Random Assignment: Redefining the Serial Rule

Anna Bogomolnaia*

January 2015

Abstract

We provide a new, welfarist, interpretation of the well-known Serial rule in the random assignment problem, strikingly different from previous attempts to define or axiomatically characterize this rule.

For each agent i we define $t_i(k)$ to be the total share of objects from her first k indifference classes this agent i gets. Serial assignment is shown to be the unique one which leximin maximizes the vector of all such shares $(t_i(k))$.

This result is very general; it applies to non-strict preferences, and/or non-integer quantities of objects, as well.

In addition, we characteize Serial rule as the unique one sd-efficient, sd-envy-free, and strategy-proof on the lexicograpic preferences extention to lotteries.

Keywords: Random assignment, Serial Rule, Leximin

We consider the random assignment problem without money. A number of indivisible objects are to be distributed among several agents, and each agent is entitled to one object only. Agents may have arbitrary preferences over objects. In the absence of transfers, allowing for randomization restores fairness.

A random assignment can be represented by a matrix P of probabilities for each agent to get each object, with rows standing for agents and columns standing for objects. Here p_{ia} is the probability for agent i to get object a.[1]

An alternative interpretation is available, leading to the same formal structure. One can assume that objects are infinitely divisible, and each agent is entitled to exactly one unit total of objects. An example would be sharing a workload between employees, each of whom owes the firm a fixed number of work hours. We obtain the same set of feasible assignments, represented by matrices P, with p_{ia} now interpreted as the share of object a allocated to agent i.[2]

*University of Glasgow, UK; e-mail: anna.bogomolnaia@glasgow.ac.uk

[1] For each agent, her ordinal ranking of objects defines an incomplete ordering of random assignments, based on stochastic dominance. We do not need to specify full preferences. A benchmark example is when agents evaluate lotteries based on expected utility.

[2] The expected utility assumption in the first model paralllels here to the assumption that agent's utility from an object is linear in the quantity she receives, and additive across the objects.

An assignment rule systematically determines a random assignment, based on reported agents' preferences. Earlier literature assumed agents report their cardinal utilities and compare lotteries by means of expected utility. Having in mind practical implementation issues, the extensive recent literature, starting from [1] (2001), concentrated on the ordinal mechanisms, when agents are only asked about their ordinal rankings of objects. Particular attention was paid to the Serial rule, which possesses many attractive properties. It was introduced in [1] (2001) for the strict preference domain, and extended in [8] (2006) to the full ordinal domain.

On the strict domain, the Serial rule is rather intuitive. Its allocation is constructed by allowing agents to "eat" shares of objects at the same unit speed, over the time interval $[0, 1]$, in decreasing order of their preferences; at any moment each agent is eating from the object she most prefers among those still available. However, when indifferences are allowed, guaranteeing efficiency becomes a non-trivial task. The Serial assignment is then constructed through a rather cumbersome multistep algorithm, involving repeated computation of the maximal flow in a network, and the exact nature of the final allocation becomes difficult to grasp.

The main purpose of this paper is to provide a welfarist re-definition of the Serial rule. It is simple and intuitive, and applies to strict as well as non strict preferences. It works equally well when objects are available in arbitrary quantities[3] (as long as each agent is entitled to exactly one unit total). The Serial assignment is shown to be the unique one which leximin maximizes the vector of total shares of objects, consumed by our agents, above their respective different indifference thresholds.

This alternative definition may also be thought of as an axiomatic characterization of the Serial rule, based on a single, entirely new principle.

We evaluate any given assignment by the list of numbers $t_i(k)$, the total shares agent i gets of objects from her first k indifference classes, for all i and k. We show that the Serial assignment is the unique (utility wise) leximin maximizer of the vector $(t_i(k))_{i,k}$ over all feasible assignments.

Note a somewhat counter-intuitive feature: in the list $(t_i(k))_{i,k}$, an agent j with few indifference classes is represented by less numbers, then another agent l whose preferences fully rank all objects. However, a parallel characterization is available, where the length of the vector t does not depend on the number of agents' indifference classes. Define $t'_i(a)$ to be the total share of objects at least as good as a agent i gets. The Serial assignment is also the unique leximin maximizer of the $(t'_i(a))_{i,a}$.

