

# Allocation of Divisible Goods Under Lexicographic Preferences

Leonard J. Schulman<sup>1</sup> and Vijay V. Vazirani<sup>2</sup>

<sup>1</sup> Caltech, MC305-16, Pasadena CA 91125, USA, [schulman@caltech.edu](mailto:schulman@caltech.edu)

<sup>2</sup> College of Computing, Georgia Institute of Technology, Atlanta GA 30332, US, [vazirani@cc.gatech.edu](mailto:vazirani@cc.gatech.edu)

---

## Abstract

We present a simple and natural non-pricing mechanism for allocating divisible goods among strategic agents having lexicographic preferences. Our mechanism has favorable properties of strategy-proofness (incentive compatibility). In addition (and even when extended to the case of Leontief bundles) it enjoys Pareto efficiency, envy-freeness, and time efficiency.

**1998 ACM Subject Classification** F.2.2 Nonnumerical Algorithms and Problems, J.4 Social and Behavioral Sciences

**Keywords and phrases** Mechanism design, lexicographic preferences, strategyproof, Pareto optimal, incentive compatible.

**Digital Object Identifier** 10.4230/LIPIcs.FSTTCS.2015.543

## 1 Introduction

The study of principled ways of allocating divisible goods among agents has long been a central topic in mathematical economics. The method of choice that emerged from this study, the Arrow-Debreu market model [1], provides a powerful approach based on pricing and leads to the fundamental welfare theorems. However, these market-based methods have limitations when agents are assumed to be strategic, e.g., these methods are not incentive compatible. Issues of the latter kind have been studied within the area of mechanism design for the last four decades, and have played a large role in the last decade in algorithmic game theory [20].

In this paper our primary focus is on deriving a non-pricing mechanism for allocating divisible goods, that satisfies incentive-compatibility, Pareto optimality and envy-freeness. A natural approach to achieving Pareto optimality and envy-freeness is to start in a greedy fashion by assigning agents their most favored goods, and gradually moving on to their less favored choices. It is easy to come up with several ways of making this approach precise—two are described in Section 6—and achieve Pareto optimality and envy-freeness. However, it is not a priori clear that it is possible to also achieve incentive compatibility, without which a mechanism is of doubtful merit in an environment of strategic agents. In the main contribution of our paper we show that a third version of this approach, the *Synchronized Greedy (SG) mechanism*, achieves all three properties.

The SG mechanism can be seen as generalizing a mechanism introduced by Crès and Moulin [7], called *Probabilistic Serial (PS)*, in the context of a job scheduling problem, and studied further by Bogomolnaia and Moulin [5] for the allocation of indivisible goods<sup>1</sup>. The

---

<sup>1</sup> These mechanisms for allocation of indivisible goods are randomized. Our focus on divisible goods is



preference model assumed by [5] was first order stochastic dominance, which we will shorten to *sd-preference*. They showed that in this model, PS is efficient, envy-proof and weakly incentive compatible. Furthermore, they showed that in this model, no mechanism satisfies all three properties, i.e., efficiency, envy-proofness and incentive compatibility. In view of the second result, we need to relax the model in order to obtain a mechanism satisfying all three properties; we do so by resorting to the *lexicographic preference relation* and assuming that the goods are divisible.

Lexicographic preferences date back to the work of Hausner [11] and are of interest to economists for the following reasons. They yield a total order on the set of all allocations (unlike sd-preferences, say, which only form a partial order) and they can be seen as a strong-preferences limit of von Neumann-Morgenstern utilities. A preference relation that is complete, transitive and satisfies the continuity condition that preferences between allocations are preserved under limits is known to be representable by a utility function [18]. Of these, lexicographic preferences forgo continuity. What favorable properties can be achieved in the area of goods allocation using only non-pricing mechanisms is a difficult question. The present paper can be regarded as carving out a certain special case, namely the limit in which agents have very strong preferences among the goods, and providing strong positive guarantees in this case. In this limit there is an additional motivation to use non-pricing mechanisms, because very strong preferences might cause a pricing mechanism to do little more than ensure that the wealthiest agents get what they want. By focusing on non-pricing mechanisms, we can study what game-theoretic properties an allocation mechanism can achieve, without depending on what resources the agents possess or care to invest in the game.

There are many every-day examples where something like our model comes up—naturally, not in market economy transactions, but in other societal mechanisms for allocation. An important class is allocation of public resources, e.g., placement lotteries in public schools, see Kojima [16] for further examples and references. (Note also that this kind of example employs a standard reduction of the indivisible goods case to the divisible goods case by randomization.)

The recent paper of Saban and Sethuraman [22] builds on our work and solves several open problems stated in an earlier version of this paper [24]; these results are described at the end of Section 1.2. The broader challenge of the utility-functions version of the allocation problem remains largely open. The simplicity of the SG mechanism is perhaps encouraging toward the existence of allocation mechanisms maintaining favorable (maybe weaker) game-theoretic properties in this setting. Finally, we note that independent of our work, Cho [6] has also studied the use of lexicographic preferences in the context of probabilistically assigning indivisible objects to agents.

### Parameters of the problem

In the allocation problem there are  $m$  distinct divisible goods which need to be allocated among  $n$  agents. Good  $j$  ( $1 \leq j \leq m$ ) is available in the amount  $q_j > 0$ , and agent  $i$  ( $1 \leq i \leq n$ ) is to receive a specified  $r_i > 0$  combined quantity of all goods; the parameters satisfy  $\sum_j q_j \geq \sum_i r_i$ , i.e., the total supply is at least as large as the total demand. If this inequality fails, our mechanism may still be run after rescaling expectations so that each

---

just as general, since an allocation of divisible goods can be used without further modification as a randomized allocation of indivisible goods in the same quantities.

agent  $i$  is to receive the quantity  $r'_i = r_i(\sum q_j)/(\sum r_\ell)$ . So in the sequel we may assume  $\sum_j q_j \geq \sum_i r_i$ .

**Preferences: the non-Leontief case**

The *non-Leontief* case of our problem is this. An *allocation* of goods is a list of numbers  $a_{ij} \geq 0$ , with  $\sum_j a_{ij} = r_i$  and  $\sum_i a_{ij} \leq q_j$ , indicating that agent  $i$  receives quantity  $a_{ij}$  of good  $j$ . The vector  $a_{i*} = (a_{i1}, \dots, a_{im})$  is referred to as agent  $i$ 's (share of the) allocation. Each agent  $i$  has a *preference list*, which is a permutation  $\pi_i$  of the goods;  $(a_{i\pi_i(1)}, \dots, a_{i\pi_i(m)})$  is agent  $i$ 's *sorted allocation*. Agent  $i$ 's preference among allocations is induced by *lexicographic order*. That is to say, agent  $i$  *lexicographic-prefers*  $a_{i*}$  to  $b_{i*}$ , written  $a_{i*} >_i b_{i*}$ , if the leftmost nonzero coordinate of  $(a_{i\pi_i(1)}, \dots, a_{i\pi_i(m)}) - (b_{i\pi_i(1)}, \dots, b_{i\pi_i(m)})$  is positive. Furthermore, we will say that agent  $i$  prefers  $a_{i*}$  to  $b_{i*}$  in the stochastic domination order [5], or *sd-prefers*  $a_{i*}$  to  $b_{i*}$ , written  $a_{i*} >_i^{sd} b_{i*}$ , if

$$\text{for all } k = 1, \dots, m : \sum_{\ell=1}^k a_{i\pi_i(\ell)} \geq \sum_{\ell=1}^k b_{i\pi_i(\ell)},$$

with at least one of the inequalities being strict. The symbols  $\geq_i$  and  $\not>_i$  will have the obvious interpretations.

Since an agent's preferences depend only on his own share of the allocation, we speak interchangeably of an agent's preference for an allocation or an allocation share. In particular,  $a_{i*} >_i b_{i*}$  may be written more simply as  $a >_i b$ , and  $a_{i*} >_i^{sd} b_{i*}$  may be written as  $a >_i^{sd} b$ .

