A Note on Object Allocation under Lexicographic Preferences

Daniela Saban* and Jay Sethuraman†

March 7, 2014

Abstract

We consider the problem of allocating m objects to n agents. Each agent has unit demand, and has strict preferences over the objects. There are q_j units of object j available and the problem is balanced in the sense that $\sum_j q_j = n$. An allocation specifies the amount of each object j that is assigned to each agent i, when the objects are divisible; when the objects are indivisible and exactly one unit of each object is available, an allocation is interpreted as the probability that agent i is assigned one unit of object j. In our setting, agent preferences over objects are extended to preferences over allocations using the natural lexicographic order. The goal is to design mechanisms that are efficient, envy-free, and strategy-proof. Schulman and Vazirani show that an adaptation of the probabilistic serial mechanism satisfies all these properties when $q_j \geq 1$ for all objects j. Our first main result is a characterization of problems for which efficiency, envy-freeness, and strategy-proofness are compatible. Furthermore, we show that these three properties *do not* characterize the serial mechanism. Finally, we show that when indifferences between objects are permitted in agent preferences, it is impossible to satisfy all three properties even in the standard setting of "house" allocation in which all object supplies are 1.

1 Introduction

We consider the problem of allocating m objects to n agents. Each agent i has unit demand, and an amount q_j, not necessarily integer, of object j is available. We assume that $\sum_j q_j = n$, so that supply and demand are balanced. An allocation $A = (a_{ij})$ is a non-negative $n \times m$ matrix in which each row i adds up to 1 and each column j adds up to q_j. The ith row of the allocation matrix A is also referred to as agent i's allocation.[1] If the objects are divisible, a_{ij} is simply the amount of object j that agent i receives; if the objects are indivisible and exactly one unit of each object is available, we may interpret a_{ij} as the probability that agent i receives object j. [2]

Each agent i has a strict preference ordering P_i over the set of objects: for objects j and k, we say that $j >_{P_i} k$ (or simply $j >_i k$) if and only if agent i strictly prefers j to k. This strict

*Graduate School of Business, Columbia University, New York, NY; dhs2131@columbia.edu

†IEOR Department, Columbia University, New York, NY; jay@ieor.columbia.edu. Research supported by NSF grant CMMI-0916453 and CMMI-1201045.

[1]We use the term "allocation" for both the matrix A, which describes the allocation for all agents, and for any of its rows, which describes the allocation for the corresponding agent.

[2]This case corresponds to the well-studied "house allocation" problem [3, 7, 9]

preference ordering over objects can be extended to a preference ordering over allocations in many different ways. In this paper we shall restrict our attention to the lexicographic extension, first considered in this context by Cho [4] for the probabilistic assignment of indivisible objects, and, independently, by Schulman and Vazirani [8] for allocating divisible objects. Given two allocations A and A', agent i prefers the one in which he receives more of his most-preferred object; if he receives the same amount of his most-preferred object in both, he prefers the one in which he receives more of his second most-preferred object, etc. To state this formally, suppose that agent i's preference ordering over the objects is $1 >_i 2 >_i \ldots >_i m$. Given two allocations A and A', agent i lexicographically prefers A to A' if and only if the first non-zero entry in $(a_{i1} - a'_{i1}, a_{i2} - a'_{i2}, \ldots, a_{im} - a'_{im})$ is positive. Of course, agent i is indifferent between A and A' if and only if i's allocation is the same in A and A'. Note that the preference ordering is *complete* in the sense that given two different allocations for agent i, he will always prefer one to the other.

A *problem* is completely specified by the set of agents N along with their preferences, and the set of objects with the corresponding q_j. A *mechanism* is a function that maps each problem to an allocation. Typically, the goal is to design a mechanism that satisfies certain desirable properties. To state these properties formally, it is important to define the notion of *dominance*.

Definition 1 (**Dominance**). *An allocation A' dominates an allocation A if each agent i (weakly) prefers his allocation in A' to that in A.*

Note that each agent compares his allocation in A to that in A' using the lexicographic extension of his preference ordering over the objects. Using this definition of dominance, we next define *efficiency* and *envy-freeness* for allocations.

Definition 2 (**Efficiency**). *An allocation A is* efficient *if it is not dominated by any other allocation A'.*

Definition 3 (**Envy Freeness**). *An allocation A is envy-free if each agent i (weakly) prefers his allocation to the allocation of any other agent.*

A mechanism is efficient if it associates an efficient allocation with each problem, and it is envy-free if it associates an envy-free allocation with each problem.[3] In other words, a mechanism is said to be *efficient* if, in any problem, there is no way to strictly improve one agent's allocation without making another agent worse off; it is *envy-free* if each agent prefers his allocation to that of any other agent in any problem. Our final desideratum for a mechanism is strategy-proofness. Informally, a mechanism is strategy-proof if it is not possible for any agent to improve his allocation by falsifying his preferences, regardless of the preferences of the other agents. Formally:

[3] In this paper we shall be concerned with envy-free mechanisms only. To be able to compare allocations across agents, we need the agent demands to be identical, and we normalize this demand to 1; this is really the rationale for insisting a priori that every agent have unit demand. The mechanisms described here admit natural extensions to settings in which the agents have different demands.

