

A Generalization of Probabilistic Serial to Randomized Social Choice

Haris Aziz

NICTA and UNSW,
Sydney 2033, Australia
haris.aziz@nicta.com.au

Paul Stursberg

Technische Universität München,
85748 Garching bei München, Germany
paul.stursberg@ma.tum.de

Abstract

The *probabilistic serial* rule is one of the most well-established and desirable rules for the random assignment problem. We present the *egalitarian simultaneous reservation* social decision scheme — an extension of *probabilistic serial* to the more general setting of randomized social choice. We consider various desirable fairness, efficiency, and strategic properties of social decision schemes and show that *egalitarian simultaneous reservation* compares favorably against existing rules. Finally, we define a more general class of social decision schemes called *simultaneous reservation*, that contains *egalitarian simultaneous reservation* as well as the serial dictatorship rules. We show that outcomes of *simultaneous reservation* characterize efficiency with respect to a natural refinement of stochastic dominance.

Introduction

Whenever people form a society and need to coordinate their actions, the problem of collective decision-making is close at hand. Regardless of whether real, physical agents or software agents in the context of a multi-agent system are involved, they are likely to have different objectives or resources, and thus prefer different collective actions (Brandt, Conitzer, and Endriss 2013).

Hence, the group needs a clearly defined mechanism, like an electoral system or a negotiation protocol, that queries the agents' preferences and produces a collective choice from a given set of alternatives. Such a mechanism in turn is more likely to be accepted if it can guarantee certain properties of the outcome, such as efficiency (if all agents prefer one outcome over another, the latter is never selected), strategyproofness (no agent can manipulate the outcome in his favour by misrepresenting his preferences) or fairness. The aggregation of preferences respecting these requirements is therefore considered one of the most fundamental problems in both economics and computer science (Conitzer 2010).

As deterministic decisions can be inherently unfair (consider the case of two alternatives and two agent with opposite preferences), *randomized social choice* investigates mechanisms that return randomized decisions, i.e. probability distributions over a set of alternatives (Procaccia 2010; Conitzer and Sandholm 2006; Service and Adams 2012).

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Depending on the application these need not be real lotteries from which only one alternative is chosen in the end. Instead, a fraction of the time available can be devoted to each single alternative or a resource in question can be shared among agents. In this paper, however, we will always talk about the outcomes as lotteries where in the end one alternative is selected at random according to some distribution.

We can think of such mechanisms as functions that take as input the agents' preferences over a set A of alternatives and return for each such preference profile a probability distribution, or lottery, over A . In randomized social choice, these functions are known as *social decision schemes* (SDSs). We explicitly allow agents to express indifferences between alternatives (for strict preferences, on the other hand, *random dictatorship* (RD) is by many accounts the single most appealing SDS (Gibbard 1977; Aziz 2013)).

One particular setting that can be seen as a restricted subclass of randomized social choice is the *random assignment problem*, one of the most important settings in resource allocation (Bogomolnaia and Moulin 2001; Katta and Sethuraman 2006). Here, agents express preferences over a set of objects and the alternatives between which the mechanism has to choose are allocations of objects to the agents. The output thus specifies for each agent the fraction of each object that is allocated to the agent. A number of natural random assignment rules have previously been proposed. These rules include *random serial dictatorship* (RSD) (Bogomolnaia and Moulin 2001), the *popular assignment rule* (Kavitha, Mestre, and Nasre 2011; Aziz, Brandt, and Stursberg 2013) and the *uniform assignment rule* (Chambers 2004). For the assignment problem in which agents have strict preferences over objects, a particular elegant mechanism with superior fairness and efficiency properties is the *probabilistic serial rule* (PS) (Bogomolnaia and Moulin 2001; Kojima and Manea 2010; Manea 2009; Saban and Sethuraman 2013). PS also has a very intuitive description: Each agent eats his most preferred available object at a uniform rate. The fraction of an object eaten by an agent is the probability with which the object is allocated to that agent. Katta and Sethuraman (2006) have generalized PS to the *extended probabilistic serial rule* (EPS) for the setting where agents may express ties between objects.

For almost all previously mentioned random assignment

rules, a generalization to randomized social choice is known. *RSD* for instance, just as the uniform rule, generalizes very naturally (Aziz, Brandt, and Brill 2013a).¹ Recently, the popular random assignment rule has been generalized to the social choice setting under the name of *strict maximal lottery rule (SML)* (Aziz, Brandt, and Brill 2013b). For *PS*, however, no generalization to the social choice domain is known. In this paper, we present a generalization of *PS* that constitutes a new and interesting SDS with desirable axiomatic properties for the general social choice setting.