**Preferences: Leontief Bundles**

Some of our results hold in the more general setting of lexicographic preferences among Leontief bundles, and some fail in that setting; details below. A Leontief bundle is specified by a non-negative vector  $\lambda = (\lambda_1, \dots, \lambda_m) \in \mathbb{R}_+^m$  (where  $\mathbb{R}_+$  = non-negative reals). The set of goods  $j$  for which  $\lambda_j$  is positive is called the *support* of this bundle. (If the set is of size one, we refer to this as a singleton bundle; in Economics this is sometimes also called the linear case.) If  $q \in \mathbb{R}_+^m$  then the bundle  $\lambda$  may be allocated from  $q$  in any quantity  $\alpha \in \mathbb{R}_+$  such that  $\alpha\lambda_j \leq q_j$  for all  $j$ . In an instance of our problem, a list of  $M$  Leontief bundles  $\lambda^1, \dots, \lambda^M$  is specified, including among them the  $m$  singleton bundles (hence always  $M \geq m$ ). It is convenient, and in our context sacrifices no generality, to impose the convention that for every bundle  $\lambda^k$ ,  $\sum_1^m \lambda_j^k = 1$ .

The case  $m = M$ , in which all bundles are singletons, is of course a special case of the Leontief framework, but to distinguish it from the general situation we call it the "non-Leontief" case.

The framework we are concerned with is that each agent  $i$  has a preference list specified by a permutation  $\pi_i$  of the bundles. A Leontief allocation is an  $n \times M$  matrix  $\ell$  in which  $\ell_{ik}$  represents the quantity of bundle  $k$  allocated to agent  $i$ . A Leontief allocation  $l$  imposes the goods allocation  $A(l)$ , an  $n \times m$  matrix, by  $A(l)_{ij} = \sum_{k=1}^M \ell_{ik} \lambda_j^k$ . We further require that a Leontief allocation satisfy the conditions  $\sum_j A(l)_{ij} = r_i$  (thanks to the convention above this is equivalent to  $\sum_k \ell_{ik} = r_i$ ) and  $\sum_i A(l)_{ij} \leq q_j$ . We speak of  $A(l)_{i*}$  and  $l_{i*}$  as agent  $i$ 's *share* of, respectively, the goods and the Leontief bundles. The vector  $(l_{i\pi_i(1)}, \dots, l_{i\pi_i(M)})$  is agent  $i$ 's *sorted Leontief share*. Agent  $i$ 's preference among allocations is induced by *lexicographic order* on his share of the allocation. That is to say, agent  $i$  *lexicographic-prefers*  $l$  to  $l'$ , written  $l >_i l'$ , if the leftmost nonzero coordinate of  $(l_{i\pi_i(1)}, \dots, l_{i\pi_i(M)}) - (l'_{i\pi_i(1)}, \dots, l'_{i\pi_i(M)})$  is positive. Thus, for any goods allocation  $a$ , there is a favored Leontief allocation, denoted

$L^\pi(a)$ , defined by providing each agent with the best Leontief share that can be assembled from his share of the goods—to be explicit, this is obtained by starting with  $a_{i^*}$  as the available goods vector, and then, for  $k$  from 1 to  $M$ , setting  $L^\pi(a)_{i\pi_i(k)}$  to be the largest  $\alpha$  such that  $((\text{available goods vector}) - \alpha\lambda^k) \in \mathbb{R}_+^M$ , then subtracting  $\alpha\lambda^k$  from the available goods vector and iterating.

We say that agent  $i$  *sd-prefers* allocation  $a$  to  $b$ , written  $a >_i^{\text{sd}} b$ , if

$$\text{for all } K = 1, \dots, M : \sum_{k=1}^K L^\pi(a)_{i\pi_i(k)} \geq \sum_{k=1}^K L^\pi(b)_{i\pi_i(k)},$$

with at least one of the inequalities being strict.

## The two orders

Observe that “lexicographic-prefers” is a complete preference relation without indifference contours (since it is antisymmetric for distinct allocation shares), and that “sd-prefers” is an incomplete preference relation; moreover the lexicographic order is a refinement of the sd order, i.e., sd-prefers implies lexicographic-prefers. The phrase “agent  $i$  weakly X-prefers” will be used to include the possibility that agent  $i$ ’s share is identical in the two allocations.

### 1.1 Our results

The SG mechanism is deterministic, treats all agents symmetrically, and has the following properties.

#### Properties w.r.t. sd preference

- If all  $r_i$ ’s are equal, the allocation produced by the SG mechanism in response to truthful bids is envy-free in the following sense: each agent weakly sd-prefers his allocation to that of any other agent. This holds also in the Leontief case.

#### Properties w.r.t. lexicographic preference

(Since most of our paper deals with the relation “lexicographic-prefers”, we subsequently abbreviate it to “prefers”.)

- The allocation produced by the SG mechanism in response to truthful bids is Pareto efficient. This holds also in the Leontief case.
- Incentive compatibility for a single agent: In the non-Leontief case, the SG mechanism is strategy-proof if  $\min_j q_j \geq \max_i r_i$ .

We give counterexamples (a) in the absence of this inequality, (b) for the Leontief case.

- Generalizing the previous item, we have: Incentive compatibility for a coalition: The SG mechanism is group strategy-proof against coalitions of  $\ell$  agents if

$$\min_j q_j \geq \max_{S:|S|=\ell} \sum_{i \in S} r_i$$

- The running time to implement the SG mechanism is  $\tilde{O}(mn)$  in the non-Leontief case, and  $\tilde{O}(n(m^2 + M))$  in the Leontief case.
- Any Pareto efficient allocation can be produced using a suitable “variable speeds” extension of the SG mechanism. This holds also in the Leontief case. (However, the variable speeds extension does not possess the rest of the properties listed above.)

The incentive compatibility properties are the main results of this paper.

## 1.2 Literature

There has been considerable work on the strategy-proof allocation of divisible goods in Arrow-Debreu economies, starting with the seminal work of Hurwicz [12], e.g., see [8, 14, 23, 25, 26, 28]. Most of these results are negative, among the recent ones being Zhou’s result showing that in a 2-agent,  $n$ -good pure exchange economy, there can be no allocation mechanism that is efficient, non-dictatorial (i.e., both agents must receive non-zero allocations) and strategy-proof [28].

The paper that is most closely related to our work is that of Bogomolnaia and Moulin [5]. In their setting there are  $n$  agents and  $n$  indivisible goods, each agent having a total preference ordering over the goods; the desired outcome is a matching of goods with agents. A straightforward mechanism for allocating one good to each agent is *random priority (RP)*: pick a uniformly random permutation of the agents and ask each agent in turn to select a good among those left. It is easy to see that this mechanism is *ex post efficient*, i.e., the allocation it produces can be represented as a probability distribution over Pareto efficient deterministic allocations, and it is strategy-proof. However, it is not *ex ante efficient*. A random allocation is said to be *ex ante efficient* if for any profile of von Neumann-Morgenstern utilities that are consistent with the preferences of agents, the expected utility vector is Pareto efficient. It is easy to see that *ex ante efficiency* implies *ex post efficiency*.

Solving a conjecture of Gale [9], Zhou [27] showed that no strategy-proof mechanism that elicits von Neumann-Morgenstern utilities and achieves Pareto efficiency can find a “fair” solution even in the weak sense of equal treatment of equals. He further showed that the solution found by RP may not be efficient if agents are endowed with utilities that are consistent with their preferences. Hence, *ex ante efficiency* had to be sacrificed, if strategy-proofness and fairness were desired.

In the face of these choices, the work of Bogomolnaia and Moulin gave the notion of *ordinal efficiency* that is intermediate between *ex post* and *ex ante efficiency*; an allocation  $a$  is *ordinally efficient* if there is no other allocation  $b$  such that every agent *sd-prefers*  $b$  to  $a$ . They went on to show that the mechanism called *probabilistic serial (PS)*, introduced in Crès and Moulin [7], yields an *ordinally efficient* allocation. Further they show that PS is *envy-free* and *weakly strategy-proof*, defined appropriately for the partial order “*sd-prefers*”. Finally, Bogomolnaia and Moulin define an extension of PS by introducing different “*eating rates*” and show that this set of mechanisms characterizes the set of all *ordinally efficient* allocations.