A welfarist nature of our definition can be illustrated by the following interpretation. Split each agent in as many "sub-agents" as the number of indifference classes in her preferences. Agent i's first sub-agent only cares about her first indifferent class, her second sub-agent only cares about her two top indifferent classes, etc. Thus, the utility of agent i's k-th sub-agent is measured

[3] This is a natural assumption for the model with infinitely divisible objects. See also [11] (2012).

by the total amount of objects she gets from her first k indifference classes. Our result is that Serial rule maximizes the leximin ("Rawlsian") collective utility of those sub-agents. It first attempts to maximize the utility of the worst-off sub-agent, then the utility of the second worst-off one, and so on.

Rephrasing, the Serial rule is the most egalitarian (the Rawlsian maximizer) in attempting to equalize agents' shares of top ranked objects (i.e. of upper counter sets of objects) under different cutoffs. Recall that agents only report rankings of objects, not their relative valuations, so equalizing allocated shares for different upper counter sets seems to be the best available instrument for an egalitarian mechanism designer.

Note that, while the vector $(t_i(k))_{i,k}$ for Serial assignment is leximin preferred to one for any other assignment, it is not Lorenz dominant vector (in fact, for a generic preference profile, there is no assignment with Lorenz dominant vector t). The only exception is the dichotomous preference domain, as discussed in [2] (2004) and [8] (2006).

In addition to our main result, we also present a yet another characterization of Serial rule. While it is more in line with recent literature (and is heavily based on it), it is nevertheless the first one which singles out Serial rule by properties of (sd) efficiency, fairness (sd no-envy), and strategy-proofness (for the lexicographic extension of deterministic rankings to the domain of lotteries).

An axiomatic characterization of the Serial rule was elusive for a decade. Recently, several characterization results, similar in nature, were obtained. In [6] (2011), [5] (2012), [3] (2012), [7] (2012) (this last paper deals with the full domain) the Serial rule is characterized by "sd" (first order stochastic-dominance based) efficiency and envy-freeness, together with an invariance axiom, and, for the case of non-unit demand, a consistency property (see also [9] (2011) and [4] (2011)). For the profiles where all preference orderings are present, [10] (2011) notes that the only sd efficient and envy-free assignment is given by Serial rule. [4] (2011) proposes a characterization, for the case when not all objects are acceptable, partly based on the axiom they call "Rawlsian Criterion". It requires separability in that the rule should distribute "the worst still acceptable object" separately from the rest (which is the main driving force of this axiom), and partial egalitarianism in that this object should be assigned so as to maximize the total objects' share of the least served agent.

While our result can be also interpreted as an axiomatic characterization of the Serial rule, it is based on completely different ideas.

Papers [6] (2011) and [3] (2012) introduce a representation of arbitrary random assignments by "consumption processes" over time. Finally, [11] (2012) shows that the Serial rule is strategy-proof under the lexicographic preferences extension to the lotteries.

1 Model and Results

Let $N = \{1, ..., n\}$ be the set of agents, and $A = \{a_1, ..., a_m\}$ be the set of objects.[4] Let $\mathcal{R}$ be the set of all orderings over A. Each $i \in N$ has preferences $R_i \in \mathcal{R}$ over A, and a preference profile is $R = (R_i)_{i \in N}$. Given a particular R, we denote by $U_i(a) = \{b \in A : bR_ia\}$ the upper contour set over a under preferences R_i. We write U_i^k for the set of objects which are in the top k indifference classes for preferences R_i. We use the notation $M_i(B)$ for the set of objects which are the most preferred in $B \subset A$ by preferences R_i, and $M_I(B) = \bigcup\limits_{i \in I} M_i(B)$ for the set of objects which are the most preferred in $B \subset A$ by at least one of the agents in I.

