markets and those of tuition exchanges, where exchanges are

[32] Dur and Ünver (2019) also listed many other applications with similar characteristics where students exchange their tuition, teachers or other professionals exchange their positions temporarily, etc.

often long-term.[33] Assuming that both sides of the market are strategic, Dur and Ünver (2019) proposed a two-sided top trading cycles rule (2S-TTC). They showed that 2S-TTC is balanced-efficient, group strategy-proof for students, acceptable, respecting internal priorities, individually rational, and immune to quota manipulation by universities. Moreover, they proved that 2S-TTC is the unique rule that satisfies the first four properties.

In other applications the exchange is short-term, such as the Erasmus exchange program where students are visiting foreign universities for one or two semesters. In this case it seems reasonable to assume that universities care most about their outgoing students, since these students will come back and graduate at their home university. In their follow-up paper on Erasmus exchange, Dur et al. (2019) dropped the assumption of Dur and Ünver (2019) that universities have preferences over incoming students, but kept the internal priority order of universities over their outgoing students. It is assumed that this internal priority order is a non-strategic decision, although in practice it can be strategic, as universities may care about which of their students temporarily visit other universities. Dur et al. (2019) studied a generalized version of the TTC rule with both cycles and chains, allowing for small deviations from the balancedness condition. This approach is closer to Biró et al. (2022) where we also studied generalized TTC rules, while in the current paper we focused on serial rules which are closer to current practices in the Erasmus exchange program.

Andersson et al. (2021) studied a balanced exchange problem motivated by time banks. In time banks the participants exchange their services in a one-to-one fashion without monetary transfers. In practice, this is usually implemented either by bilateral agreements or through a dynamic credit system. The model of Andersson et al. (2021) is similar to ours: (i) agents have agent-specific goods and only care about the goods they receive and (ii) the outcome is required to be balanced. However, the main difference is that (Andersson et al., 2021) focused on a different preference domain where each agent (a) has dichotomous preferences over other agents' goods and (b) has a specific upper bound for each acceptable good (i.e., not one upper bound for the size of bundles). A bundle is acceptable if and only if it contains only acceptable goods and respects the associated upper bounds. An acceptable bundle is preferred to another acceptable bundle if the former contains more goods from other agents. For this setting Andersson et al. (2021) proposed a rule that is individually rational and maximizes the total number of acceptable goods exchanged in a balanced way, which guarantees Pareto-efficiency. They showed that their rule is also strategy-proof and that the underlying graph algorithm can be implemented efficiently. As shown in Biró et al. (2022), the three properties (individual rationality, Pareto-efficiency, and strategy-proofness) are incompatible in our model, except for very specific capacity configurations.

Manjunath and Westkamp (2021) studied a balanced exchange problem motivated by shift reallocation. Their model is different from ours in that each agent is assumed to be endowed with heterogeneous goods, in the sense that each agent (worker) can be endowed with different goods (shifts), and not all the goods of an agent may be acceptable to another agent. They studied a restricted trichotomous preference domain: all desirable goods are ranked first, in the most preferred indifference class, followed by all undesirable goods endowed to the agent, leaving the undesirable goods of others for the third and lowest-ranked indifference class. These assumptions are natural in the context of shift exchanges studied by Manjunath and Westkamp (2021), since the acceptability of a shift mainly depends on its timing and not on whose pre-assigned shift it was. In contrast to both (Andersson et al., 2021) and our paper, they dispense with the assumption that a bundle is acceptable if it contains only acceptable goods. Similarly to Andersson et al. (2021), their main result is an efficiently computable rule that is individually rational, Pareto-efficient, and strategy-proof. However, their property of Pareto-efficiency is slightly weaker than the maximal volume property of Andersson et al. (2021), and the two algorithms and the proofs for strategy-proofness are also different.

Student exchange programs, time banks, and shift reallocation are all real-life applications that are captured by our model with responsive preferences or can be studied using a slightly adapted model. Another relevant application is financial clearing. Banks or companies often have cyclic liabilities or debts that can cause liquidity problems or even create systemic risk. In a financial clearing (or *portfolio compression*) the parties involved agree to clear the same amount of debt in a cycle of liabilities (see for example Csóka and Herings, 2018; D'Errico and Roukny, 2021; and Schuldenzucker and Seuken, 2020). Each party has natural preferences over all possible clearances. For instance, each party may want to secure payments from riskier partners first. The search for clearing cycles can be coordinated by private companies or national agencies (as in e.g. Gavrila and Popa, 2021). Any proposed set of clearing cycles constitutes a circulation in the market, and vice versa: any circulation can be decomposed into clearing cycles (see Veraart, 2020). Multiple-Serial-IR rules could serve as appropriate preference-based solutions in these markets for which the particular selection order may be based on an objective criterion such as the financial vulnerability of the companies. If the participants agree to accept any clearing cycle, which ensures that dropping manipulations cannot occur, then the Multiple-Serial-IR rule becomes strategy-proof, given that it is swapping-proof (see Proposition 19).

References

Andersson, T., Cseh, Á., Ehlers, L., Erlanson, A., 2021. Organizing time banks: lessons from matching markets. Am. Econ. J. Microecon. 13 (1), 338–373.
Aziz, H., Biró, P., Lang, J., Lesca, J., Monnot, J., 2019. Optimal reallocation under additive and ordinal preferences. Theor. Comput. Sci. 790, 1–15.
Aziz, H., Gaspers, S., Mackenzie, S., Walsh, T., 2015. Fair assignment of indivisible objects under ordinal preferences. Artif. Intell. 227, 71–92.
Biró, P., Klijn, F., Pápai, S., 2022. Balanced exchange in a multi-unit Shapley-Scarf market. Barcelona School of Economics Working Paper 1342.