Definition 4 (Strategy-proofness). *For each agent i, and for each profile of preferences $P_{-i} = (P_1, P_2, \ldots, P_{i-1}, P_{i+1}, P_{i+2}, \ldots, P_n)$, let A_i and A'_i be the allocations of agent i when the mechanism is applied to the profiles (P_i, P_{-i}) and (P'_i, P_{-i}) respectively. A mechanism is* strategy-proof *if $A_i \geq_{P_i} A'_i$ for all i, P_i, P_{-i}, and P'_i.*

Bogomolnaia and Moulin [3] proposed an important new mechanism for the house allocation problem that they called the probabilistic serial (PS) mechanism. The specific preference model they assumed was the one based on first order stochastic dominance (hereafter, the *sd*-preference model): an agent i prefers an allocation A to an allocation A' if A stochastically dominates A', or equivalently, agent i's expected utility of A is greater than that of A' for *every* utility function consistent with his preference ordering. Although this notion leads only to a partial ordering over allocations for each agent, they were able to show that the outcome of the PS mechanism was strongly envy-free in the sense that each agent's allocation stochastically dominates the allocation of any other agent. Furthermore, they proved that the PS mechanism was efficient (no other allocation matrix stochastically dominates the PS outcome), but only weakly strategy-proof (by misreporting his preference ordering, an agent cannot obtain an allocation that stochastically dominates his true allocation). Moreover, they show that no mechanism satisfies efficiency, envy-freeness and strategy-proofness in the *sd*-preference model.

Given that no mechanism is efficient, envy-free, and strategy-proof in the *sd*-preference model, one can consider relaxations of these requirements. The model of this paper, due to Cho [4] and, independently, Schulman & Vazirani [8], can be seen as a relaxation of the stochastic dominance preference model, in which agents compare allocations based on lexicographic preferences. The key motivation for considering lexicographic preferences is three-fold: first, lexicographic preferences form a complete relation so that *every* pair of allocations can be compared; second, it is a weakening of the stochastic dominance preference model in the sense that stochastic dominance implies lex-dominance; and finally, it can be interpreted as a limiting case of agents with von-Neumann-Morgenstern utilities.[4]

Cho [4] observed that the probabilistic serial mechanism of Bogomolnaia and Moulin [3] satisfies all three properties when $m = n$, $q_j = 1$ for all j, and the objects are indivisible. Independently, Schulman and Vazirani [8] studied the lexicographic preference model when agents' demands and objects' supplies are arbitrary[5] and the objects are divisible. Their algorithm—called synchronized greedy (SG)—is efficient in all cases, envy-free when the agents have unit (or equal) demands, and strategy-proof when the maximum amount demanded by an agent is at most the minimum supply of an object. Specialized to our setting then, the SG mechanism

[4] Lexicographic preferences have a long history in the economics literature, dating back at least to Hausner [5], who formulated the theory of lexicographic expected utility. The key point of departure from the standard expected utility theory is that continuity is not required; thus, preferences over lotteries that do not satisfy continuity can be represented by a *vector*-valued utility function, and preferences over lotteries translate to lexicographic dominance of the corresponding utility vectors. We refer the reader to Cho [4] for a more extensive discussion.

[5] Earlier, Heo [6] extended the probabilistic serial mechanism to the more general case of arbitrary object supplies and arbitrary agent demands (not necessarily 1), but focused on *sd*-preference model. She showed that the serial rule (as she termed it) is efficient, and envy-free (if agents have unit demands).

is efficient and envy-free, and is strategy-proof whenever $q_j \geq 1$ for all j.

Our first result, proved in Section 2, is that when this last condition is not satisfied *no* mechanism can satisfy all three properties. Specifically, we show that if there is an object j with $q_j < 1$, then there is an instance involving as many agents as objects for which no strategy-proof mechanism can be both efficient and envy-free. In proving this result we assume that we are free to *choose* all the parameters of the problem: the number of agents, the number (and supply) of objects, and the preference ordering of the agents. A natural question is if a similar result is possible for a given number of agents and a fixed set of objects with a given supply vector, with only the agent preferences varying. For this setting, we characterize when it is possible to find a mechanism satisfying all three properties: we show that either the synchronized greedy mechanism is strategy-proof for every preference profile (and hence satisfies all three properties) or *no* mechanism can satisfy all three properties. One way to interpret this result is that whenever the SG mechanism fails strategy-proofness, so does every other mechanism that is envy-free and efficient. Our next two results answer questions left open by Schulman and Vazirani [8]: In Section 3, we show that the synchronized greedy mechanism is not the only mechanism with these properties by constructing a "hybrid" mechanism that incorporates elements of a greedy mechanism and the SG mechanism. This new mechanism satisfies an invariance property called *bounded invariance* (defined formally in Section 3), in addition to being efficient, envy-free and strategy-proof. An important consequence is that envy-freeness, efficiency, and bounded invariance do not characterize the SG mechanism when agents have lexicographic preferences; this is in sharp contrast to the *sd*-preference model, where these properties *do* characterize the serial mechanism (see Bogomolnaia and Heo [2]). In Section 4 we allow for the agents to be indifferent between objects. The natural extension of the serial mechanism to this setting, developed by Katta and Sethuraman [7], is efficient and envy-free but not strategy-proof. We show that *no* mechanism can be efficient, envy-free and strategy-proof, even for the special case of unit q_j.

2 Strict Preferences: Strategy-proofness

Schulman and Vazirani show that the SG mechanism is efficient and envy-free, and is strategy-proof whenever $q_j \geq 1$ for all j. We start this section by showing that, in the absence of this last condition, *no* mechanism can satisfy all three properties.

Proposition 1. *For any $\epsilon \in (0, 1)$, there exists a problem such that $q_j \leq 1 - \epsilon$ for which no strategy-proof mechanism is both efficient and envy-free.*

Proposition 1 is a special case of Theorem 1. Nevertheless, we present a detailed proof of this result because the proof describes ideas that are used repeatedly throughout the paper.

Proof. Suppose the supply of object a is $1-\epsilon$. Consider an instance of the problem with k agents and k objects, where $k \geq 1/\epsilon$ and $k \geq 3$. The supply of object c is $1 + \epsilon$, and the supply of every other object (including object b) is 1. We will now list the implications of a strategy-proof mechanism that finds an efficient and envy-free allocation.