The set $\mathcal{P}$ of feasible assignments consists of $|N| \times |A|$ matrices $P = (p_{ia})_{i \in N, a \in A}$, where p_{ia} is interpreted as the probability for agent i to get object a or as the share of object a assigned to agent i.[5] We denote by P_i its row corresponding to agent i, and by P^a its column corresponding to object a.

An assignment rule is a correspondence $f : \mathcal{R}^n \rightrightarrows \mathcal{P}$, which is essentially single-valued. Specifically, $f(R) \subset P$, $f(R) \neq \varnothing$, and all agents are indifferent between any two $P, P' \in f(R)$. This last requirement means $\sum\limits_{a \in U^k(R_i)} p_{ia} = \sum\limits_{a \in U^k(R_i)} p'_{ia}$ for all i and k. Restricted to the strict domain, an assignment rule is always a (single-valued) function.

Fix a preference profile R. For an arbitrary assignment P we define $t_i(a) = t_i^P(a) = \sum\limits_{b \in U_i(a)} p_{ib}$, the total share of objects, at least as good as a, which agent i receives under P.

We can interpret any assignment P as the result of "consumption process" over time interval $[0, 1]$. Each agent i "consumes" shares of objects at unit speed and up to quotas assigned to her by P, in decreasing order of her preferences. Since P specifies how much of each object an agent can consume, during this process she might be often forced to leave her preferred object when it is still available. We fix an arbitrary ordering $a_1, ..., a_m$ of objects, and assume that, when an agent is indifferent between several objects, she consumes first those which come earlier in this ordering[6]. Let $P[t]$ be the partial assignment obtained by this process by time t. The time $t_i^P(a)$ will be exactly the last moment in this process when agent i consumes objects at least as good for her as a.[7]

For each agent, we define a vector $t_i = (t_i(1), ..., t_i(K_i))$, where K_i is the number of indifference classes in R_i. Here $t_i(k) = \sum\limits_{b \in U_i^k} p_{ib}$ is the total share

[4] For the clarity of presentation, assume that each object has quota 1 and is valued above getting nothing ("null object"). Those assumptions can be relaxed — see the comment after the proof.

[5] Birkoff theorem tells us that the convex hull of zero-one bistochastic matrices (deterministic assignments) contains all bistochastic matrices. Thus, such matrices P are exactly those which can be represented as lotteries over deterministic assignments.

[6] Alternatively, we could consider the set of all utility-equivalent consumption processes. In this case, we do not need to specify in which order an agent consumes her equivalent objects.

[7] Note that $P[t]$ is continuous in t on $[0, 1]$.

agent i gets of objects from her first k indifference classes, or the moment at which she stops consuming such objects in consumption process $P[t]$.[8]

When preference profile R is strict, the Serial rule assignment $S = S(R)$ is defined as the result of the consumption process where any agent at any moment consumes from the best for her good among those not yet exhausted. For an arbitrary preference profile R, we extend the notion of Serial assignment $S(R)$ (a subset of $\mathcal{P}$), based on the same principle. We propose the following, rather informal, construction to do it[9].

We start from the set of all feasible assignments, and then repeatedly shrink this set, each time specifying the allocation of some objects. At each step, we only care about distributing to each agent the best for her among still available objects; and we pursue the egalitarian goal to guarantee the best treatment for the worst off agents.

At step 1, we discard all assignments which do not maximize $\min_i t_i(1)$. I.e., we only keep the assignments under which the agent who receives the smallest share of her top indifference class is treated the best (assignments for which the smallest share of top objects is the largest possible). It is easy to see, that all those assignments fully distribute the same set B_1 of "bottleneck" (the most demanded) objects, and to the same agents, those who are the worst off; each such agent gets share r_1 of her top objects. We are left with a reduced set of objects $A_2 = A \backslash B_1$ to allocate.

