The most ordinally-egalitarian of random voting rules

Anna Bogomolnaia[*]

January 2015

Abstract

Aziz and Stursberg [1] propose an "Egalitarian Simultaneous Reservation" rule (ESR), a generalization of Serial rule, one of the most discussed mechanisms in random assignment problem, to the more general random social choice domain. We provide an alternative definition, or characterization, of ESR as the unique most ordinally-egalitarian one.

Specifically, given a lottery p over alternatives, for each agent i we define $t_i^p(k)$ to be the total share in p of objects from her first k indifference classes. ESR is shown to be the unique one which leximin maximizes the vector of all such shares $(t_i^p(k))_{i,k}$.

Serial rule is known to be characterized by the same property (see [2]). Thus, we provide an alternative way to show that ESR, indeed, coincides with Serial rule on the assignment domain. Moreover, since both rules are defined as the unique most ordinally-egalitarian ones, out result shows that ESR is "the right way" to think about generalizing Serial rule.

Keywords: Random Social choice, Random assignment, Serial Rule, Leximin

1 Introduction

We consider the classical "voting" problem, when n agents have to jointly choose one common alternative from a given set $A = \{a_1, ..., a_w\}$. Or goal is investigate plausible systematic preference aggregating mechanisms (rules) for this problem, which do not use monetary transfers.

When preferences over A differ substantially, it might be difficult to choose an alternative agents would consider a good compromise. One way to overcome this problem is to allow for an outcome to be a lottery over A, or a vector of "shares" of alternatives, rather then a unique alternative. Potentially, any probability distribution $p = (p_1, ..., p_n) \in \Delta A$ can be jointly chosen as an outcome. We may interpret p as a real lottery to be performed. Hence, the final ex-post outcome still would be a single "pure" alternative. However, agents might regard the

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

process (if not the outcome) as more fair. Alternatively, p may be interpreted as a vector of "time-shares", fractions of total time each alternative is in place. Which interpretation is more appropriate, depends on the particular economic situation. We abstract from it and concentrate on the formal model, which encompasses both.

A prominent paper by Gibbard [9] on this random social choice model restricts attention to the "ordinal" mechanisms. Agents are assumed to have strict preferences over A, and are only asked about their orderings of pure alternatives. It is implicitly assumed though, that they have cardinal utilities over the alternatives, and compare lotteries based on the expected utility. This assumption gives rise to a strong requirement of "strategy-proofness": a rule is non-manipulable only if an agent can never gain (no matter what are her cardinal utilities behind) by altering her ordinal input. Gibbard [9] characterizes all strategy-proof mechanisms. Follow-up works were also mainly concerned with non-manipulability.

When we allow for indifferences in agents' preferences, assignment of private goods becomes a particular case of voting model (each agent is indifferent between all assignments which give her the same thing). Ordinal random assignment model, with indivisible goods and no monetary transfers, became a very active area of research in recent 10-15 years. Bogomolnaia and Moulin [3] proposed to look at the ordinal random assignment mechanisms (agents are only required to report preferences over deterministic alternatives). They introduced, for the strict preference domain, a new "Serial rule", which is computed allowing agents to acquire ("eat") shares of objects simultaneously with the same constant speed, in decreasing order of their preferences. While it only satisfies a week version of non-manipulability, its fairness and efficiency properties are very strong. In particular, it is not only anonymous, but also envy-free. Serial rule became one of the most studied ordinal mechanisms (another one is Random Priority). Several generalizations were proposed. Bogomolnaia, Moulin [4] introduced Egalitarian rule for the dichotomous assignment domain. Katta, Sethuraman [10] extended the definition of Serial rule to the full domain (indifferences are allowed), which is also an extension of the above Egalitarian rule from dichotomous domain. Several recent papers (see, for example [6], [7], and, for the full domain, [8]) provided axiomatic characterizations of Serial rule, mostly by means of efficiency, envy-freeness, and some type of monotonicity with respect to certain changes in preferences. Bogomolnaia [2] proposes an alternative definition, or a characterization, of Serial rule, both for strict and full domain. Given an arbitrary random assignment, one can calculate, for each agent i, her total share $t_i(k)$ of goods from her k best indifference classes. Serial rule happens to be the unique one which leximin maximizes vector $t = (t_i(k))_{i,k}$. Thus, arguably, Serial rule emerges as the (unique) most egalitarian random rule, when social planner is restricted to ordinal information only. This result also serves as a justification of the extension of Serial rule to the full domain, proposed in [10]. It shows that this generalization is indeed in the same spirit as Serial (or Egalitarian) rule, aiming at the ordinally-egalitarian goal.