[33] An example is the French teacher re-allocation scheme (Combe et al., 2022).

Budish, E., Cachon, G.P., Kessler, J.B., Othman, A., 2017. Course match: a large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. Oper. Res. 65 (2), 314–336.
Cechlárová, K., Eirinakis, P., Fleiner, T., Magos, D., Mourtos, I., Potpinková, E., 2014. Pareto optimality in many-to-many matching problems. Discrete Optim. 14, 160–169.
Combe, J., Tercieux, O., Terrier, C., 2022. The design of teacher assignment: theory and evidence. Rev. Econ. Stud. https://doi.org/10.1093/restud/rdac002, forthcoming.
Csóka, P., Herings, J.-J., 2018. Decentralized clearing in financial networks. Manag. Sci. 64 (10), 4681–4699.
D'Errico, M., Roukny, T., 2021. Compressing over-the-counter markets. Oper. Res. 69 (6), 1660–1679.
Dur, U.M., Ünver, M.U., 2019. Two-sided matching via balanced exchange. J. Polit. Econ. 127 (3), 1156–1177.
Dur, U.M., Kesten, O., Ünver, M.U., 2019. Flexible top trading cycles and chains mechanism: maintaining diversity in erasmus student exchange. Working paper. North Carolina State University.
Ehlers, L., Klaus, B., 2003. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. Soc. Choice Welf. 21 (2), 265–280.
Fujita, E., Lesca, J., Sonoda, A., Todo, T., Yokoo, M., 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, pp. 907–913.
Gavrila, L.-I., Popa, A., 2021. A novel algorithm for clearing financial obligations between companies – an application within the Romanian Ministry of Economy. Algorithmic Finance 9 (1–2), 49–60.
Gourvès, L., 2019. Agreeable sets with matroidal constraints. J. Comb. Optim. 37 (3), 866–888.
Hall, P., 1935. On representatives of subsets. J. Lond. Math. Soc. 10 (1), 26–30.
Hatfield, J.W., 2009. Strategy-proof, efficient, and nonbossy quota allocations. Soc. Choice Welf. 33 (3), 505–515.
Hatfield, J.W., Kominers, S.D., Nichifor, A., Ostrovsky, M., Westkamp, A., 2013. Stability and competitive equilibrium in trading networks. J. Polit. Econ. 121 (5), 966–1005.
Hosseini, H., Larson, K., 2019. Multiple assignment problems under lexicographic preferences. In: Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems. AAMAS 2019, Montreal, Canada, May 13–17, 2019.
Klaus, B., 2008. The coordinate-wise core for multiple-type housing market is second-best incentive compatible. J. Math. Econ. 44 (9–10), 919–924.
Klaus, B., Miyagawa, E., 2001. Strategy-proofness, solidarity, and consistency for multiple assignment problems. Int. J. Game Theory 30, 421–435.
Kojima, F., 2013. Efficient resource allocation under multi-unit demand. Games Econ. Behav. 82, 1–14.
Konishi, H., Quint, T., Wako, J., 2001. On the Shapley-Scarf economy: the case of multiple types of indivisible goods. J. Math. Econ. 35 (1), 1–15.
Lesca, J., Todo, T., 2018. Service exchange problem. In: Proceedings of IJCAI-18, pp. 354–360.
Ma, J., 1994. Strategy-proofness and the strict core in a market with indivisibilities. Int. J. Game Theory 23 (1), 75–83.
Manea, M., 2007. Serial dictatorship and Pareto optimality. Games Econ. Behav. 61 (2), 316–330.
Manjunath, V., Westkamp, A., 2021. Strategy-proof exchange under trichotomous preferences. J. Econ. Theory 193, 105197.
Milgrom, P.R., 2000. Putting auction theory to work: the simultaneous ascending auction. J. Polit. Econ. 108 (2), 245–272.
Pápai, S., 2000. Strategyproof multiple assignment using quotas. Rev. Econ. Des. 5, 91–105.
Pápai, S., 2001. Strategy-proof and nonbossy multiple assignments. J. Public Econ. Theory 3 (3), 257–271.
Pápai, S., 2003. Strategyproof exchange of indivisible goods. J. Math. Econ. 39 (8), 931–959.
Pápai, S., 2007. Exchange in a general market with indivisible goods. J. Econ. Theory 132 (1), 208–235.
Postlewaite, A., 1979. Manipulation via endowments. Rev. Econ. Stud. 46 (2), 255–262.
Satterthwaite, M., Sonnenschein, H., 1981. Strategy-proof allocation mechanisms at differentiable points. Rev. Econ. Stud. 48 (4), 587–597.
Schrijver, A., 2003. Combinatorial Optimization: Polyhedra and Efficiency. Algorithms and Combinatorics, vol. 24. Springer-Verlag, Berlin.
Schuldenzucker, S., Seuken, S., 2020. Portfolio compression in financial networks: incentives and systemic risk. In: Proceedings of the 21st ACM Conference on Economics and Computation.
Shapley, L.S., Scarf, H.E., 1974. On cores and indivisibility. J. Math. Econ. 1 (1), 23–37.
Svensson, L.G., 1999. Strategy-proof allocation of indivisible goods. Soc. Choice Welf. 16 (4), 557–567.
Todo, T., Sun, H., Yokoo, M., 2014. Strategyproof exchange with multiple private endowments. In: Proceedings of the 28th AAAI Conference on Artificial Intelligence, pp. 805–811.
Veraart, L.A.M., 2020. When does portfolio compression reduce systemic risk? Working paper, SSRN 3688495.