At each step k, we start with previously reduced set of assignments and reduced set A_k of objects. We then only keep assignments which maximize $\min_i t_i(k_i)$, where k_i is the best for agent i indifference class in which some objects are still available (i.e., are in A_k). In other words, given an assignment, for each agent i we calculate $t_i(k_i)$ — the total share of objects, at least a good for her as her best objects in A_k. We then only keep assignments which maximize the total share $t_i(k_i)$ of the worst treated agent i. Again (see the algorithm below), all those assignments fully distribute the same set B_k of bottleneck objects to the same group of the worst off agents. Each such agent i gets share r_k of objects at least as good as ones in k_i. Hence, the set of still unallocated objects is further reduced, to $A_{k+1} = A_k \backslash B_k$.

This process ends in a finite number of steps with a set of utility equivalent allocations (again, this easily follows from the algorithm below), which we shell call Serial assignment.

The standard formal definition of the Serial assignment on the full domain, used in the literature, is obtained by the following algorithm.[10]

[8] Note two different usages: $t_i^P(a)$ with $a \in A$ for the strict preferences, and $t_i^P(k)$ with an integer k for the preferences with indifferences.

[9] The only definition proposed and used in the literature so far is the algorithmic one, provided by [8] (2006), which we present below. Our description almost immediately leads to our re-definition of the Serial rule (Theorem 1), and thus can be considered as yet another version of our main result.

[10] This is a concise description of the algorithm introduced in [8] (2006), who first extended the Serial rule for the full domain. We refer the reader to their paper for the extensive discussion and illustrating examples.

(1) Set $A_1 = A$, $c(i,1) = 0$ for all $i \in N$, $k = 1$

(2) Step k: At each step k, we find the largest share λ_k such that each agent can consume at least λ_k of her most desired objects from A_k, the set of objects still available at this step. Each agent is then guaranteed the share λ_k of her preferred objects. Maximality of λ_k implies that some objects are exhausted at this point. We define A_{k+1} to be the set of not yet exhausted objects and proceed to step $k+1$.

For this purpose, we construct the following directed network, with the set of edges $A_k \cup N \cup \{s\} \cup \{t\}$. (i) From the "source" s we draw an arc of capacity $c(i,k)+\lambda$ toward each agent $i \in N$. (ii) From each agent $i \in N$ we draw an arc of infinite capacity toward each of her best in A_k objects, those from $M_i(A_k)$. (iii) From each object in A_k we draw an arc of capacity 1 toward the "sink" t.

Here $\lambda \geq 0$ is a parameter. For each λ we find the maximal flow through this network, which can be sent from the source s to the sink t. When we continuously increase λ from zero on, there is the last moment $\lambda = \lambda_k$ after which the maximal flow is less then the total "out" capacity of the source, $\sum\limits_{i \in N} (c(i,k)+\lambda)$.

Maximal flow through a network is known to be equal to the capacity of a minimal cut. A cut is a partition of the network's nodes into $S_s \ni s$ and $S_t \ni t$, and its capacity is the sum of capacities of all arcs going from S_s to S_t. In our network, all cuts of finite capacity are such that $S_s = \{s\} \cup X \cup W$ where $X \subset N$ and $M_X(A_k) \subset W \subset A_k$. The capacity of such cut is $\sum\limits_{i \in N \setminus X} (c(i,k)+\lambda) + |W|$, so in a minimal cut it has to be $M_X(A_k) = W$. When $\lambda \leq \lambda_k$, the maximal flow is equal to the "out" capacity of the source ($S_s = \{s\}$ is a minimal cut). For $\lambda > \lambda_k$, $S_s = \{s\}$ stops to be a minimal cut. Hence (by continuity of our process in λ), when $\lambda = \lambda_k$, there is another minimal cut with $S_s = \{s\} \cup X_k$ where $X_k \subset N$, $X_k \neq \emptyset$. It has to be that $X_k \in \arg\min\limits_{X \subset N} \left(\sum\limits_{i \in N \setminus X} (c(i,k)+\lambda) + |M_{X_k}(A_k)| \right)$. If the above arg min is not a singleton, we choose X_k to be the largest set in the sense of inclusion.[11]