Katta and Sethuraman [15] generalize the setting of Bogomolnaia and Moulin to the “*full domain*”, i.e., agents may be indifferent between pairs of goods. Thus, each agent partitions the goods by equality and defines a total order on the equivalence classes of her partition (the agent is equally happy with any good received from an equivalence class). For this setting, they give a randomized mechanism that is a generalization (different from ours) of PS and achieves the same game-theoretic properties as PS.

A mechanism that probabilistically allocates indivisible goods can also be viewed as one that fractionally allocates divisible goods. Under the latter interpretation, the SG mechanism is equivalent to PS for the case that  $m = n$  and the quantity of each good and the requirement of each agent is one unit. An important difference is that Bogomolnaia and Moulin analyze PS under an incomplete preference relation (stochastic dominance) in which “*most*” allocation shares are incomparable; whereas we analyze SG under a complete preference relation (lexicographic) that is a refinement of stochastic dominance. The statement that a mechanism’s allocation is Pareto efficient w.r.t. lexicographic preferences is considerably stronger than the same statement w.r.t. stochastic dominance preferences, because each

agent's share is dominated by more alternative shares in the lexicographic order, than it is in the sd order; so, fewer allocations are Pareto efficient in the lexicographic than in the sd order. Our results should be viewed therefore as demonstrating that the PS mechanism and its natural generalization, SG, have far stronger game-theoretic properties than even envisioned in [5].

For somewhat related questions primarily regarding exchange economies, see Barberà and Jackson [4], Nicolo [19], Ghodsi et al. [10], and Li and Xue [17]. Finally, we remark only that the problem of allocating a *single* divisible good among multiple agents with known privileges is considerably different; the principal issue studied in that problem is how to make the division in a manner that is fair w.r.t. the given privileges. This is known as the bankruptcy problem and has a long history, e.g., see [21, 2]. Despite an interesting resemblance between the PS mechanism and some of the mechanisms used in the solutions of that problem [13], the issues at stake in the bankruptcy literature are distinct from those in our paper and its predecessors.

Saban and Sethuraman [22] solve some of the open problems stated in an earlier version of this paper. They consider the special case that all  $r_i = 1$ . First they show that our condition  $\min_j q_j \geq \max_i r_i$  is tight in the sense that for any  $q_1 < 1$  there exists an  $n$ , a finite list  $q_2, \dots, q_n$ , and agent preferences such that no mechanism is efficient, envy-free and strategyproof. They also show that if  $q_1 < 1$ , and list  $q_2, \dots, q_n$  and the agent preferences are given, then SG achieves all three properties if and only if any mechanism achieves all three properties. Finally for the generalized setting of Katta and Sethuraman, where agents can be indifferent between objects, they show that no mechanism can satisfy all three properties.

Since the PS rule is not strategyproof, recent work has studied the situation where agents are strategic. A Nash equilibrium for the PS rule is a preference profile for which no agent has an incentive to report a different profile. [3] show that a pure Nash equilibrium is guaranteed to exist; however determining whether a given preference profile is a Nash equilibrium is coNP-complete.

## 2 The Synchronized Greedy Mechanism

The mechanism is simple. Each agent  $i$  submits a preference list  $\sigma_i$ . The submitted list may or may not, of course, agree with his true preference list  $\pi_i$ .

(A simple case to consider is that of  $M = m = n$  and all  $q_j = r_i = 1$ . Because of the restriction that each preference list must include all  $m$  singleton bundles, each agent's preference list in this case is a permutation of the  $m$  goods. Despite being quite special, this case, or the slightly more general case in which  $M = m \leq n$  and all  $r_i$  are equal, is already interesting to analyze and is well motivated by the examples, mentioned earlier, involving sharing of tasks or of scarce public resources.)

The mechanism simulates the following physical process. Consider each good  $j$  as a "liquid", and each agent as a receptacle of capacity  $r_i$ . The mechanism starts out at time 0 by (for all  $i$  in parallel) pouring bundle  $\lambda^{\sigma_i(1)}$  into receptacle  $i$  at rate  $r_i$  units of liquid per unit time. Each good  $j$  is therefore being drained at rate  $\sum_i r_i \lambda_j^{\sigma_i(1)}$ . (Note that since  $\sum_j \lambda_j = 1$ , the total liquid being added to receptacle  $i$  per unit time is  $r_i$ , as desired.)

This continues until one of the goods, say  $j$ , is exhausted. For all agents who were currently being allocated bundles with  $j$  in their support, their favorite Leontief bundle has now been exhausted. (We say that a Leontief bundle has been *exhausted* at a given time if any of the goods in its support has been exhausted, and otherwise that the bundle is *available*.) All such agents,  $i$ , are immediately allocated the next available bundle on their preference list,

and the pouring of bundles continues. The algorithm continues in this way, allocating to an agent from the next available bundle whenever the current bundle has been exhausted. Since the singleton bundles are included in all preference lists, all agents continuously receive goods at rate  $r_i$  until time 1, at which time they simultaneously complete their full allocation.

Observe that the Leontief allocation  $l$  constructed by SG satisfies  $l = L^\pi(A(l))$  because the bundles are provided to each agent greedily based on the availability of goods.

This continuous process can easily be converted into a discrete algorithm with the run time cited earlier: maintain a priority queue of goods, keyed by termination times. Each time a good is exhausted, each agent is assigned its next unexhausted bundle, and an updated termination time for each good is computed using the coefficients of the active bundles.

Observe that if an agent prefers bundle  $\lambda$  to bundle  $\lambda'$ , and  $\text{support}(\lambda) \subseteq \text{support}(\lambda')$ , then  $\lambda'$  may be removed from the agent's preference list. It cannot be allocated to the agent by SG nor can it be part of any Pareto efficient allocation to the agent.

### 3 Properties of the Synchronized Greedy Mechanism

#### 3.1 Pareto Efficiency

Let  $l^\sigma$  be the allocation created by the SG mechanism in response to bids  $\sigma$  declared by the agents. As before  $\pi$  denotes the truthful bids.

► **Theorem 1.** *The allocation produced by the SG mechanism in response to truthful bids is Pareto efficient w.r.t. lexicographic preference. That is to say, for all  $l \neq l^\pi$ ,  $\exists i$   $l <_i l^\pi$ .*

**Proof.** For agent  $i$  and for  $K \geq 1$  let  $t_{iK} = \frac{1}{r_i} \sum_{k=1}^K l_{i\pi_i(k)}^\pi$ . If agent  $i$  receives a positive quantity of his  $K$ 'th-most-favored bundle, then  $t_{iK}$  is the time when that bundle is exhausted in SG. If the agent receives nothing from the bundle then the bundle is exhausted in SG no later than  $t_{iK}$ .

Suppose for contradiction the existence of  $l$  s.t.  $\forall i$   $l \geq_i l^\pi$ , and for some  $i$ ,  $l >_i l^\pi$ . Let  $t$  be minimum s.t.  $\exists i, K$  s.t.  $t = t_{iK} < \frac{1}{r_i} \sum_{k=1}^K l_{i\pi_i(k)}$ . Note, if  $t_{i'K'} < t$  then  $t_{i'K'} = \frac{1}{r_{i'}} \sum_{k=1}^{K'} l_{i'\pi_{i'}(k)}$ .