We start with the following preference profile.

$$
\begin{array}{c|ccccc}
1 & \mathrm{a} & \mathrm{b} & \cdots & \\
2 & \mathrm{b} & \mathrm{c} & \cdots & \mathrm{a} \\
3 & \mathrm{b} & \mathrm{c} & \cdots & \mathrm{a} \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
\mathrm{k} & \mathrm{b} & \mathrm{c} & \cdots & \mathrm{a}
\end{array}
$$

That is, the last $k-1$ agents have identical preferences.

First, by (lex) efficiency, object a must be completely allocated to agent 1. Otherwise, part of a must be allocated to some agent $i \neq 1$, and agent 1 must be allocated a positive amount of some other object $j \neq a$; this is because each agent has a demand of 1 and less than 1 unit of a is available. As every agent other than 1 has a as his least-preferred object, agents 1 and i can exchange a (small) positive amount of their allocations of j and a respectively, making each of them better off, and thus resulting in a violation of efficiency.

Second, we must have $a_{ib} = 1/k$ for all agents i. To see why this holds, consider the following preference profile

$$
\begin{array}{c|ccccc}
1 & \mathrm{b} & \mathrm{a} & \cdots & \\
2 & \mathrm{b} & \mathrm{c} & \cdots & \mathrm{a} \\
3 & \mathrm{b} & \mathrm{c} & \cdots & \mathrm{a} \\
\vdots & \vdots & \vdots & \vdots & \vdots \\
\mathrm{k} & \mathrm{b} & \mathrm{c} & \cdots & \mathrm{a}
\end{array}
$$

in which only agent 1's preference ordering has been changed. By envy-freeness, every agent must be allocated $1/k$ units of b. Note that $1/k \leq \epsilon$, so that agent 1 still needs at least $1-\epsilon$ units of the other objects for his allocation to be complete. Again by efficiency, agent 1 must be allocated all of a. So agent 1 is allocated $1-\epsilon$ units of a and $1/k$ units of b at this preference profile. By strategy-proofness, the allocation at the original preference profile must be the same for agent 1, otherwise agent 1 can manipulate at one of these profiles.

Finally, consider the allocation obtained for the preference profile

$$
\begin{array}{c|ccccc}
1 & \mathrm{a} & \mathrm{b} & \cdots & & \\
2 & \mathrm{b} & \mathrm{a} & \mathrm{c} & \cdots & \\
3 & \mathrm{b} & \mathrm{c} & \cdots & & \mathrm{a} \\
\vdots & \vdots & \vdots & \vdots & & \vdots \\
\mathrm{k} & \mathrm{b} & \mathrm{c} & \cdots & & \mathrm{a}
\end{array}
$$

which differs from the first preference profile in only agent 2's preference ordering. Strategy-proofness for agent 2 implies $a_{2b} = 1/k$, otherwise agent 2 can manipulate at one of these profiles. Now, envy-freeness implies $a_{jb} = 1/k$ for all $j \neq 1$, which, in turn, implies $a_{1b} = 1/k$. This, however, poses a problem: efficiency would require that agent 1 be allocated all of object a, but agent 2 would envy agent 1 in any such allocation. □

The impossibility result proved in Proposition 1 relies on our ability to choose not just the preferences of the agents, but also the number of agents, and the number of objects (and their availabilities). A natural question then is the following. Suppose we are given a set of agents and objects, and suppose the q_j are also fixed for the objects; a problem is then specified by the preference profile for the given set of agents over the given set of objects. In this setting, when is it possible to design a strategy-proof mechanism that is always efficient and envy-free? Our next result settles this question and establishes the robustness of the probabilistic serial mechanism in this setting.

We start with an observation about strategy-proofness in this setting. This result, whose proof is omitted here, was proved by Cho [4].

Claim 1. *Consider a preference profile $P = (P_1, P_2, \ldots, P_n)$, and suppose P'_i is obtained from P_i by swapping some pair of adjacent objects in P_i, and suppose $P' = (P'_i, P_{-i})$. Let A_i and A'_i be the allocations to agent i when a mechanism π is applied to profiles P and P' respectively. Then π is strategy-proof if and only if $A_i >_{P_i} A'_i$ for all i, P_i, P'_i, P_{-i}.*

Theorem 1. *Suppose there are n agents (each with unit demand) and m objects available in quantities $q_1, q_2 \ldots, q_m$. Then there is a strategy-proof mechanism that is efficient and envy-free if and only if*

$$q_a + \frac{q_b}{n} > 1$$

for every pair of distinct objects a and b. Furthermore, whenever this condition is satisfied, the probabilistic serial mechanism is strategy-proof.

Proof. Assume there is a distinct pair of objects a and b violating the condition of the claim, i.e., $q_a + q_b/n \leq 1$. We proceed as in the proof of Proposition 1. Consider the preference profile $P = (P_1, \ldots, P_n)$, where $P_1 = (a, b, \ldots)$ and $P_i = (b, c, \ldots, a)$ for each $i \neq 1$. Efficiency requires that object a be allocated to agent 1 completely. Strategy-proofness requires that agent 1 be allocated at least q_b/n of object b: otherwise agent 1 could report $(b, a, \ldots)$ as his preference ordering and get q_b/n units of b (by envy-freeness) and q_a units of a (by efficiency). Now, consider what happens when agent 2 reports $(b, a, \ldots)$ as his preference ordering. Strategy-proofness implies agent 2 receives exactly q_b/n units of b; efficiency requires that agent 1 receive all of a; and in that case, 2 will envy agent 1. This argument also highlights the key role the condition plays in obtaining the impossibility: if $q_a + q_b/n > 1$, agent 1 will not be allocated q_b/n units of b (as his demand is only 1), in which case there is no violation of envy-freeness. This establishes the necessity of the stated condition.