We hence must have $\sum\limits_{i \in N} (c(i,k)+\lambda_k) = \sum\limits_{i \in N \setminus X_k} (c(i,k)+\lambda_k) + |M_{X_k}(A_k)|$, or

$$\lambda_k = \frac{|M_{X_k}(A_k)| - \sum\limits_{i \in X_k} c(i,k)}{|X_k|} = \min_{X \subset N} \frac{|M_X(A_k)| - \sum\limits_{i \in X} c(i,k)}{|X|}.$$

Agents from X_k constitute the "bottleneck" of our algorithm at Step k. When each agent from X_k is assigned her "in" capacity $c(i,k)+\lambda_k$ of her top good, the set $B_k = M_{X_k}(A_k)$ of their top goods is completely exhausted. Maximality of X_k implies that it is feasible to give all agents from $N \setminus X_k$ shares, strictly larger then their "in" capacity $c(i,k)+\lambda_k$, of their top in A_k goods, using only goods from $A_k \setminus M_{X_k}(A_k)$.

[11] It is easy to check that a union of two such "minimal" X will also be minimal.

Given this observation, we assign to each agent $i \in X_k$ the share λ_k of her top in A_k goods. We define the set of available goods for the next step, $A_{k+1} = A_k \backslash M_{X_k}(A_k)$. Finally, we set the new capacity of agents from the bottleneck, $i \in X$, to be $c(i, k+1) = 0$, and we increase the capacity of agents, not involved in the bottleneck, $j \in N \backslash X_k$, to be $c(j, k+1) = c(j, k) + \lambda_k$.

Since at each step at least one good is exhausted, the algorithm will finish in $K \leq m$ steps, with each agent getting exactly one unit total of goods' shares, and with $\sum\limits_{1 \leq k \leq K} \lambda_k = 1$. It is easy to see that $r_k = \sum\limits_{1 \leq j \leq K} \lambda_j$ (where r_k is the share of at least as good as k_i-class objects for worst off agents i in Step k from our informal description above). Also, r_k is the time moment in the consumption process $S[t]$ when step k of the algorithm stops.

Note that there could be multiple maximal flows (and hence multiple assignments) at each step, but all of them are utility-equivalent (and any assignment, utility-equivalent to one constructed this way, can also be constructed this way).

Definition 1 The leximin order L on $\mathbb{R}^n$ is defined as follows. For any $x = (x_1, ..., x_n) \in \mathbb{R}^n$, let $x^* = (x_1^*, ..., x_n^*) \in \mathbb{R}^n$ be a permutation of the coordinates of vector x in the increasing order: $x_1^* \leq ... \leq x_n^*$. We say that xLy if there is a $j \in \{1, ..., n\}$ such that $x_j^* > y_j^*$, while $x_i^* = y_i^*$ for all $i < j$.

Theorem 1

For all preference profiles R, the Serial assignment $S(R)$ is exactly the set of feasible assignments which leximin maximize the vector[12] of shares $(t_1, ..., t_n)$.

Proof.

This statement is very intuitive on the strict domain, where the above vector coincides with the vector of $t_i(a)$, for all $i \in N$, $a \in A$. So we present the argument for strict preferences first.

Fix a strict preference profile R. Let $P \neq S$, and let $\tau \in [0, 1]$ be the last moment t such that $P[t] = S[t]$. There is an agent j and an object b, such that the object b is still available at time τ, but under the consumption process corresponding to P this agent j switches at the moment τ to some object c which she values less then b. Note that all $t_i^S(a) = t_i^P(a)$ for all i, a such that $t_i^S(a) \leq \tau$, while $t_i^P(a) > \tau$ implies $t_i^S(a) > \tau$. However, $t_j^S(b) > \tau$ while $t_j^P(b) = \tau$.

Now, fix an arbitrary preference profile R. Let $P \notin S(R)$ and let $\tau \in [0, 1]$ be the last moment t such that $P[t]$ is utility equivalent to $\{S[t] : S \in S(R)\}$. Let r be such that $\sum\limits_{1 \leq k \leq r} \lambda_k \leq \tau < \sum\limits_{1 \leq k \leq r+1} \lambda_k$. Thus, the consumption process $P[t]$ coincides utility-wise with the Serial process $S[t]$ during its first r steps (in both assigned and guaranteed shares), but deviates from the Serial assignment algorithm during its step $r + 1$.