For every one of the bundles  $b \in \{\pi_i(1), \dots, \pi_i(K)\}$  there is a good  $j(b)$  that appears positively in  $b$  and which is exhausted by time  $t$ . Since  $t_{iK} < \frac{1}{r_i} \sum_{k=1}^K l_{i\pi_i(k)}$  while  $t_{iK'} = \frac{1}{r_i} \sum_{k=1}^{K'} l_{i\pi_i(k)}$  for all  $K' < K$ , some agent  $i' \neq i$  receives strictly less of good  $j(\pi_i(K))$  in  $l$  than in  $l^\pi$ . Since  $j(\pi_i(K))$  is exhausted in SG by time  $t$ , this means that there is some  $K''$  such that  $\frac{1}{r_{i'}} \sum_{k=1}^{K''} l_{i'\pi_{i'}(k)} < \frac{1}{r_{i'}} \sum_{k=1}^{K''} l_{i'\pi_{i'}(k)}^\pi \leq t$ . This contradicts the minimality of  $t$ . ◀

#### 3.2 Strategy-Proofness

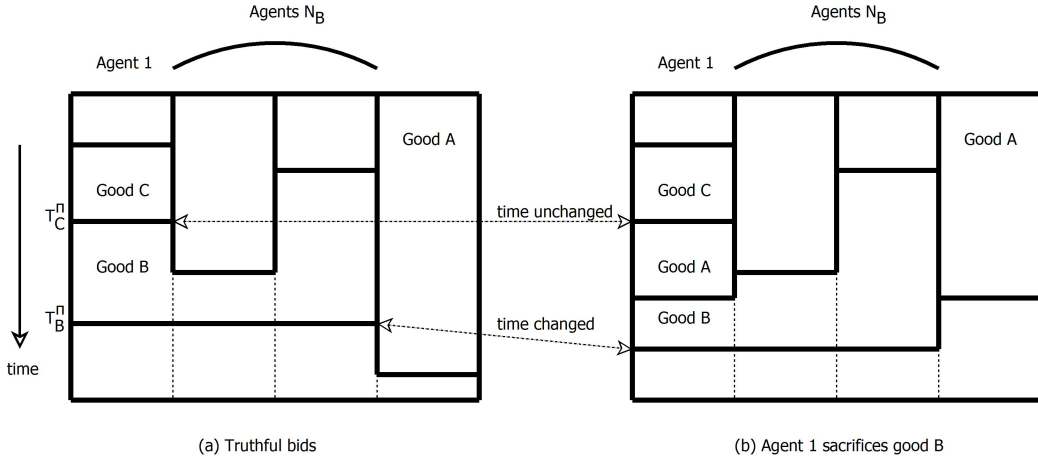
A mechanism is said to be *strategy-proof* if for every agent and for every list of bids by the remaining agents, the agent cannot obtain a strictly improved allocation by lying.

► **Theorem 2.** *In the non-Leontief case, the SG mechanism is strategy-proof if  $\min q_j \geq \max r_i$ .*

**Proof.** Without loss of generality focus on agent 1. For the remainder of this proof  $\pi_2, \dots, \pi_n$  are arbitrary bids by the agents 2,  $\dots$ ,  $n$ , but  $\pi_1$  is agent 1's truthful bid. We need to show that for any bid  $\sigma_1$  (and write  $\sigma = (\sigma_1, \pi_2, \dots, \pi_n)$ ),  $a_{1*}^\sigma \leq a_{1*}^\pi$ . The theorem is trivial if  $a^\sigma = a^\pi$ .

The theorem is also trivial if agent 1, bidding truthfully, receives only his top choice. So we may suppose that agent 1 does not receive the entire allocation of any one good.





■ **Figure 1** The mechanism with truthful vs. lying bids of Agent 1.

We may also suppose that if  $a_{1j}^\sigma = 0$  and  $a_{1j'}^\sigma > 0$ , then  $\sigma_1^{-1}(j) > \sigma_1^{-1}(j')$ . (Define  $\sigma_1^{-1}(j)$  to be the  $s$  such that  $\sigma_1(s) = j$ . Define  $\pi_1^{-1}(j)$  analogously.) In other words, all the requests in  $\sigma_1$  that come up empty may as well be deferred to the end.

Let  $G(j) = \{j' : \pi_1^{-1}(j') \leq \pi_1^{-1}(j) \text{ and } a_{1j'}^\pi > 0\}$ . These are the goods that agent 1 weakly prefers to good  $j$  and receives a positive quantity of in the allocation  $a^\pi$ .

Say that agent 1 *sacrifices* good  $j$  in  $\sigma$  if:

1.  $a_{1j}^\pi > 0$ ,
2.  $\sigma_1^{-1}(j) > |G(j)|$ , and
3.  $\pi_1^{-1}(j) < \pi_1^{-1}(j')$  if  $j'$  also satisfies (1),(2).

That is to say,  $j$  is the most-preferred good which agent 1 receives a positive quantity of in  $\pi$ , but requests later in  $\sigma$  than in  $\pi$ .

For a collection of bids  $\rho$  let  $T_j^\rho$  be the time at which good  $j$  is exhausted if the mechanism is run with bids  $\rho$ .

Agent 1 must sacrifice some good, call it  $B$ , since otherwise the allocation will not change. See Figure 1. We will show that agent 1 receives strictly less of  $B$  in  $\sigma$  than in  $\pi$ , and that this is not compensated for by getting more of more-preferred goods.

► **Lemma 3.** *If  $D$  is a good and  $T_D^\sigma < T_B^\pi$ , then  $T_D^\sigma \leq T_D^\pi$ .*

**Proof.** Supposing the contrary, let  $D$  be a counterexample minimizing  $T_D^\pi$ . Since  $T_D^\sigma < T_B^\pi$ ,  $D \neq B$ . Now let  $i$  be any agent (who may or may not be agent 1) for whom  $a_{iD}^\pi > 0$ . Due to the minimality of  $D$ , each of the goods  $j$  which  $i$  prefers in  $\pi$  to  $D$ , has  $T_j^\sigma \leq T_j^\pi$ . Therefore  $i$  requests  $D$  at a time in  $\sigma$  that is at least as soon as the time  $i$  requests it in  $\pi$ .

Since this holds for all  $i$  who received a positive allocation of  $D$  in  $\pi$ , the lemma follows. ◀

Let  $N_B$  be the set of agents  $i \neq 1$  for whom  $a_{iB}^\pi > 0$ . The condition on  $r_i$ 's and  $q_j$ 's ensures that this set is nonempty.

Due to the lemma, for each agent in  $N_B$ , the request time for  $B$  in  $\sigma$  is weakly earlier than it is in  $\pi$ . Now let  $C$  be the good such that  $\pi_1^{-1}(C)$  is maximal subject to  $\pi_1^{-1}(C) < \pi_1^{-1}(B)$  and  $a_{1C}^\pi > 0$ . Due to the lemma, all goods  $j'$  such that  $\pi_1^{-1}(j') \leq \pi_1^{-1}(C)$  have  $T_{j'}^\sigma \leq T_{j'}^\pi$ . Next we show:

► **Proposition 4.** *If  $\pi_1^{-1}(j') \leq \pi_1^{-1}(C)$ , then  $a_{1j'}^\sigma = a_{1j'}^\pi$ .*






Figure 2 Failure of strategy-proofness without the hypothesis of Theorem 2.

**Proof.** Supposing the contrary, let  $\pi_1^{-1}(j')$  be minimal such that  $\pi_1^{-1}(j') \leq \pi_1^{-1}(C)$  and  $a_{1j'}^\sigma \neq a_{1j'}^\pi$ . There are two possibilities to consider.