Contributions We present a new SDS called *egalitarian simultaneous reservation (ESR)*. We show that it is equivalent to *EPS*, *PS*, and the egalitarian assignment rule (Bogomolnaia and Moulin 2004) for their respective subdomains. Our proof for the equivalence result relies on a recent characterization of *EPS* (Heo and Yilmaz 2012). *ESR* also generalizes the egalitarian SDS defined solely for dichotomous preferences by Bogomolnaia, Moulin, and Stong (2005).

We show that from an axiomatic point of view, *ESR* compares favorably to *RSD* and *SML* — two prominent SDSs for the general preference domain. We prove that *ESR* is *DL*-efficient and hence *SD*-efficient. It is monotonic (reinforcing an alternative in the agent’s opinions can only increase its probability), a ‘clearly desirable condition’ (Fishburn 1982), and satisfies the fair outcome share property. *ESR* is weak *SD*-strategyproof for strict preferences, but not for dichotomous preferences.

Finally, based on a parametrization of the *ESR* mechanism, we define a more general class of SDSs called *simultaneous reservation (SR)*, the outcomes of which characterize all *DL*-efficient outcomes. Our characterization provides a new perspective on *DL*-efficiency, a concept that has recently attracted interest in randomized social choice (Schulman and Vazirani 2012; Saban and Sethuraman 2013; Cho 2012). We show that serial dictatorship is a special class of SDSs within the class of *SR*.

Preliminaries

Social choice and lotteries Consider the social choice setting in which there is a set of agents $N = \{1, \dots, n\}$, a set of alternatives $A = \{a_1, \dots, a_m\}$ and a preference profile $\succsim = (\succsim_1, \dots, \succsim_n)$ such that each \succsim_i is a complete and transitive relation over A . We write $a \succsim_i b$ to denote that agent i values alternative a at least as much as alternative b and use \succ_i for the strict part of \succsim_i , i.e., $a \succ_i b$ iff $a \succsim_i b$ but not $b \succsim_i a$. Finally, \sim_i denotes i ’s indifference relation, i.e., $a \sim_i b$ iff both $a \succsim_i b$ and $b \succsim_i a$. The relation \succsim_i results in equivalence classes $E_i^1, E_i^2, \dots, E_i^{k_i}$ for some k_i such that $a \succ_i a'$ iff $a \in E_i^l$ and $a' \in E_i^{l'}$ for some $l < l'$. Often, we will use these equivalence classes to represent the preference relation of an agent as a preference list $i : E_i^1, E_i^2, \dots, E_i^{k_i}$. For example, we will denote the preferences $a \sim_i b \succ_i c$ by

¹For social choice with strict preferences, this generalization is equivalent to *RD*.

the list $i : \{a, b\}, \{c\}$. We denote the set of all equivalence classes of agent i by \mathcal{E}_i and the union over all agents by \mathcal{E} .

An agent i ’s preferences are *dichotomous* iff he partitions the alternatives into just two equivalence classes, i.e., $k_i = 2$. An agent i ’s preferences are *strict* iff \succsim_i is antisymmetric, i.e. all equivalence classes have size 1.

Let $\Delta(A)$ denote the set of all *lotteries* (or *probability distributions*) over A . The support of a lottery $p \in \Delta(A)$, denoted by $\text{supp}(p)$, is the set of all alternatives to which p assigns a positive probability, i.e., $\text{supp}(p) = \{x \in A \mid p(x) > 0\}$. We will write $p(a)$ for the probability of alternative a and we will represent a lottery as $p_1 a_1 + \dots + p_m a_m$ where $p_j = p(a_j)$ for $j \in \{1, \dots, m\}$. For $A' \subseteq A$, we will (slightly abusing notation) denote $\sum_{a \in A'} p(a)$ by $p(A')$.

A *social decision scheme* is a function $f : \mathcal{R}^n \rightarrow \Delta(A)$. If f yields a set rather than a single lottery, we call f a *correspondence*. Two minimal fairness conditions for SDSs are *anonymity* and *neutrality*. Informally, they require that the SDS should not depend on the names of the agents or alternatives respectively.