[12] Remember that each t_i is itself a vector, $t_i = (t_i(1), ..., t_i(K_i))$, where $t_i(k) = \sum\limits_{b \in U_i^k} p_{ib}$. Hence, $(t_1, ..., t_n)$ is the vetor of length $\sum\limits_{i \in N} K_i$.

In (any) Serial consumption process $S[t]$, along the time interval $T_r = \left[\sum_{1\le k\le r} \lambda_k, \sum_{1\le k\le r+1} \lambda_k\right]$ each agent i consumes her best goods from A_{r+1}, which belong to her $l_r(i)$-th indifference class among the whole set of objects A.

However, the consumption process $P[t]$ differs utility-wise, from the moment τ. Thus, there must be at least one agent who under $P[t]$ immediately after τ stops consuming objects from her $l_r(i)$-th indifference class..We now check that under $P[t]$ there is an agent j, who at time τ starts consuming objects from lower indifference classes then $l_r(j)$-th class (one she consumes from under $S[t]$).

Consider two cases:

(i) $\sum_{1\le k\le r} \lambda_k < \tau$ (divergence from $S[t]$ occurs in the middle of algorithm's step $r+1$). Since under $S[t]$ everyone consumes from the best available objects (same ones before and after τ), any agent who under $P[t]$ does differently utility-vise is getting objects from lower indifference classes.

(ii) $\sum_{1\le k\le r} \lambda_k = \tau$ (divergence occurs at the start of algorithm's step $r+1$).

??

However, the consumption process $P[t]$ differs utility-wise, from the moment τ. Thus, as in the case of strict preferences, there is an agent j and an object b, such that the object b is still available at time τ, but under $P[t]$ this agent j switches at the moment τ to some object c which she values strictly less then b. Let object b be from the l-th indifference class for this agent j. Again, $t_i^S(k) = t_i^P(k)$ for all i and their indifference classes k, such that $t_i^S(k) \le \tau$, while $t_i^P(k) > \tau$ implies $t_i^S(k) > \tau$. However, $t_j^S(l) > \tau$ while $t_j^P(l) = \tau$. $\square$

Remark 1: The same line of argument allows us to prove a parallel characterization. Let $t_i'(a)$ to be the total share of objects at least as good as a agent i gets. The vector $t' = (t_i'(a))_{i,a}$ has fixed length nm, no matter whether preferences are strict or not.

Theorem 1'

For all preference profiles R, Serial assignment $S(R)$ is also the set of leximin maximizers of the $t = (t_i'(a))_{i,a}$.

Remark 2: It is easy to see that our result admits the following extensions.

(1) When agents find some objects "unacceptable", we introduce "null object", with as many copies as there are agents. We assume that each agent is indifferent between the "null object" and any her unacceptable object. The argument above still goes through.

(2) When different objects are available in different, even non-integer, quantities, our result is valid provided that each agent it entitled to the same unit total share of objects. This extension is important when we think of the second interpretation of our model (assignment with infinitely divisible goods).

Remark 3: While the vector $(t_i(k))_{i\in N,k}$ is leximin maximized at the Serial assignment, it is not, in general, Lorenz dominant on the set of such vectors for different feasible assignments (this is only true for the dichotomous domain of preferences, see [2] (2004)).

We conclude by providing a yet another characterization of the Serial rule. This one is much more in line with recent axiomatic work on random assignment, but apparently it was overlooked. It is however strongly related to the new definition of the Serial rule we propose. Indeed, on the strict domain Theorem 1 defines Serial rule as one leximin maximizing the vector of objects' shares all agents get. Thus, it can be regarded as the most fair rule when social planner, as well as agents, have lexicographic-type preferences. Our second result abstracts from social planner's preferences, and characterizes Serial rule as the unique appealing rule (efficient, fair, and incentive compatible) when agents' preferences over lotteries are lexicographic. Note that this is the first in the literature characterization of a specific assignment rule by those three basic classical requirements, without need for any additional technical axioms.