Returning to the random voting problem, now over full domain (indifferences

allowed), Aziz and Stursberg [1] propose a random social choice (voting) rule which they call Egalitarian Simultaneous Reservation (ESR). This rule is based on the same "simultaneous eating" ideas as Serial rule. Indeed, ESR is introduced as a joint generalization of two rules. One is Serial rule, introduced in [3], [4], and [10] for strict, dichotomous, and full private goods domain, and another is Egalitarian random voting rule, defined in [5] on the dichotomous voting preference domain. Aziz and Stursberg [1] show that on the assignment domain ESR coincides with Serial rule and on the dichotomous domain it coincides with the Egalitarian voting rule. In order to show the equivalence with Serial rule for assignment domain, they rely on a recent characterization result, [8], which singles out Serial rule on full domain by ordinal efficiency, envy-freeness, and "limited invariance" (the proof in [8] is long and complicated). Aziz and Stursberg [1] check that ESR satisfies those three properties on the assignment domain. The checks, while not very complicated, are also rather long and non-intuitive. Note also, that the notion of envy-freeness looses its meaning beyond the domain of private goods' assignment. While ESR remains ordinally efficient and satisfies limited invariance on the full domain, those properties do not single it out. This, while ESR is one of possible generalizations of Serial rule, the question remains whether it is the most appropriate one.

Current work proposes an alternative definition (or a characterization) of ESR, along the same lines as the alternative definition of Seral rule, proposed in Bogomolnaia [2]. It shows that ESR is the unique leximin maximizer of the vector $t^p = (t_i^p(k))_{i,k}$ (where $t_i^p(k)$ is the total share agent i gets in lottery p from her top k indifference classes). An immediate corollary, given the result in Bogomolnaia [2], is an alternative, much more straightforward, way to show that ESR is indeed an extension of Serial rule to the larger domain.

Our main result can be illustrated by the following interpretation. Fix a lottery p. 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 by the total amount of objects she gets from her first k indifference classes, i.e. $t_i^p(k)$. Our result is that ESR 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 ESR allocation 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.

All the above allows us to argue that ESR is indeed the "right way" to generalize Serial rule to the general random voting domain. Both original and generalized rules have exactly the same nature: each gives the unique most "ordinally-egalitarian" way to compromise between agents.

3

2   Model and Results

Given are a set of agents $N = \{1, ..., n\}$ and a set of alternatives $A = \{a_1, ..., a_w\}$. Each agent $i \in N$ has arbitrary preferences R_i, or $\succeq_i$, over A, which are represented by a partition of A into K_i indifference classes $E_i^1, ..., E_i^{K_i}$ ($a \succ_i b$ if and only if $a \in E_i^l$, $b \in E_i^r$, and $l < r$). Denote by $\mathcal{R}$ the set of all preference orderings over A.

Let $\Delta(A)$ be the set of all lotteries on A. A random social choice rule is a correspondence $f : \mathcal{R}^n \rightrightarrows \Delta(A)$, which is essentially single-valued. Specifically, at any preference profile $R = (R_1, ..., R_n)$, $f(R) \subset \Delta(A)$, $f(R) \neq \varnothing$, and all agents are indifferent between all outcomes in $f(R)$.

Fix a preference profile R, once and forever. Let $\mathcal{E}_i = \left\{ E_i^1, ..., E_i^{K_i} \right\}$ be the set of agent's i indifference classes, and let $\mathcal{E} = \bigcup \mathcal{E}_i$ (the union of all indifference classes over all agents).