- (a)  $a_{1j'}^\sigma < a_{1j'}^\pi$ . This is not possible because then  $a_{1*}^\sigma < a_{1*}^\pi$ .
- (b)  $a_{1j'}^\sigma > a_{1j'}^\pi$ . Note:

► **Lemma 5.** *Let  $j_1, j_2$  be such that  $\pi_1^{-1}(j_1) \leq \pi_1^{-1}(B)$ ,  $\pi_1^{-1}(j_2) \leq \pi_1^{-1}(B)$ ,  $a_{1j_1}^\pi > 0$ , and  $\pi_1^{-1}(j_1) < \pi_1^{-1}(j_2)$ . Then  $\sigma_1^{-1}(j_1) < \sigma_1^{-1}(j_2)$ .*

**Proof.** Consider the least  $j_1$  that is part of a pair  $j_1, j_2$  violating the lemma. Then  $j_1$  satisfies conditions (1),(2) above, contradicting that  $B$  is the good sacrificed by agent 1. ◀

It follows that  $T_{j'}^\sigma \geq \sum_{j'': \pi_1^{-1}(j'') \leq \pi_1^{-1}(j')} a_{1j''}^\sigma$ . Due to the minimality of  $j'$ , this means that if  $a_{1j'}^\sigma > a_{1j'}^\pi$ , then  $T_{j'}^\sigma > T_{j'}^\pi$ , contradicting our earlier conclusion. This completes demonstration of the Proposition. ◀

A consequence of the Proposition is that  $T_C^\sigma = T_C^\pi$ .

Since agent 1 sacrifices  $B$ , his request time for  $B$  in  $\sigma$  is strictly greater than his request time for  $B$  in  $\pi$ .

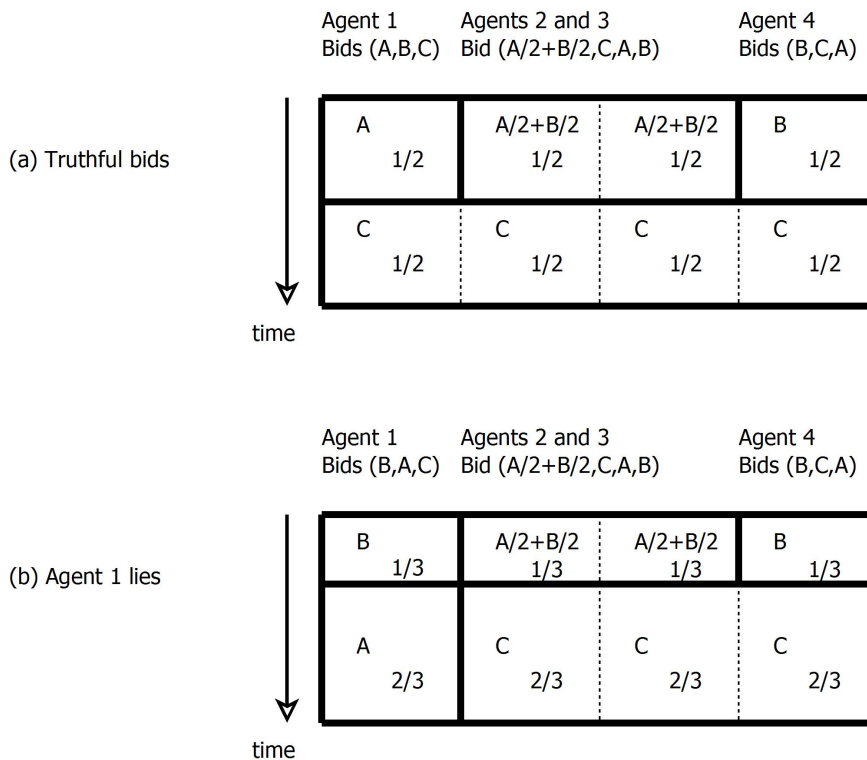
Recall that  $N_B$  is nonempty. At time  $T_B^\pi$ , the agents of  $N_B$  have received as least as much of  $B$  in  $\sigma$  as they have in  $\pi$ , and the latter is positive. On the other hand, at the same time  $T_B^\pi$ , agent 1 has received strictly less of  $B$  in  $\sigma$  than he has in  $\pi$ . In order for agent 1 to receive at least as much of  $B$  in  $\sigma$  as in  $\pi$ , he would have to receive all of  $B$  that is allocated after time  $T_B^\pi$ ; however, that is not possible, because the set of agents receiving  $B$  after  $T_B^\pi$  includes  $N_B$ . Thus  $a_{1*}^\sigma < a_{1*}^\pi$ . ◀

### 3.3 Necessity of a Hypothesis on $\{r_i\}, \{q_j\}$ s

We next provide an example in which strategy-proofness fails in the absence of the condition  $\max r_i \leq \min q_j$ . For convenience now let  $r_1 \geq \dots \geq r_n$  and  $q_1 \leq \dots \leq q_m$ .

► **Example 6.** Let  $n = 2$  and  $m = 3$ . Let  $r_1 = r_2 = 3/2$ ; label the goods  $A, B, C$ , let  $q_A = q_B = q_C = 1$ , and let the preference lists be  $\pi_1 = (A, B, C)$ ,  $\pi_2 = (B, C, A)$ . If agent 1 bids truthfully he receives the sorted allocation  $(1, 0, 1/2)$ . If instead he bids  $(B, A, C)$  (while agent 2 bids truthfully), he receives the improved sorted allocation  $(1, 1/2, 0)$ . See Figure 2.

This example does not limit the theorem sharply, because it uses  $r_1 = (3/2)q_1$  rather than  $r_1$  arbitrarily close to  $q_1$ . Jeremy Hurwitz has pointed out that one may construct similar examples whenever  $r_1 \geq q_1 / (1 - q_2 / \sum q_j)$ ; this would appear to be a tight bound.



■ **Figure 3** Failure of strategy-proofness in the Leontief case.

### 3.4 Failure of strategy-proofness for the Leontief case

Theorem 2 has no equivalent for general Leontief bundles. Consider the following four-agent system with  $r_1 = r_2 = r_3 = r_4 = 1$  and three goods in supply  $q_A = q_B = 1, q_C = 2$ . Agent 1's desired Leontief bundles are in the preference order  $(A, B, C)$  (this agent is interested only in singleton bundles); agent 2 and 3's desired Leontief bundles are in the order  $(\frac{1}{2}A + \frac{1}{2}B, C, A, B)$ ; agent 4's Leontief bundles are in the order  $(B, C, A)$ .

Under truthful bidding agent 1 receives the sorted goods allocation  $(1/2, 0, 1/2)$ . By bidding instead  $(B, A, C)$ , agent 1 receives the improved sorted goods allocation  $(2/3, 1/3, 0)$ . See Figure 3.

### 3.5 Group Strategy-Proofness

A mechanism is *group strategy-proof* against a family  $F$  of subsets of agents if for every "coalition"  $S \in F$  and for any list of bids by the agents outside of  $S$ , the agents of  $S$  cannot obtain an improved allocation by lying, where by "improved allocation" we mean that no agent of  $S$  obtains a worse allocation and at least one obtains a strictly better allocation.

We now provide the following generalization of Theorem 2:

► **Theorem 7.** *In the non-Leontief case, the SG mechanism is group strategy-proof against the family of subsets  $S$  for which  $\min_j q_j \geq \sum_{i \in S} r_i$ .*

► **Corollary 8.** *In the non-Leontief case, the SG mechanism is group strategy-proof against coalitions of  $\ell$  agents if  $\min_j q_j \geq \max_{S: |S|=\ell} \sum_{i \in S} r_i$ .*

The proof of Theorem 7 follows a structure similar to that of Theorem 2 but the argument is complicated by the fact that different agents in  $S$  can sacrifice different goods, and some of the agents may actually be better off due to their untruthful bids (as they may benefit from the interactions among the several lies). The proof needs to effectively “chase through” an unbounded iteration of good transfers relative to  $a^\pi$ , and show that *some* agent in the coalition is worse off than in  $\pi$ . Fortunately, this can be done without explicitly pursuing the iteration.

**Proof.** Let  $S$  be a minimal counterexample. That is,

- (a)  $\min_j q_j \geq \sum_{i \in S} r_i$ ;
- (b) With  $\pi_i$  representing in this proof the truthful preferences for  $i \in S$  and arbitrary preferences for  $i \notin S$ , there are bids  $\sigma_i$  for  $i \in S$  such that every  $i \in S$  “is a willing participant in the coalition  $S$ ”, namely (with  $\sigma_\ell = \pi_\ell$  for  $\ell \notin S$ )  $a_{i^*}^\sigma \geq_i a_{i^*}^\pi$ ;
- (c) For some  $i \in S$ ,  $a_{i^*}^\sigma >_i a_{i^*}^\pi$ ;
- (d) No strict subset of  $S$  satisfies (a),(b),(c).