To prove its sufficiency, it is enough to show that the probabilistic serial mechanism is strategy-proof whenever $q_a + q_b/n > 1$ for every distinct pair of objects a and b (by earlier results of Bogomolnaia & Moulin [3] and Schulman & Vazirani [8], we know that the probabilistic serial mechanism is envy-free and efficient for every $(q_1, q_2, \ldots, q_m)$). And by Claim 1, it is enough to consider the case when some agent i swaps a pair of objects a and b that are adjacent in his preference ordering.

Fix the preference profile P_{-i} for all the agents except agent i. Suppose agent i ranks a immediately before b in P_i, and swaps these two objects in P_i'. Let A and A' be the allocations for the profiles $P = (P_i, P_{-i})$ and $P' = (P_i', P_{-i})$ respectively. It is clear that the probabilistic serial algorithm makes the same choices under both P and P' until agent i's best remaining object is a or b. Let t be the first epoch at which this event happens (note that t is the same for both problems P and P'). If object a or object b is completely consumed by time t, there is no change in the final allocation, so we may assume that both a and b are available at time t. In this case, we shall show that $a_{ia} > a'_{ia}$; that is, by delaying the request for object a, agent i will obtain a smaller quantity of it.

If $a_{ij} = 0$ for all objects j that agent i ranks worse than a, the result is immediate: $a_{i\ell} = a'_{i\ell}$ for all $\ell >_{P_i} a$, and

$$a_{ia} = 1 - \sum_{\ell:\ell >_{P_i} a} a_{i\ell} \; = \; 1 - \sum_{\ell:\ell >_{P_i} a} a'_{i\ell} \; = \; 1 - \sum_{\ell:\ell >_{P_i'} a} a'_{i\ell} + a'_{ib} \; > \; 1 - \sum_{\ell:\ell >_{P_i'} a} a'_{i\ell} \; \geq \; a'_{ia}.$$

Here we make use of the fact that at time t a positive fraction of both a and b are available, which implies $a'_{ib} > 0$.

Suppose agent i is the only agent to whom object a is allocated in P. Then $a_{ia} = q_a$ (and $q_a \leq 1$ necessarily in this case). We note that this can happen only if a is the top ranked object of agent i in P—otherwise some other object ℓ is the top ranked object, in which case i will be assigned at least q_ℓ/n units of that object, and $q_\ell/n + q_a > 1$ by the condition of the claim. Therefore, when i's preference changes to P', b will be the top-ranked object for agent i. So $a'_{ib} \geq q_b/n$; the condition $q_b/n + q_a > 1$ implies then that $a'_{ia} \leq 1 - q_b/n < q_a$.

The only remaining case is when agent i is not the only agent to whom object a is allocated, and that agent i is allocated a positive fraction of some object that he ranks worse than a. When the profile is P, agent i consumes a in the interval $[t, t_a)$ where t_a is the epoch at which a is completely consumed by the agents in P. By assumption $t > 0$ and $t_a < 1$, and that there is at least one other agent consuming a in the interval $[t', t_a)$ for some $t' \in (t, t_a)$. When P_i changes to P_i', agent i will spend some positive fraction of time consuming b; this implies that every agent other than i will request a (weakly) earlier in P' than in P. Moreover, agent i starts consuming a strictly after time t in problem P'. These two observations imply that $a_{ia} > a'_{ia}$ as required. $\square$

3 Strict Preferences: An Alternative Mechanism

We now address one of the open questions posed by Schulman and Vazirani [8]. We start with a description of the *synchronized greedy mechanism* (or simply, the serial rule) when agent demands are not necessarily identical. Suppose the demand of agent i is r_i and the supply of object j is q_j and suppose $\sum_i r_i = \sum_j q_j$. The serial rule starts at time 0 for all agents in parallel. Each agent consumes his top-ranked object at rate r_i. An object may be simultaneously consumed by several agents. This continues until some object is completely consumed. Then instantaneously any agent i whose most-preferred object ran out switches to his next best available object, and begins consuming it at the same rate r_i. In general, whenever an object runs out, all the agents who were in the process of consuming it switch to their next most-preferred object that has not yet been fully allocated. Each agent i is, at any time, consuming exactly one object, at rate r_i. At time 1, all the agents simultaneously complete their full allocation. Schulman and Vazirani show that this mechanism is efficient, strategy-proof if $\max_i r_i \leq \min_j q_j$, and envy-free whenever $r_i = r$ for all agents i.

We next define *bounded invariance*. A mechanism satisfies the bounded invariance property if for any agent i and any object j, any change in agent i's preferences involving only objects he ranks below j does not modify the amount of object j each agent obtains. Formally:

Definition 5 (Bounded Invariance). *For each agent i, each object j, let P_i and P_i' be two preference orderings which only differ in the order of the objects that i likes less than j. For each profile of preferences $P_{-i} = (P_1, P_2, \ldots, P_{i-1}, P_{i+1}, P_{i+2}, \ldots, P_n)$, let A and A' be the allocations obtained when the mechanism is applied to the profiles (P_i, P_{-i}) and (P_i', P_{-i}) respectively. A mechanism satisfies* bounded invariance *if for all agents i, preferences P_i, P_{-i}, objects j, and P_i', $a_{kj} = a_{kj}'$ for all agents k.*

We now turn to two natural questions. First, is the SG mechanism the only mechanism that is efficient, strategy-proof and envy-free whenever $\max_i r_i \leq \min_j q_j$, and $r_i = r$ for all agents i?[6] Second, Bogomolnaia and Heo [2] showed that, for the *sd*-preference model, the serial rule is characterized by efficiency, envy-freeness, and bounded invariance. Schulman and Vazirani asked if this result was true for the lexicographic preference model. In Algorithm 1 we present a strategy-proof mechanism that is both efficient and envy-free under the required conditions. Furthermore, we argue that the mechanism in Algorithm 1 also satisfies the bounded invariance property, thus answering both questions in the negative.