Another desirable property is monotonicity: An alternative preference profile \succsim' reinforces a compared to \succsim iff for some agent $i \in N$ and alternative a we have $\succsim_j = \succsim'_j$ for all agents $j \neq i$ and \succsim'_i is a transitive and complete relation with $b \succsim'_i b' \Leftrightarrow b \succsim_i b'$ for all $b, b' \neq a$ and $a \succ b \Rightarrow a \succ' b$. A correspondence f is *monotonic* iff whenever \succsim' reinforces a compared to \succsim then for every $p \in f(\succsim)$ we can find $p' \in f(\succsim')$ with $p'(a) \geq p(a)$.²

Lottery extensions In order to reason about the outcomes of SDSs, we need to determine how agents compare lotteries. A *lottery extension* extends preferences over alternatives to (possibly incomplete) preferences over lotteries. Given \succsim_i over A , a *lottery extension* \mathcal{X} extends \succsim_i to $\succsim_i^{\mathcal{X}}$ over the set of lotteries $\Delta(A)$. We now define some particular lottery extensions that we will later refer to.

- Under *stochastic dominance (SD)*, an agent prefers a lottery that, for each alternative $x \in A$, has a higher probability of selecting an alternative that is at least as good as x . Formally, $p \succsim_i^{SD} q$ iff $\forall y \in A : \sum_{x \in A : x \succsim_i y} p(x) \geq \sum_{x \in A : x \succsim_i y} q(x)$.
- In the *downward lexicographic (DL)* extension, an agent prefers the lottery with higher probability for his most preferred equivalence class, in case of equality, the one with higher probability for the second most preferred equivalence class, and so on. Formally, $p \succsim_i^{DL} q$ iff for the smallest (if any) l with $p(E_i^l) \neq q(E_i^l)$ we have $p(E_i^l) > q(E_i^l)$.
- The *sure thing (ST)* extension is defined as follows. Let $\delta(p, q) = \{x \in A : p(x) \neq q(x)\}$. Then, $p \succsim_i^{ST} q$ iff $p = q$ or $\forall x \neq y \in A : (\{x, y\} \cap \delta(p, q) \neq \emptyset \text{ and } p(x)q(y) > 0) \Rightarrow x \succ_i y$.

²Note that monotonicity by this definition implies monotonicity with respect to the reinforcement of a in more than one agent’s preference relation (by reinforcing a in the agents’ preferences one after the other).

SD (Bogomolnaia and Moulin 2001) is particularly important because $p \succsim^{SD} q$ iff p yields at least as much expected utility as q for any von-Neumann-Morgenstern utility function consistent with the ordinal preferences (Cho 2012). DL refines SD to a complete relation based on the natural lexicographic relation over lotteries (Schulman and Vazirani 2012; Abdulkadiroğlu and Sönmez 2003; Cho 2012) whereas ST (Aziz, Brandt, and Brill 2013b) is even coarser than SD . Based on the lottery extensions, we define corresponding notions of efficiency and strategyproofness.

Efficiency and strategyproofness Let \mathcal{X} be any lottery extension. A lottery p is \mathcal{X} -efficient iff there exists no lottery q such that $q \succsim_i^{\mathcal{X}} p$ for all $i \in N$ and $q \succ_i^{\mathcal{X}} p$ for some $i \in N$. An SDS is \mathcal{X} -efficient iff it always returns an \mathcal{X} -efficient lottery. A standard efficiency notion that cannot be phrased in terms of lottery extensions is *ex post efficiency*. A lottery is *ex post efficient* iff it is a lottery over Pareto optimal alternatives. It is the case that DL -efficiency \implies SD -efficiency \implies *ex post efficiency* \implies ST -efficiency.

An SDS f is \mathcal{X} -manipulable iff there exists an agent $i \in N$ and preference profiles \succsim and \succsim' with $\succsim_j = \succsim'_j$ for all $j \neq i$ such that $f(\succsim') \succ_i^{\mathcal{X}} f(\succsim)$. An SDS is *weakly \mathcal{X} -strategyproof* iff it is not \mathcal{X} -manipulable, it is \mathcal{X} -strategyproof iff $f(\succsim) \succsim_i^{\mathcal{X}} f(\succsim')$ for all \succsim and \succsim' with $\succsim_j = \succsim'_j$ for all $j \neq i$. Note that SD -strategyproofness is equivalent to strategyproofness in the Gibbard sense. It is known that SD -strategyproof \implies DL -strategyproof \implies weak SD -strategyproof \implies weak ST -strategyproof.

Egalitarian Simultaneous Reservation

Starting from the entire set $\Delta(A)$, the ESR algorithm proceeds by gradually restricting the set of possible outcomes. The restrictions enforced are lower bounds for the probability of certain equivalence classes while it is always maintained that a lottery exists that satisfies all these lower bounds. Each equivalence class E is represented by a *tower* where at any time t , the height of this tower's *ceiling* $\ell_t(E)$ represents the lower bound in place for the probability of this subset at that time. During the course of the algorithm, agents will climb up these towers and try to push up the ceilings, thereby increasing the lower bounds for certain subsets.