Note by minimality that in  $\sigma$ , every agent  $i \in S$  bids untruthfully (differently from  $\pi$ ) and this has an effect, namely, if  $i$  reverts to bidding according to  $\pi$  then the allocation is different than in  $\sigma$ .

If  $a_{i^{\pi_i(1)}}^\pi = r_i$  for all  $i \in S$ , that is, with truthful bids these agents receive only their top choices, then none of them can be strictly rewarded by submitting a different bid.

Otherwise (i.e., if  $a_{i^{\pi_i(1)}}^\pi < r_i$  for some  $i \in S$ ), then thanks to the hypothesis, under the truthful bids  $\pi$ , every good has a positive allocation outside  $S$ .

We may simplify the argument slightly by supposing that for each agent  $i \in S$ , if  $a_{i^j}^\sigma = 0$  and  $a_{i^j}^\sigma > 0$ , then  $\sigma_i^{-1}(j) > \sigma_i^{-1}(j')$ . In other words, all the requests that come up empty may as well be deferred to the end.

Let  $G(i, j) = \{j' : \pi_i^{-1}(j') \leq \pi_i^{-1}(j) \text{ and } a_{i^j}^\sigma > 0\}$ .

Say that agent  $i$  *sacrifices* good  $j$  in  $\sigma$  if:

1.  $a_{i^j}^\sigma > 0$ ,
2.  $\sigma_i^{-1}(j) > |G(i, j)|$ , and
3.  $\pi_i^{-1}(j) < \pi_i^{-1}(j')$  if  $j'$  also satisfies (1),(2).

Some good must be sacrificed by some agent, since otherwise the allocation will not change. (However, while every agent in  $S$  is untruthful, not every  $i \in S$  necessarily sacrifices a good; setting  $\sigma_i(j) > \pi_i(j)$  might have an effect even if  $a_{i^j}^\sigma = 0$  because of increased availability of  $j$  due to bidding changes of other agents.)

Of all the sacrificed goods let  $B$  be one for which  $T_B^\pi$  is minimal.

► **Lemma 9.** *If  $D$  is a good and  $T_D^\pi < T_B^\pi$ , then  $T_D^\sigma \leq T_D^\pi$ .*

**Proof.** Supposing the contrary, let  $D$  be a counterexample minimizing  $T_D^\pi$ . By the minimality of  $B$ ,  $D$  cannot be a sacrificed good.

Now let  $i$  be any agent (inside or outside of  $S$ ) for whom  $a_{i^D}^\pi > 0$ . Due to the minimality of  $D$ , each of the goods  $j$  which  $i$  truthfully prefers to  $D$ , has  $T_j^\sigma \leq T_j^\pi$ . Therefore  $i$  requests  $D$  at a time in  $\sigma$  that is at least as soon as the time  $i$  requests it in  $\pi$ .

Since this holds for all  $i$  who received a positive allocation of  $D$  in  $\pi$ , the lemma follows. ◀

Let  $O_B \subseteq S$  be the set of agents who sacrifice  $B$ , and let  $N_B$  be the set of agents  $i$  for whom  $a_{i^B}^\pi > 0$  but who do not sacrifice  $B$ . Due to the lemma, for each agent in  $N_B$ , the request time for  $B$  in  $\sigma$  is weakly earlier than it is in  $\pi$ . Now consider an agent  $i \in O_B$ . Let  $C$  be the good such that  $\pi_i^{-1}(C)$  is maximal subject to  $\pi_i^{-1}(C) < \pi_i^{-1}(B)$  and  $a_{i^C}^\pi > 0$ . Due to the lemma, all goods  $j'$  such that  $\pi_i^{-1}(j') \leq \pi_i^{-1}(C)$  have  $T_{j'}^\sigma \leq T_{j'}^\pi$ . Next we show:

► **Proposition 10.** *If  $\pi_i^{-1}(j') \leq \pi_i^{-1}(C)$ , then  $a_{ij'}^\sigma = a_{ij'}^\pi$ .*

**Proof.** Supposing the contrary, let  $\pi_i^{-1}(j')$  be minimal such that  $\pi_i^{-1}(j') \leq \pi_i^{-1}(C)$  and  $a_{ij'}^\sigma \neq a_{ij'}^\pi$ . There are two possibilities to consider.

- (a)  $a_{ij'}^\sigma < a_{ij'}^\pi$ . This is not possible because  $i$  is a willing participant in the coalition.
- (b)  $a_{ij'}^\sigma > a_{ij'}^\pi$ . Note:

► **Lemma 11.** *Let  $j_1, j_2$  be such that  $\pi_i^{-1}(j_1) \leq \pi_i^{-1}(B)$ ,  $\pi_i^{-1}(j_2) \leq \pi_i^{-1}(B)$ ,  $a_{ij_1}^\pi > 0$ , and  $\pi_i^{-1}(j_1) < \pi_i^{-1}(j_2)$ . Then  $\sigma_i^{-1}(j_1) < \sigma_i^{-1}(j_2)$ .*

**Proof.** Identical to the proof of Lemma 5 with agent  $i$  in place of agent 1. ◀

It follows that  $T_{j'}^\sigma \geq \sum_{j'': \pi_i^{-1}(j'') \leq \pi_i^{-1}(j')} a_{ij''}^\sigma$ . Due to the minimality of  $j'$ , this means that if  $a_{ij'}^\sigma > a_{ij'}^\pi$ , then  $T_{j'}^\sigma > T_{j'}^\pi$ , contradicting our earlier conclusion. This completes demonstration of the Proposition. ◀

A consequence of the Proposition is that  $T_C^\sigma = T_C^\pi$ .

Since agent  $i$  sacrifices  $B$ , his request time for  $B$  in  $\sigma$  is strictly greater than his request time for  $B$  in  $\pi$ .

Since we are in the case that every good has a positive allocation outside  $S$ ,  $N_B$  is nonempty. At time  $T_B^\pi$ , the agents of  $N_B$  have received as least as much of  $B$  in  $\sigma$  as they have in  $\pi$ , and the latter is positive. On the other hand, at the same time  $T_B^\pi$ , the agents of  $O_B$  have received strictly less of  $B$  in  $\sigma$  than they have in  $\pi$ . In order for the agents of  $O_B$  to receive collectively at least as much of  $B$  in  $\sigma$  as in  $\pi$ , they would have to receive all of  $B$  that is allocated after time  $T_B^\pi$ ; however, that is not possible, because the set of agents receiving  $B$  after  $T_B^\pi$  includes  $N_B$ . Therefore there is some  $i \in O_B$  for whom  $a_{iB}^\sigma < a_{iB}^\pi$ . This contradicts the requirement that  $i$  be a willing participant in the coalition  $S$ . ◀

► **Example 12.** Example 6, in which strategy-proofness failed absent the hypothesis of Theorem 2, can be extended in a straightforward manner to one in which the group strategy-proof property fails to hold absent the hypothesis of Corollary 8. Again use  $m = 3$ , but instead of two agents, use  $n = 2\ell$  agents, the first half having the same preference order  $(A, B, C)$  as agent 1 in the earlier example, and the second half having the same preference order  $(B, C, A)$  as agent 2 in the earlier example. If all agents bid truthfully, then the first  $\ell$  agents each receive the sorted allocation  $(1, 0, 1/2)$ ; however if they lie and bid  $(B, A, C)$ , while the remainder bid truthfully, then each lying agent receives the improved sorted allocation  $(1, 1/2, 0)$ .

## 4 Characterizing All Pareto Efficient Allocations

Bogomolnaia and Moulin [5] extended their mechanism by allowing players to receive goods at time-varying rates. Specifically, for each agent  $i$  there is a speed function  $\eta_i$  mapping the time interval  $[0, 1]$  into the nonnegative reals, such that for all  $i$ ,  $\int_0^1 \eta_i(t) dt = r_i$ . Subject to these speeds, goods flow to agents in order of the preference lists they bid, just as before. They showed that this extension characterizes all ordinally efficient allocations.