The Hybrid mechanism, shown in Algorithm 1, treats the most-preferred objects differently than the rest of the objects. Any object that is over-demanded—meaning supply is at most the demand—is fully allocated to the agents who demand it, and their demands are adjusted down by this amount; any agent who demands an object that is not over-demanded is allocated an amount $s \leq 1$, where s is the smallest allocation given to the agents who consume over-demanded objects. If every agent's demand is met, the algorithm stops; otherwise, the SG mechanism is applied to the residual problem.

[6] Note that, by normalizing, we may assume without loss of generality that $r_i = 1$ $\forall 1 \leq i \leq n$.

Algorithm 1 A Hybrid Mechanism

Input: Agents' strict preference lists, $r_i = 1$ for all agents, $q_j \geq 1$ for all objects.

Output: A lex-efficient, envy-free, strategy-proof allocation.

Step 1) (*greedy (partial) allocation*) For each object j, let S_j be the set of agents who have j as a top choice and $s = \min_j \frac{q_j}{|S_j|}$. Assign each object j as follows:

(a) if $q_j \leq |S_j|$, assign $\frac{q_j}{|S_j|}$ units of j to each $i \in S_j$.

(b) otherwise, assign s units of j to each $i \in S_j$

Update the remaining capacities of the agents and quantities of the objects accordingly. Remove from the problem all agents who have completed their allocation and all objects that have been fully assigned.

Step 2) Complete the allocation using the synchronized greedy mechanism.

Theorem 2. *The Hybrid Mechanism is efficient, envy-free and strategy-proof whenever $r_i = 1$ for all i and $q_j \geq 1$ for all j.*

Before we present the proof of Theorem 2, we argue that both the Hybrid mechanism and the SG mechanism satisfy bounded invariance. Suppose we apply the Hybrid mechanism to the profiles $P = (P_i, P_{-i})$ and $P' = (P'_i, P_{-i})$ as in the definition of bounded invariance. It is clear that the mechanism makes the same choices under both P and P' at least until object j is fully allocated, as both preference profiles agree for all objects that i likes at least as much as j. Therefore, object j will be identically allocated in both problems. The same argument can be used to show that the SG mechanism satisfies bounded invariance.

Remark. One fact worth noting is that the SG mechanism (equivalently, the serial rule) is *sd*-envyfree, which means that the allocation of an agent stochastically dominates the allocation of any other agent. Trivially, *sd*-envyfreeness implies envy-freeness in the lexicographic sense. Furthermore, it is easy to see that the Hybrid mechanism is *not sd*-envyfree. Indeed, Bogomolnaia [1] shows that the SG mechanism is the only one satisfying efficiency, *sd*-envyfreeness, and bounded invariance.

We conclude this section by presenting the proof of Theorem 2.

Proof of Theorem 2. For each agent i and each object j, let r'_i be the remaining requirements of agent i and q'_j be the remaining quantity available of object j after Step 1 is completed. As all agents with $r'_i = 0$ and all objects with $q'_j = 0$ are removed from the problem, we will just focus on the agents and objects with $r'_i > 0$ or $q'_j > 0$ respectively.

Recall that $\sum_j q_j = n$ and each $q_j \geq 1$, so $0 < s \leq 1$ in Step 1 of the algorithm. If in fact, when $s = 1$, every agent is allocated one unit of his most-preferred object, and the Theorem follows trivially. Therefore we will assume that $0 < s < 1$ in Step 1 of the algorithm.

Efficiency. Efficiency of the Hybrid mechanism follows from a simple observation that it is a particular case of the "extended" SG mechanism of [8], in which each agent i has an arbitrary speed function $\eta_i(\cdot) : [0,1] \to \mathbb{R}_+$ such that $\int_0^1 \eta_i(t)dt = 1$. The speed function of an agent i in the Hybrid mechanism depends on whether i's most-preferred object j is over-demanded (meaning $q_j \leq |S_j|$) or not (meaning $q_j > |S_j|$). In the latter case, the speed function is 1 over the interval $[0,1]$; in the former case, the speed function is $\frac{q_j}{s|S_j|}$ in the interval $[0,s]$, and $\frac{1-q_j/|S_j|}{1-s}$ in the interval $(s,1]$. Note that the speed function of the agents whose most-preferred object is under-demanded is at least 1 initially, and is at most one in the interval $[s,1]$. Schulman and Vazirani [8] showed that *any* mechanism in this class is efficient, and this implies the efficiency of the Hybrid mechanism.

Envy-freeness. Suppose agent i has j as his most-preferred object. We first show that the only agents that i potentially envies are those in S_j, i.e., an agent will never envy the allocation of someone who did not obtain the same object as him in Step 1 of the algorithm. If $q_j \leq |S_j|$, this follows trivially as all agents outside S_j will get zero units of j, while i gets a positive amount. On the other hand, if $q_j > |S_j|$, each agent in S_j obtains s units of j in Step 1 and will start Step 2 with j as his top choice. In Step 2, the "eating speed" of the agents in S_j is at least as much as the "eating speed" of any other agent in the problem. Therefore, the agents who are in S_j will be allocated at least s extra units of j than those who are outside of S_j. Thus, the only agents that i potentially envies are those in S_j. However, i cannot envy any agent $i' \in S_j$ because their eating speeds are identical throughout.