Given an arbitrary lottery $p = (p_1, ..., p_w)$, where $p_j = p(a_j)$, we will write $p(E) = \sum_{a \in E} p(a)$. For any agent i and any $k \leq K_i$, we define $t_i(k) = t_i^p(k) = \sum_{r \leq k} p(E_i^r) = \sum_{a_j \in \cup_{r \leq k} E_i^r} p_j$, the total share agent i gets in lottery p of objects from her first k indifference classes. Define vector $t_i = t_i^p = (t_i(1), ..., t_i(K_i))$, where K_i is the number of indifference classes in R_i.

Aziz and Stursberg [1] propose a random social choice rule which they call Egalitarian Simultaneous Reservation (ESR).

ESR starts from the set of all feasible lotteries, and then repeatedly shrinks this set over time interval $[0, 1]$. This is done by sequentially introducing (the largest feasible) lower bounds on probabilities of certain indifference classes from $\mathcal{E}$ in the final lottery. At each step, ESR only cares to guarantee each agent the largest share of objects from her top indifference class still available for distributing. As we will see, ESR pursues an egalitarian goal to guarantee the best treatment for the worst off agents. Aziz and Stursberg's definition is given by the following algorithm.

Each indifference class $E = E_i^r$ is represented by a "tower" growing over time, with the "ceiling" $l_t(E) \in \mathbb{R}_+$ at time t. This ceiling represents the minimal guarantee for the probability the subset E of alternatives is to receive in the final lottery. During the course of the algorithm, agents climb up those towers (all with the same constant speed) and in doing that push up the ceilings, therefore increasing the lower bounds on probabilities of corresponding subsets.

All towers have ceilings zero at time $\tau = 0$. Algorithm proceeds in stages. In stage 1, each agent starts by climbing the tower corresponding to her top indifference class. A tower's height is frozen, if increasing it would result in nonexistence of a lottery satisfying all lower bounds for all towers. Once a tower is frozen, any agent at its ceiling falls off it, and moves to the bottom of the tower corresponding to her next indifference class[1]. Any agent in the middle of a frozen tower will continue to climb it until the frozen ceiling. Then she will

---

[1]It might be an already frozen tower!

fall off and move to her next best tower. A stage ends whenever some agent falls off some ceiling.

Notation: Let $l^{(k)} : \mathcal{E} \to \mathbb{R}_+$ be the height of the ceiling in the tower $S \in \mathcal{E}$ at the end of stage k. Function $E^{(k)} : N \to \mathcal{E}$ represents the towers each agent i climbs during stage k. Let $h_i^{(k)}$ be the height where agent i is in her tower at the beginning of stage k; $\lambda^{(k)}$ be the duration of stage k; $N^{(k)} \subset N$ be the set of agents falling off at the end of stage k; $\tau^{(k)}$ be the total time passed by the beginning of stage k.

Assign $l^{(0)}(S) = h_i^{(0)} = \tau^{(0)} = 0$, and $E^{(0)}(i) = E_i^1$ (the top indifference class of agent i), for all $i \in N$ and all $S \in \mathcal{E}$.

Stage k. Compute $\lambda^{(k)}$ and $N^{(k)}$ (see the linear program below the algorithm's description), as well as towers to be frozen. Let $\tau^{(k+1)} = \tau^{(k)} + \lambda^{(k)}$, and agents' heights for the beginning of the next stage be $h_i^{(k+1)} = h_i^{(k)} + \lambda^{(k)}$ if $i \notin N^{(k)}$, $h_i^{(k+1)} = 0$ if $i \in N^{(k)}$ (i.e., if agent i fell off). For any $S \in \mathcal{E}$ define $l^{(k+1)}(S) = \max \left\{ l^{(k)}(S), \max \left\{ h_i^{(k+1)} | E^{(k)}(i) = S \right\} \right\}$ (new ceilings, defined as maximum of the old ceiling and maximal of agents' heights over those who were climbing this tower during stage k). For all agents who fell off, $i \in N^{(k)}$, define $E^{(k+1)}(i)$ to be the agent's i next indifference class. For remaining agents, define $E^{(k+1)}(i) = E^{(k)}(i)$.