In this section, we obtain an analogous characterization of all Pareto efficient allocations by a similar extension of our mechanism. Specifically, we prove that for *any* Pareto efficient allocation of bundles, there exist speeds such that the extended SG mechanism produces that allocation. We prove this after first noting that the extended SG mechanism always results in Pareto efficient allocations.

In this section when  $\eta_i$  ( $1 \leq i \leq n$ ) are fixed, we let  $a^\pi$  (with the  $\eta$ 's implicit) be the goods allocation produced by the extended SG mechanism with these speeds and truthful bids. We let  $l^\pi = L^\pi(a^\pi)$  be the corresponding allocation of bundles.

#### 4.1 Pareto Efficiency

► **Theorem 13.** *Let  $\eta_i$ ,  $1 \leq i \leq n$ , be any speed functions. Then the allocation  $l^\pi$  is Pareto efficient.*

**Proof.** The argument is the same as for Theorem 1 with the proviso that the definition  $t_{iK} = \frac{1}{r_i} \sum_{k=1}^K l_{i\pi_i(k)}^\pi$  is replaced by  $t_{iK} = \inf\{y : \int_0^y \eta_i(t) dt \geq \sum_{k=1}^K l_{i\pi_i(k)}^\pi\}$ . ◀

#### 4.2 Characterizing All Pareto Efficient Allocations

If the last result mirrored the First Welfare Theorem, the next mirrors the Second Welfare Theorem:

► **Theorem 14.** *Let  $\pi$  be the collection of agent preference lists over bundles, and let  $l$  be a Pareto efficient allocation. There exist speed functions  $\eta_i$ ,  $1 \leq i \leq n$ , such that  $l = l^\pi$ .*

**Proof.** As before the bundles are  $(\lambda^k)_{k=1}^M$ , where for each  $k$ ,  $\sum_{j=1}^m \lambda_j^k = 1$ , and  $\lambda_j^k \geq 0$  for all  $j$ .

Construction of the speeds  $\eta_i$  is simple. Let a “partial bundle allocation” be a list  $\hat{l}_{ik}$ , each  $\hat{l}_{ik} \geq 0$ , such that for every  $i$ ,  $\sum_{k,j} \hat{l}_{ik} \lambda_j^k \leq r_i$ , and for every  $j$ ,  $\sum_{i,k} \hat{l}_{ik} \lambda_j^k \leq q_j$ .

Initialize  $t = 0$  and initialize each agent  $i$  with the empty partial allocation  $\hat{l}_{ik} = 0$  for all  $i, k$ .

Initialize  $c_j$  to be the quantity of good  $j$  that is allocated in  $l$ . (Necessarily  $c_j \leq q_j$  and  $\sum c_j = \sum r_i$ . If  $\sum q_j > \sum r_i$  then for some  $j$ ,  $c_j < q_j$ .)

Then repeat the following until  $t = 1$ .

Find an agent  $i$  for whom there is an  $\ell$  such that  $\hat{l}_{i\pi_i(\ell)} < l_{i\pi_i(\ell)}$ , and such that for all  $\ell' < \ell$ , the bundle  $\pi_i(\ell')$  has been exhausted (that is to say, there is a good  $j$  such that  $\lambda_j^{\pi_i(\ell')} > 0$  and  $c_j = 0$ .) To see that there is such an  $i$ , suppose the contrary, and consider all the agents for whom  $\sum_{k,j} \hat{l}_{ik} \lambda_j^k < r_i$ . For each of them there is a favorite bundle which has not yet been exhausted. Evidently none of these agents is to be allocated in  $l$  any additional quantity of this favorite bundle. However since these favorite bundles have not yet been exhausted, we can allocate to every player a slight additional positive amount of his favorite unexhausted bundle, without exhausting any additional goods. Any extension of this new partial bundle allocation to a full bundle allocation, strictly Pareto dominates  $l$ , contrary to assumption.

Now set  $\delta = (l_{i\pi_i(\ell)} - \hat{l}_{i\pi_i(\ell)}) / \sum r_i$ . For  $t < t' < t + \delta$ , make the settings  $\eta_i(t') = \sum r_i$  and, for  $i' \neq i$ ,  $\eta_{i'}(t') = 0$ . Then increment  $\hat{l}_{i\pi_i(\ell)}$  by  $\delta \sum r_i$ , and decrement each  $c_j$  by the corresponding amount, namely, decrement  $c_j$  by  $\lambda_j^{\pi_i(\ell)} \delta \sum r_i$ . Finally, increment  $t$  by  $\delta$ .

This process terminates in finitely many iterations because in each iteration some agent completes its allocation of some bundle. ◀

Examination of the above proof reveals:

► **Corollary 15.** *There is a polynomial time algorithm for checking whether a given allocation is Pareto efficient.*

### 4.3 No Incentive Compatibility for the Variable Speeds Variant

We note that the synchrony imposed among agents by the SG mechanism is key to its incentive compatibility and envy-freeness properties (indeed, the properties hold even if the basic mechanism is extended with the *same* speed function for all agents). If different agents have different speed functions under the extended SG mechanism, Theorems 2 and 7, showing incentive compatibility, fail to hold. The argument breaks down as soon as it uses termination times, in Lemma 3. Below is a counter-example for strategy-proofness; a similar idea gives counter-examples for group strategy-proofness and envy-freeness.

► **Example 16.** Assume  $m = n = 4$  and that all  $r_i = q_j = 1$ . Let the speed function for agent 1 be 1 over the interval  $[0, 1]$ . The speeds of agents 2, 3, and 4 equal 1 over the interval  $[0, 1/2]$ , 0 over the interval  $(1/2, 5/6]$ , and 3 over the interval  $(5/6, 1]$ . The preference orders of agents 1 and 2 are  $(1, 2, 3, 4)$ , and the preference orders of agents 3 and 4 are  $(2, 4, 3, 1)$ . If all agents bid truthfully, agent 1 receives the sorted allocation  $(1/2, 0, 1/2, 0)$ . On the other hand, if agent 1 bids  $(2, 1, 3, 4)$  while the rest bid truthfully, then agent 1 receives the better sorted allocation  $(1/2, 1/3, 1/6, 0)$ .

## 5 Envy-Freeness w.r.t. stochastic dominance preference

(This section is the only part of the paper where we use sd preference.)

Given a bundle allocation  $l$ , let  $\bar{l}$  denote the *relative allocation*, where  $\bar{l}_{ij} = l_{ij}/r_i$ .

► **Theorem 17.** *Under truthful bidding, every agent  $i$  weakly sd-prefers his relative allocation  $\bar{l}_{i*}^\pi$  to the relative allocation  $\bar{l}_{i' *}^\pi$  of any other agent  $i'$ .*

**Proof.** Fix any  $1 \leq k \leq M$ . We are to show that

$$\frac{1}{r_i} \sum_{\ell=1}^k l_{i\pi_i(\ell)}^\pi \geq \frac{1}{r_{i'}} \sum_{\ell=1}^k l_{i'\pi_i(\ell)}^\pi.$$

Let  $t$  be the time at which the last of the bundles  $\pi_i(1), \dots, \pi_i(k)$  is exhausted. So  $tr_i = \sum_{\ell=1}^k l_{i\pi_i(\ell)}^\pi$ . No other agent can receive any of these bundles after time  $t$ , so  $tr_{i'} \geq \sum_{\ell=1}^k l_{i'\pi_i(\ell)}^\pi$ . ◀

## 6 Other Greedy Mechanisms

As stated in the Introduction, obtaining an efficient and envy-free non-pricing mechanism for allocating divisible goods is easy, but additionally satisfying incentive compatibility is harder. In this section we present two greedy mechanisms which satisfy the first two properties but not the third. To simplify description of the mechanisms, assume that  $m = n$  and that all  $r_i = q_j = 1$ ; it is straightforward to generalize the mechanisms beyond this restriction, and our counterexamples are possible even with it.