Strategy-Proofness. Finally, we show that the mechanism is strategy-proof whenever $\min q_j \geq 1$. Recall that, by Claim 1, it is sufficient to show that no agent can benefit by swapping adjacent objects in his preference profile.

Fix an agent i, and consider two preference profiles P and P' that differ only in i's report. Suppose also that a is ranked just above b in P, whereas their order is swapped in P'.

We first consider the case when a is *not* the most-preferred object of agent i. In this case, the outcome of Step 1 of the Hybrid mechanism is the same for profiles P and P'; Step 2 is simply an application of the SG mechanism to the problem that remains after Step 1, and so the result would follow if we can show that the SG mechanism is strategy-proof on this resulting instance. By the result of Schulman and Vazirani [8], it is enough to show that $\max_{i:r'_i>0} r'_i \leq \min_{j:q'_j>0} q'_j$, where r'_i and q'_j are the residual demand and supply of agent i and object j respectively. First, note that all agents obtain at least s units of an object in Step 1 and this bound is tight, hence $\max_{i:r'_i>0} r'_i = 1-s$. Second, for all objects j with $S_j = \emptyset$, $q'_j = q_j \geq 1 > 1-s$. If $q_j \leq |S_j|$, object j is fully allocated in Step 1 and is removed from the problem. The only objects that remain are those with $q_j > |S_j|$ and $|S_j| \geq 1$. For such objects, exactly s units of it are assigned to each agent in S_j, so the remaining quantity $q'_j = q_j - s|S_j| > (1-s)|S_j| \geq (1-s)$. Thus, $\max_{i:r'_i>0} r'_i \leq \min_{j:q'_j>0} q'_j$ as required.

We now consider the case when a is the most-preferred object of agent i and b is his second most-preferred object. Let α (resp. β) be the amount of a agent 1 gets when he reports a (resp. b) as his top-choice. We prove the result by showing that $\alpha > \beta$. Define $S'_a = S_a \backslash \{i\}$, $S'_b = S_b \cup \{i\}$, $S'_j = S_j$ for all objects $j \neq a, b$ and let $s' = \min_j \frac{q_j}{|S'_j|}$.

Case 1: $q_a \leq |S_a| - 1$**.** In this case a will be fully distributed in Step 1 of the Hybrid mechanism, in both P and P'. Thus $\alpha = q_a/|S_a| > 0$, whereas $\beta = 0$.

Case 2: $|S_a| - 1 < q_a \leq |S_a|$**.** In this case $\alpha = q_a/|S_a|$. If agent i ranks b first followed by a, each agent in S'_a will get s' units of a in Step 1 and will have an eating speed $(1-s')$ in Step 2; agent i will consume only b in Step 1, and his maximum eating speed in Step 2 is $(1-s')$. So if agent i is able to obtain β units of a in Step 2, each of the agents in S'_a will obtain *at least* β units of a in Step 2 as well. As only q_a units of a are available, we must have $(s'+\beta)|S'_a| + \beta \;=\; s'|S'_a| + \beta|S_a| \leq q_a$. Dividing both sides by $|S_a|$ and observing that $s' > 0$, it follows that $\beta < \alpha$.

Case 3: $q_a > |S_a|$**.** Here, $q_a/|S_a| > 1$, and so a cannot be the object that determines s in the original problem P. As $q_j/|S'_j| \leq q_j/|S_j|$ for every $j \neq a$, we must have $s' \leq s$. Furthermore, note that β—the quantity of a that agent i obtains by misreporting—is maximized whenever i obtains s' units of b and b is fully allocated at the end of Step 1, so we will assume that is the case. In this case, $s' = q_b/|S'_b|$.

Let $F = \{j :\; j \neq b \text{ and } q_j > |S_j|\}$ and $E = \{j :\; j \neq b \text{ and } q_j \leq |S_j|\}$. The sets F and E partition the set of objects other than object b; furthermore, only the objects in F remain in the problem in Step 2, and $a \in F$. We analyze how the problem obtained in Step 2 changes because of agent i's misreport.

- For every $j \in F$, the amount of j at the beginning of Step 2 changes from $q_j - s|S_j|$ to $q_j - s'|S'_j|$, and all the agents in S'_j have eating speeds $(1-s')$ instead of $(1-s)$. In particular, for $j \neq a$, the quantity available increases by $(s-s')|S_j|$ units, whereas for object a the quantity increases by $(s-s')|S_a| + s'$. These agents all eat at a faster speed of $(1-s')$ rather than $(1-s)$.

- Object b is fully allocated, and the agents in S'_b have speed $(1-s')$. In the original problem b may or may not have been fully allocated, and the agents in S_b will have had an eating speed of at most $(1-s)$.

- An agent whose most-preferred object is in E will have the same speed in both P and P', as his allocation in Step 1 will be the same in both problems.

- Except possibly for b, only the objects in F remain in Step 2, in both P and P'.

For each object $j \in F$ (including a), let T_j be the epoch at which j is fully allocated in problem P and T'_j be the epoch at which it is fully allocated in P'. Then, $\beta = (1-s')T'_a$ and $\alpha = (1-s)T_a + s$. To show that the mechanism is strategy-proof, we need to show that $\beta < \alpha$, or equivalently:

$$T'_a \quad < \quad \frac{(1-s)}{(1-s')}T_a + \frac{s}{(1-s')} \quad = \quad T_a + \frac{(s-s')}{(1-s')}(1-T_a) + \frac{s'}{(1-s')}.$$

Let $\mathcal{T} = \frac{(1-s)}{(1-s')}T_a + \frac{s}{(1-s')}$, and let X (resp. Y) be the total amount of a consumed by agents in S_a (resp. outside S_a) during the second step in problem P. Note that a total of $X + Y + (s-s')|S_a| + s'$ units of a are available at the beginning of the second step of P'. We shall establish that $T'_a < \mathcal{T}$ as follows: (i) Show that by time $\mathcal{T}$, the agents in S_a (including i) would have consumed *more* than $X + (s-s')|S_a| + s'$ units of a in P', and (ii) show that by time $\mathcal{T}$, agents outside S_a would have consumed *at least* Y units of a in P'. Together, these two statements imply that a will be exhausted by time $\mathcal{T}$, and therefore $T'_a < \mathcal{T}$ as desired.