Calculation of $\lambda^{(k)}$ and $N^{(k)}$ comes from solving the following Linear Program.

First, we find $\lambda^{(k)}$ by solving $\max \lambda$ subject to:
$\sum_{a \in S} p(a) \geq l^{(k)}(S)$ for all $S \in \mathcal{E}$; $\sum_{a \in E^{(k)}(i)} p(a) \geq h_i^{(k)} + \lambda$ for all $i \in N$; $\sum_{a \in N} p(a) \leq 1$, $p(a) \geq 0$ for all $a \in A$; $\lambda \geq 0$.

Then, for all $i \in N$ we run this program again, imposing $\lambda = \lambda^{(k)}$, and maximizing the slack in the inequality $\sum_{a \in E^{(k)}(i)} p(a) \geq h_i^{(k)} + \lambda$; whenever objective function value is zero, we add corresponding i to $N^{(k)}$.

Thus, resulting $N^{(k)}$ is the largest (by inclusion) set of agents who create the "bottleneck" (prevent increasing $\lambda^{(k)}$). I.e., $\lambda^{(k)}$ and $N^{(k)}$ are computed as the largest ones so that there still exists a lottery $p \in \Delta(A)$ with $p(S) = \sum_{a \in S} p(a) \geq l^{(k+1)}(S)$ for all $S \in \mathcal{E}$.

Freeze (forever) ceilings of all towers which were climbed by at least one agent from $N^{(k)}$.

Algorithm finishes when all towers are frozen. Once it is finished, ESR at the profile R is the set of "remaining" lotteries p, satisfying inequalities $\sum_{a \in S} p(a) \geq l^{(\overline{k})}(S)$ for the last step $\overline{k}$ (and hence for all steps). It is straightforward to show (see [1]) that for any profile R this set is essentially single-valued. Hence, ESR is indeed a random social choice rule. Moreover, it is easy to see that for any profile R, $p \in ESR(R)$ and (q utility equivalent to p) imply $q \in ESR(R)$.

Thus, $ESR(R)$ includes all lotteries utility equivalent to any one obtained by the algorithm above.

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

Theorem 1

For all preference profiles R, $ESR(R)$ is exactly the set of lotteries over $\Delta(A)$ which leximin maximize the vector[2] of shares $t = (t_1, ..., t_n)$.

Proof.

Fix an arbitrary preference profile R. Since they are utility equivalent, all $p \in ESR(R)$ result in the the same vector t^p and same permutation-vector t^{p*}. Denote the corresponding permutation $\pi : t^p \to t^{p*}$.

Since $k < r$ implies $t_i^p(k) \leq t_i^p(r)$, we can assume that π is such that $\pi\left(t_i^p(k)\right) < \pi\left(t_i^p(r)\right)$ for any i, whenever $k < r$.

In the algorithm, each agent i climbs towers corresponding to her indifference classes $E_i^1, ..., E_i^{K_i}$ successively in increasing order, and without skipping. She only falls off a tower and moves to the next one once she reaches a frozen ceiling. Assume at the end of stage k, at the time $\tau^{(k+1)}$, agent i falls off tower E_i^m. It means that by the time $\tau^{(k+1)}$ towers corresponding to $E_i^1, ..., E_i^m$ are frozen at their "final" heights $l\left(E_i^1\right), ..., l\left(E_i^m\right)$ which correspond to their maximal feasible minimal guarantees. I.e., in a final lottery probabilities of those indifference classes will be exactly $l\left(E_i^1\right), ..., l\left(E_i^m\right)$. Since our agent i goes through each tower from bottom to top and never stops, we obtain that the total share she gets of objects from her first m indifference classes $t_i^p(m) = l\left(E_i^1\right) + ... + l\left(E_i^m\right) = \tau^{(k+1)}$.

Moreover, if instead at the beginning of stage $k + 1$ agent i is at the height $h_i^{(k+1)} > 0$ in tower E_i^m, then $t_i^p(m - 1) = l\left(E_i^1\right) + ... + l\left(E_i^{m-1}\right) = \tau^{(k+1)} - h_i^{(k+1)}$.