**Mechanism 1:** The mechanism proceeds iteratively. In round  $i$ , it considers the  $i$ th-favorite goods of all agents who still have not been allocated a full unit of goods. Among such agents, if the  $i$ th-favorite good of a set  $S$  of agents is good  $j$ , the remaining quantity of good  $j$  is allocated equally among the agents in  $S$ , subject to no agent getting more than a total of one unit of goods. (Some of good  $j$  may remain after the round.)

**Mechanism 2:** The mechanism has a notion of time, similar to SG. Goods allocation starts at time 0 and is completed at time 1. During this interval each agent receives goods at rate 1. The interval is punctuated by finitely many critical instants at which some of the agents switch which good they are receiving. The first critical instant is 0 and the others are the times at which some nonempty set of agents  $T$  finishes receiving their promised allocation of a good. At such an instant, the mechanism identifies, for each of the agents in  $T$ , the next-favorite good on their list that has not yet been fully promised to other agents. The mechanism promises each agent in  $T$  some of that good, in the following fashion: let  $T_j$  be the subset of  $T$  requesting good  $j$  and let  $u$  be the amount of good  $j$  that has not been previously promised. Then each agent in  $T_j$  is promised an equal share of  $u$  subject to no agent exceeding a total of one unit of goods. (The next critical instant affecting these agents is of course easily computed.) The mechanism then proceeds to the next critical instant.

The proofs given above, for showing that the SG mechanism is efficient and envy-free, extend easily to showing that Mechanisms 1 and 2 are also efficient and envy-free. Here, however, are counterexamples to incentive compatibility:

► **Example 18** (Mechanism 1). Let  $m = n = 4$ ; name the goods  $A, \dots, D$ . Agent 1's preference list is  $A, B, C, D$ ; agents 2 and 3 have preferences  $A, C, B, D$ ; and agent 4's favorite good is  $B$ . If the agents bid truthfully then in round 1, agent 4 is allocated all of good  $B$ , while the first three agents are each allocated a third of good  $A$ . In the second round agent 1 is left out while agents 2 and 3 are allocated half of good  $C$ . In round 3 no allocations are made, and in round 4 good  $D$  is allocated among the first three agents. The allocation to agent 1 is therefore  $(A : 1/3, D : 2/3)$ . If instead agent 1 submits the preference list  $A, C, B, D$  then she is treated the same as agents 2 and 3, and her allocation is  $(A : 1/3, C : 1/3, D : 1/3)$ , which she prefers.

The counterexample for the second mechanism is more involved.

► **Example 19** (Mechanism 2). Let  $m = n = 8$ ; name the goods  $A, \dots, H$ . We specify only the essential components of the preference orders. The preference order of agent 1 is alphabetical,  $(A, \dots, H)$ . Agents 2, 3, 4 have the preference order  $(A, G, H, F, \dots)$ . Agents 5, 6, 7 have the preference order  $(B, C, E, F, \dots)$ . Agent 8 has the preference order  $(B, D, \dots)$ . If all agents report their preferences truthfully, agent 1 gets the allocation  $(A : 1/4, C : 1/4, D : 1/4, F : 1/4)$ ; if agent 1 lies and reports the order  $(A, C, E, D, \dots)$  she gets the allocation  $(A : 1/4, C : 1/4, D : 1/4, E : 1/4)$ , which she prefers.

## 7 Discussion

Our main open problem is the one mentioned in the Introduction, i.e., achieving approximate versions of the properties of the SG mechanism but when agents' preferences are representable by utility functions.

Another natural open question concerns the existence of mechanisms to produce lexicographically most equitable allocations, having favorable algorithmic and game-theoretic properties (esp., incentive compatibility). The SG mechanism is not very equitable: see the full paper [24].

**Acknowledgments.** We are indebted to Hervé Moulin for generously sharing his deep understanding of this research domain. Thanks also to Jeremy Hurwitz for stimulating discussions and to Amin Saberi for pointing us to useful references.



Schulman was supported in part by NSF Grants 1038578 and 1319745. Part of this work was done while visiting the Simons Institute for the Theory of Computing at UC Berkeley.

Vazirani was supported in part by NSF Grants CCF-0914732 and CCF-1216019, and a Google Research Grant. He would like to thank the Social and Informational Sciences Laboratory at Caltech for their hospitality; part of this work was done while he was Distinguished SISL Visitor during 2011-12. He would also like to thank the Guggenheim Foundation for a Fellowship during 2011-12.

---

### References

---

- 1 K. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–290, 1954.
- 2 R. Aumann and M. Maschler. Game theoretic analysis of a bankruptcy problem from the Talmud. *J. Economic Theory*, 36:195–213, 1985.
- 3 H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Equilibria under the probabilistic serial rule. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI’15*, pages 1105–1112, 2015.
- 4 S. Barberà and M. O. Jackson. Strategy-proof exchange. *Econometrica*, 63(1):51–87, 1995.
- 5 A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295–328, 2001.
- 6 W. J. Cho. Probabilistic assignment: A two-fold axiomatic approach. 2012.
- 7 H. Crès and H. Moulin. Scheduling with opting out: improving upon random priority. *Operations Research*, 49(4):565–577, 2001.
- 8 P. Dasgupta, P. Hammond, and E. Maskin. The implementation of social choice rules. *Review of Economic Studies*, 46:153–170, 1979.
- 9 D. Gale. College course assignments and optimal lotteries. Mimeo, University of California, Berkeley, 1987.
- 10 A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: fair allocation of multiple resource types. In *Proc. 8<sup>th</sup> USENIX conference on networked systems design and implementation*, 2011.
- 11 M. Hausner. Multidimensional utilities. In R. L. Davis, R. M. Thrall, and C. H. Coombs, editors, *Decision Processes*, pages 167–180. Wiley, 1954.
- 12 L. Hurwicz. On informationally decentralized systems. In C. B. McGuire and R. Radner, editors, *Decision and Organization*, pages 297–336. North-Holland, Amsterdam, 1972.
- 13 M. Kaminski. Hydraulic rationing. *Mathematical Social Sciences*, 40:131–155, 2000.
- 14 M. Kato and S. Ohseto. Toward general impossibility theorems in pure exchange economies. *Social Choice and Welfare*, 19:659–664, 2002.
- 15 A. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131:231–250, 2006.
- 16 F. Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009.
- 17 J. Li and J. Xue. Egalitarian division under Leontief preferences. Manuscript, 2012.
- 18 A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- 19 A. Nicolo. Efficiency and truthfulness with Leontief preferences. *Review of Economic Design*, 8(4):373–382, 2004.
- 20 N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 209–241. Cambridge University Press, 2007.

- 21 B. O'Neill. A problem of rights arbitration from the Talmud. *Mathematical Social Sciences*, 2(4):345–371, 1982.
- 22 D. Saban and J. Sethuraman. A note on object allocation under lexicographic preferences. *J. Mathematical Economics*, 50:283–289, 2014.
- 23 M. Satterthwaite and H. Sonnenschein. Strategy-proof allocation mechanisms at differentiable points. *Review of Economic Studies*, 48:587–597, 1981.
- 24 L.J. Schulman and V. V. Vazirani. Allocation of divisible goods under lexicographic preferences. In arXiv, 1206.4366, 2012.
- 25 S. Serizawa. Inefficiency of strategy-proof rules for pure exchange economies. *Journal of Economic Theory*, 106:219–241, 2002.
- 26 S. Serizawa and J. Weymark. Efficient strategy-proof exchange and minimum consumption guarantees. *Journal of Economic Theory*, 109:246–263, 2003.
- 27 L. Zhou. On a conjecture by Gale about one-sided matching problems. *J. Econ. Theory*, 52:123–135, 1990.
- 28 L. Zhou. Inefficiency of strategy-proof allocation mechanisms in pure exchange economies. *Social Choice and Welfare*, 8(3):247–254, 1991.