To show (i), note that the $X = |S_a|(1-s)T_a$ units of a consumed by the agents in S_a during the second step of P can now be consumed by those agents at a faster rate in exactly $\frac{(1-s)}{(1-s')}T_a$ units of time. The additional $(s-s')|S_a| + s'$ units can be consumed in $\frac{(s-s')|S_a|+s'}{(1-s')|S_a|} < \frac{s}{(1-s')}$ units of time, as $|S_a| \geq 2$. This proves statement (i).

To prove (ii), we show that $T'_j \leq T_j + \frac{(s-s')}{(1-s')}$ for all $T_j < T_a$, by induction on the epochs at which the objects are fully allocated in the original problem. Suppose $T_{j_1} \leq T_{j_2} \leq \ldots \leq T_a$. If $j_1 \in E$ or $j_1 = b$, the result is trivially true, as these objects are fully allocated (weakly) *earlier* in P' than in P. If $j_1 \in F$ and $|S_{j_1}| = \emptyset$, the same quantity of j_1 is available in both problems, and as each agent consumes weakly faster in P' than in P, $T'_{j_1} \leq T_{j_1} < T_{j_1} + \frac{(s-s')}{(1-s')}$ and the result follows. Finally, if $j_1 \in F$ and $|S_{j_1}| \neq \emptyset$, note that $(s-s')|S_{j_1}|$ additional units available in P' than in P. Every agent who has an object in E as his top-choice and j_1 as his second choice has the same eating-speed in both problems, and therefore will consume the same amount of j_1 by time T_{j_1}. Moreover, at least $|S_{j_1}|$ agents will obtain j_1 at rate $(1-s')$ in P', and can collectively consume the same number of units as before plus the additional $(s-s')|S_{j_1}|$ by time $T_{j_1} + \frac{(s-s')}{(1-s')}$. Therefore, it follows that by time $T_{j_1} + \frac{(s-s')}{(1-s')}$ the object j_1 must be fully consumed.

Now suppose this inequality holds for every object up to j_{k-1} and consider j_k. Each agent can start to consume j_k at most $\frac{(s-s')}{(1-s')}$ units later in P' than in P. The later the agents start consuming this object, the longer it will remain in the problem. Therefore, the worst case is that every agent starts consuming j_k exactly $\frac{(s-s')}{(1-s')}$ units of time later in P' than in P. If the same amount of object j_k is available in both P and P', $T'_{j_k} \leq T_{j_k} + \frac{(s-s')}{(1-s')}$follows by observing that the eating speed of any agent in P' is (weakly) greater than his eating speed in P. If j_k is available in different amounts, the result is obtained by the above argument plus the fact that the additional amount of $(s-s')|S_{j_k}|$ will run out before time $T_{j_k} + \frac{(s-s')}{(1-s')}$.

The above result implies that every agent outside S_a will start consuming a *at most* $\frac{(s-s')}{(1-s')}$ units of time later in P' than in P. As the eating speeds of the agents are (weakly) greater in P', by time $T_a + \frac{(s-s')}{(1-s')} \leq T_a + \frac{s'}{(1-s')} < \mathcal{T}$ they would have collectively consumed *at least*

Y units of a, where the first inequality follows from the following observation: by definition $s' = \frac{q_b}{|S'_b|} = \frac{q_b}{|S_b|+1}$, and $s \leq \frac{q_b}{|S_b|}$ plus $|S_b| \geq 1$ implies $s' \geq (s - s')$.

□

4 An impossibility result for the full domain

In this section we consider a version of the problem when agents may be indifferent between objects. When preferences are strict and when the q_j are all 1, the probabilistic serial mechanism is envy-free, efficient, and strategy-proof. When indifferences are permitted, there is a natural generalization of the probabilistic serial mechanism (see Katta and Sethuraman [7]) that is always envy-free and efficient; however, it is easy to show that this mechanism is not strategy-proof. We next show that *no* mechanism can be efficient, envy-free and strategy-proof, even for the special case of unit q_j.

Theorem 3. *Suppose there are at least 6 agents and 6 objects and suppose $r_i = q_j = 1$ for all i and j. If the agents are allowed to report indifferences, then no strategy-proof mechanism can be both efficient and envy-free.*

Proof. We start by proving the claim when there are exactly 6 agents and 6 objects. Consider the following preference profile:

1	a	{b,c,d}	...			
2	a	b	c	d	e	f
3	a	b	c	e	f	d
⋮	⋮					⋮
6	a	b	c	e	f	d

Agents 3 through 6 have identical preferences. Only agent 1, who likes objects b, c, and d equally—reports indifferences. As with earlier proofs, we explicitly list the implications of efficiency, envy-freeness and strategy-proofness.

First, by envy-freeness, each agent must receive exactly 1/6 units of object a as it is the unique top-ranked object for all agents.

Second, by efficiency, $a_{1b} = a_{1c} = 0$. For otherwise, agent 1 is allocated a positive amount of b or c, and some other agent j has a positive amount of d; agents 1 and j can then engage in a Pareto-improving exchange, contradicting efficiency. Therefore, by envy-freeness, agents 2 through 6 each receive 1/5 units of b and 1/5 units of c.