We do not formally define all the properties, as our proof heavily relies on the existing results. The exception is the property of limited invariance, which is explicitly used in the proof. An assignment rule $P(R)$ satisfies limited invariance, if the following is true for any initial preference profile R, any object a, and any agent i. When this agent i rearranges her preference ordering of objects she ranks below a, it does not affect her share of the object a. Formally, if her preferences R_i' are such that $U(R_i',a) = U(R_i,a)$ and $R_i'|_{U(R_i,a)} = R_i|_{U(R_i,a)}$, then it has to be $p_{ia}(R) = p_{ia}(R_i, R_{-i})$.

Definition 2 An agent has lexicographic preferences over $\triangle(A)$, the set of all probability distributions over objects, if she has strict preferences over A, say, $a_1 \succ a_2 \succ ... \succ a_m$, and prefers $p = (p_a)_{a\in A} \in \triangle(A)$ to $q = (q_a)_{a\in A} \in \triangle(A)$ as long as there is a $j \in \{1,...,n\}$ such that $p_{a_j} > q_{a_j}$, while $p_{a_i} = q_{a_i}$ for all $i < j$.

Theorem 2

On the strict ordinal preference domain, Serial rule is the only one which is sd-efficient, sd-envy-free, and strategy-proof on the lexicographic preference domain extension for lotteries.

Proof.

Serial rule is well-known to be sd-efficient and sd-envy-free. [11] (2012) shows that the Serial rule is strategy-proof on the lexicographic preference domain extension[13]. [5] (2012) shows that Serial rule is the only one which is sd-efficient, sd-envy-free, and satisfies limited invariance. Thus, it is enough to show that strategy-proofness on the lexicographic domain implies limited invariance.

Indeed, assume that a rule fails limited invariance. Hence, there is a preference profile R, agent i, object a, and preferences R_i', with $U(R_i',a) = U(R_i,a)$

[13] They need a condition that the smallest quantity of a good available is at least as large a the largest quota to which an agent is entitled. This condition is clearly satisfied in our setting, as all agents are entitled to 1 unit and each good exists in (at least) one unit.

and $R'_i|_{U(R_i,a)} = R_i|_{U(R_i,a)}$ and , such that $\sum\limits_{b\in U(R_i,a)} p_{ib}(R) \neq \sum\limits_{b\in U(R_i,a)} p_{ib}(R_i, R_{-i})$. Given such R and i, assume a to be the best for agent i object among those for which this inequality is true. Then, for any c which is better for i then a (at either R_i or R'_i), we have $\sum\limits_{b\in U(R_i,c)} p_{ib}(R) = \sum\limits_{b\in U(R_i,c)} p_{ib}(R_i, R_{-i})$. Hence, for those c we have $p_{ic}(R) = p_{ic}(R_i, R_{-i})$, while for a we obtain $p_{ic}(R) \neq p_{ic}(R_i, R_{-i})$.

Thus, there is a manipulation for agent i (either at profile R, or at profile $R' = (R_i, R_{-i})$) which increases her total share of good a, keeping the same her shares of goods better then a. The result of this manipulation will be lexicographically preferred to telling the truth, so strategy-proofness is violated. $\square$

Remark 4: This result cannot be extended to the full (non-strict) ordinal domain of preferences. While the argument in the proof goes through, Serial rule is not strategy-proof on the lexicographic extension of the full domain. Here is a counterexample (borrowed from [8] (2006)), with 3 agents and 3 objects. Consider the preference profile R, where agent 1 has preferences $a \sim_1 b \succ_1 c$, agent 2 has preferences $a \succ_2 b \succ_2 c$, and agent 3 has preferences $a \succ_3 c \succ_3 b$. If agent 1 reports preferences R'_1 with $a \succ_{1'} b \succ_{1'} c$, Serial assignment $S(R')$ gives her lexicographically better row of probabilities then $S(R)$ would, as can be seen below:

$$S(R) = \begin{array}{cccc} & a & b & c \\ 1 & 0 & 3/4 & 1/4 \\ 2 & 1/2 & 1/4 & 1/4 \\ 3 & 1/2 & 0 & 1/2 \end{array} \quad \text{and } S(R')S = \begin{array}{cccc} & a & b & c \\ 1 & 1/3 & 1/2 & 1/6 \\ 2 & 1/3 & 1/2 & 1/6 \\ 3 & 1/3 & 0 & 2/3 \end{array}$$