We will check by induction (in stages of the algorithm), that at each stage k the algorithm to find $ESR(R)$ considers, for each agent i, the smallest "not yet fixed" $t_i^p(m_i)$ (here $m_i = m_i(k)$, we suppress the stage index), and maximizes ("fixes") the smallest of those. More specifically, it finds a largest μ such that it is feasible to have $\min_{i \in N} t_i^p(m_i) \geq \mu$ and "fixes" all $t_i^p(m_i)$ which cannot be larger then μ to be $t_i^p(m_i) = \mu$. The theorem follows immediately from this statement.

In stage 0, the algorithm starts by eliminating all lotteries, except those which maximize the smallest among $t_i^p(1) = p(E_i^1) = \sum_{a \in E_i^1} p(a)$, i.e. the smallest element in t^p. In other words, it "fixes" (freezes) all $t_i^p(1)$ with $i \in N^{(k)}$, i.e. all $t_i^p(1)$ which cannot feasibly be more then $\lambda^{(1)}$, at the level $t_i^p(1) = \lambda^{(1)}$.

---

[2] Recall that each t_i is itself a vector, $t_i = (t_i(1), ..., t_i(K_i))$, where $t_i(k) = \sum_{a_j \in \cup_{r \leq k} E_i^r} p_j$.

Hence, $t = (t_1, ..., t_n)$ is the vector of length $\sum_{i \in N} K_i$.

In each next stage $k+1$, the algorithm pays attention to exactly one tower $E_i^{m_i}$ for each agent i (the tower $E^{(k+1)}(i)$ she climbs during this stage), corresponding to her m_i-th indifference class ($m_i = m_i(k+1)$, we suppress the stage index), while all towers corresponding to agents' better indifference classes are already frozen. Hence, all $t_i^p(m)$ for all i and all $m < m_i$, are "fixed" to $t_i^p(m) = l\left(E_i^1\right) + ... + l\left(E_i^m\right)$, their best minimal guarantees (by induction hypothesis).

In stage $k+1$ we find the largest λ such that $p\left(E_i^{m_i}\right) = p\left(E^{(k+1)}(i)\right) = \sum_{a \in E^{(k+1)}(i)} p(a) \geq h_i^{(k+1)} + \lambda$ for all i. But $h_i^{(k+1)} = \tau^{(k+1)} - t_i^p(m_i - 1)$, so the constraints can be rewritten as $p\left(E_i^{m_i}\right) \geq \tau^{(k+1)} - t_i^p(m_i - 1) + \lambda$, or the $t_i^p(m_i) = p\left(E_i^{m_i}\right) + t_i^p(m_i - 1) \geq \tau^{(k+1)} + \lambda$.

Thus, algorithm finds the largest amount $\tau^{(k+2)} = \tau^{(k+1)} + \lambda^{(k+1)}$ such that all $t_i^p(m_i)$ are at least $\tau^{(k+2)}$. It then freezes all towers $E_i^{m_i} = E^{(k+1)}(i)$ whose height cannot feasibly increase, i.e., "fixes" all $t_i^p(m_i)$ which cannot be larger then $\tau^{(k+2)}$ to be $t_i^p(m_i) = \tau^{(k+2)}$. $\quad\square$

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 nw, no matter whether preferences are strict or not.

Theorem 1a

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

Serial rule for random assignment problem is known to be characterized by the same property as in our Theorem 1 (see [2]). We hence obtain:

Corollary

On the assignment domain, ESR coincides with Serial rule.

References

[1] Aziz, H., and P. Stursberg, "A Generalization of Probabilistic Serial to Randomized Social Choice", *Proceedings of the 28-th AAAI Conference on Artificial Intelligence* (2014), 559–565

[2] Bogomolnaia, A., "Random Assignment: Redefining the Serial Rule" (2014), mimeo

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

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

[5] Bogomolnaia, A., H. Moulin, and R. Stong, "Collective choice under dichotomous preferences", *Journal of Economic Theory*, vol. 122(2) (2005), 165–184

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

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

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

[9] Gibbard, A., "Manipulation of Schemes That Mix Voting with Chance", *Econometrica,* vol. 45 (1977), 665–681

[10] 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