Third, by strategy-proofness, 1 must be allocated 5/6 units of d. To see this, consider what happens when 1 reports the ordering a, b, c, d, e, f (the same as agent 2). In this case, by envy-freeness, objects a, b, and c will each be shared equally by all the agents; object d, which is the last choice of agents 3 through 6 cannot be assigned to them, and so will be shared equally by agents 1 and 2 because of efficiency considerations. Thus agent 1 will be assigned 1/6 of a, 1/6 of b, 1/6 of c, and 1/2 of d under the new report. So agent 1 must be assigned 5/6 units of d for the mechanism to be strategy-proof. This, in turn, implies that agent 2 is assigned exactly 1/6 units of d (anything less will be inefficient). As $\frac{1}{6} + \frac{1}{5} + \frac{1}{5} + \frac{1}{6} < 1$, agent 2 must own some of either e or f. In fact we can say something stronger: by strategy-proofness, agent 2 must obtain exactly 1/5 units of e. Why? If 2 reports his preference ordering to be a, b, c, e, d, f (i.e. if 2 swaps e and d in his preference ordering), he obtains exactly 1/5 of object e. By envy-freeness, he cannot obtain more than 1/5—recall that agents 2-6 get identical shares of a, b, and c; if 2 is assigned more than 1/5 of e, at least one of the agents 3-6 has to be assigned less than 1/5 of e resulting in that agent envying agent 2).

Finally, consider what happens when agent 3 reports the ordering a, b, c, e, d, f. Strategy-proofness implies $a_{3e} = 1/5$ as otherwise agent 3 can manipulate at one of the profiles. In addition, by the same arguments as before, agent 1 must still be allocated 5/6 units of d. Envy-freeness then implies that $a_{4e} = a_{5e} = a_{6e} = 1/5$ which then forces $a_{2e} = 1/5$ as well. We now look at agent 2's allocation when 3 reports this new ordering. If $a_{2d} = 1/6$ as before then 3 envies 2; if $a_{2d} < 1/6$, then efficiency is violated: at least one of the agents 3 through 6 is allocated a positive amount of d who can give that up for an equal amount of e. This establishes the claim for the case of 6 agents and 6 objects.

Suppose there are $n > 6$ agents, and let G be the set of all objects with $|G| = n$. Partition G into $G_1 = \{a, b, c, d, e, f\}$, $G_2 = G \backslash G_1$, and consider the following class of preference profiles:

1	a	{b,c,d}	...				G_2
2	a	b	c	d	e	f	G_2
3	a	b	c	e	f	d	G_2
⋮	⋮						⋮
6	a	b	c	e	f	d	G_2
7	G_2	...					G_1
⋮	⋮						⋮
n	G_2	...					G_1

Agents 1 through 6 have the same preferences over the objects in G_1 as in the earlier proof; each of these agents arbitrarily order the objects in G_2 after the objects in G_1. Each agent $i = 7, 8, \ldots, n$ has arbitrarily order the objects in G_2 before ordering the objects in G_1. Efficiency dictates that the objects in G_1 be allocated to the agents 1 through 6, and that the objects in G_2 to the remaining agents. □

5 Discussion

Our results shed additional insight into allocation mechanisms when agents have lexicographic preferences. Specifically, we show that envy-freeness, efficiency, and strategy-proofness do not characterize the serial rule if agents have lexicographic preferences, even in the restricted setting of unit demands and supplies. In addition, we show that the rule is not characterized by envy-freeness, efficiency and a property called bounded invariance. This is in stark contrast to the standard setting when agents have "sd-preferences": there, efficiency, envy-freeness and bounded-invariance, characterize the serial mechanism. An interesting open question is to characterize the serial rule by natural additional properties, or by a different combination of properties.

When agents have "sd-preferences" the generalization of the serial rule described by Katta and Sethuraman [7] finds an allocation that is envy-free and efficient. There, they showed that the extended serial rule fails strategy-proofness, but so does every rule that is efficient and envy-free. We develop a parallel result, but for the lexicographic preference model. Specifically, we show that when the agents are allowed to be indifferent between objects, no mechanism can satisfy all three properties, even when $r_i \equiv 1$ and $q_j \equiv 1$. It remains an interesting open question to characterize the problems for which it is possible to design an envy-free, efficient, and strategy-proof mechanism when indifferences are permitted.

Acknowledgements

We are grateful to two anonymous referees for their detailed comments on an earlier version of the paper.

References

[1] Anna Bogomolnaia. Random assignment: Redefining the serial rule. 2012.

[2] Anna Bogomolnaia and Eun J. Heo. Probabilistic assignment of objects: Characterizing the serial rule. *J. Economic Theory*, 147(5):2072–2082, 2012.

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

[4] Wonki J. Cho. Probabilistic assignment: A two-fold axiomatic approach. 2012.

[5] M. Hausner. Multidimensional utilities. In R. L. Davis R. M. Thrall, C. H. Coombs, editor, *Decision Processes*, pages 167–180. Wiley, 1954.

[6] Eun J. Heo. Probabilistic assignment with multiple demands: a generalization of the serial rule and a characterization. 2010.

[7] Akshay-Kumar Katta and Jay Sethuraman. A solution to the random assignment problem on the full preference domain. *J. Economic Theory*, 131(1):231–250, 2006.

[8] Leonard J. Schulman and Vijay V. Vazirani. Allocation of divisible goods under lexicographic preferences. *CoRR*, abs/1206.4366, 2012.

[9] Tayfun Sonmez and Utku Ünver. Matching, allocation, and exchange of discrete resources. In Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, editors, *Handbook of Social Economics, Volume 1A*, pages 781–852. North Holland, 2010.