Remark 5: The axioms of sd-efficiency and sd-envy-freeness in Theorem 2 are based on the incomplete first order stochastic dominance relation, induced on the set of lotteries by ordinal preference reports over deterministic objects. One might be tempted to consider instead efficiency and envy-freeness for the lexicographic extension to lotteries. Notice that the requirement of efficiency for lexicographic extension (and, indeed, for any extension of deterministic rankings into lotteries which respects first order stochastic dominance) is at least as strong as sd-efficiency[14], while lexicographic no-envy is weaker then sd-no-envy. Unfortunately, when we weaken no-envy in this way, Serial rule is not singled out by our axioms anymore. The reader can easily check (going stage by stage), that the following "Boston mechanism" is lexicographically: envy-free, efficient and strategy-proof.

Stage 1: Start with the full set of objects $A_0 = A$ and fully distribute each object, which is the top choice for at least one person, equally between those agents who value it the best. Let $A_1 \subset A_0$ be the set of remaining unassigned objects.

Stage k: Start with the set A_k of not yet assigned objects, and with those agents who so far got less then full share 1 of objects. We distribute each object a, which is the best in A_k for at least one such agent, equally between those

[14] In fact, lexicographic efficiency is equivalent to sd-efficiency: both of them are equivalent to acyclicity condition from [1] (2001).

agents I_a who value it the best, subject to the constraint that no agent should get more then 1 unit total of objects' shares. Specifically, each agent i who ranks object $a \in A_k$ as her top in A_k, and who was already assigned the total quantity $SH_{k-1}(i) < 1$ of objects in preceding stages, gets $sh_k(i) = \min\{\lambda, 1 - SH_{k-1}(i)\}$ of object a, where λ is such that $\sum_{i \in I_a} sh_k(i) = 1$. Let $A_{k+1} \subset A_k$ be the set of remaining unassigned objects.

References

[1] Bogomolnaia, A., and H. Moulin, "New Solution to the Random Assignment Problem", *Journal of Economic Theory*, vol. 100 (2001), 295-328

[2] Bogomolnaia, A., and H. Moulin, "Random Matching under Dichotomous Preferences", *Econometrica*, vol. 72 (2004), 1, 257-279

[3] Bogomolnaia, A., and E.-J. Heo, "Probabilistic Assignment of Objects: Characterizing the Serial Rule", *Journal of Economic Theory,* vol. 147 (2012), 2072-2082

[4] Hashimoto, T., and D. Hirata, "Characterizations of the Probabilistic Serial Mechanism" (2011), mimeo

[5] Hashimoto, T., D. Hirata, O. Kesten, M. Kurino, and U. Ünver "Two Axiomatic Approaches to the Probabilistic Serial Mechanism" (2014), *Theoretical Economics* 9, 253–277

[6] Heo, E.-J., "Probabilistic Assignment with Quotas: A Generalization and a Characterization of the Serial Rule" (2010), mimeo

[7] Heo, E.-J., and O. Yilmaz, "A Characterization of the Extended Serial Correspondence" (2012), mimeo

[8] Katta, A.-K., and J. Sethuraman "A solution to the random assignment problem on the full preference domain", *Journal of Economic Theory*, vol. 131 (2006), 231-250

[9] Kesten, O., M. Kurino, and U. Ünver, "Fair and Efficient Assignment via the Probabilistic Serial Mechanism" (2010), mimeo

[10] Liu, Q., and M. Pycia, "Ordinal Efficiency, Fairness, and Incentives in Large Markets" (2011), mimeo

[11] Schulman L.J., and V.V. Vazirani, "Allocation of Divisible Goods under Lexicographic Preferences" (2012), mimeo