Each tower starts with the height of its ceiling set to 0. Every agent starts climbing up the tower that corresponds to his most preferred equivalence class. Whenever an agent hits the ceiling, he tries to push it up. He continues climbing, pushing up the ceiling at the same time. Note that the ceiling will only be pushed up as fast as the agent pushing it can climb. Two agents pushing up a ceiling at the same time therefore does not increase the speed of it being pushed up. When it cannot be pushed up any further without compromising the existence of a lottery satisfying all current lower bounds, we say that set E is *tight* and has been *frozen*. At this point, the agent *bounces off the ceiling* and drops back to the floor, moving on to the tower corresponding to his next most preferred equivalence class. We can think of the algorithm proceeding in stages where a stage ends whenever some agent bounces off the ceiling. The algorithm ends

when all the equivalence classes have been frozen at which point some lottery satisfying the lower bounds is returned. A formal description of ESR is given as Algorithm 1.³ In general, the result is not unique, although often it is.

Example 1. Consider the following preference profile:

- 1 : $\{a\}, \{b\}, \{e\}, \{c, d\}$
- 2 : $\{a\}, \{c\}, \{d\}, \{b, e\}$
- 3 : $\{b, d\}, \{a, c, e\}$
- 4 : $\{c, e\}, \{a, b, d\}$
- 5 : $\{c\}, \{a, b, e\}, \{d\}$

At time $1/3$, all agents bounce off their respective ceilings. The lower bounds in place at this time are $\ell(\{a\}) = 1/3$; $\ell(\{b, d\}) = 1/3$; $\ell(\{c, e\}) = 1/3$; $\ell(\{c\}) = 1/3$. All agents move to their next equivalence class. After an additional $1/3$ time, the following lower bounds have been added: $\ell(\{b\}) = 1/3$; $\ell(\{a, c, e\}) = 1/3$; $\ell(\{a, b, d\}) = 1/3$; $\ell(\{a, b, e\}) = 1/3$. Note that during this period, agent 2 could not push up the ceiling in tower $\{c\}$ any further, nor did he bounce off as he first had to reach the ceiling. Eventually, $ESR(\succsim) = \frac{1}{3}a + \frac{1}{3}b + \frac{1}{3}c$.

Observation 1. For strict preferences, ESR returns the uniform lottery over all alternatives that are most preferred by some agent. Hence, ESR is not equivalent to random dictatorship (RD) for strict preferences (Gibbard 1977) where RD is defined as $RD(\succsim) = \sum_{i \in N} \frac{1}{n} \max_{\succsim_i} (A)$.

Although ESR is technically a correspondence, each agent is completely indifferent between all possible outcomes. A correspondence is called *essentially single-valued* iff for every agent i , equivalence class E of i and lotteries p and q that are outcomes of the correspondence, $p(E) = q(E)$. Furthermore, ESR is computable in polynomial time.

Theorem 1. ESR is essentially single-valued.

Proof. Let p and q be two outcomes and E any equivalence class of some agent i where p and q differ, w.l.o.g. let $p(E) > q(E)$. Consider the point in time t when i bounced off the ceiling in tower E . As q is an outcome, $\ell_t(E) \leq q(E) < p(E)$. As p is an outcome, it does fulfill all other guarantees, in particular those guarantees that are present at time t . But this contradicts with i bouncing off the ceiling of E at time t (as p is a feasible lottery that yields a higher probability for E). \square

Theorem 2. The ESR algorithm runs in polynomial time.

Proof. First, consider `computeLambda`. The function solves $|N| + 1$ linear programs where each constant in each LP is either zero, one, or a polynomial-size rational computed from a previous LP. These LPs (and thus `computeLambda`) can be solved in polynomial time. Furthermore, optimality of λ^* ensures that N^* contains at least one element. Thus, during each iteration of Algorithm 1 in line 11, at least one agent moves to his next equivalence

³A few variants of ESR naturally come into mind, but neither the variant in which the ceiling of an equivalence class moves at the speed proportional to the number of agents pushing it nor the variant in which agents do not enter a new tower at the bottom but at the current height of its ceiling is monotonic.

Input: (N, A, \succsim)
Output: lottery p over A .

$\ell^{(k)} : \mathcal{E} \rightarrow \mathbb{R}$: height of ceiling in tower $S \in \mathcal{E}$ at the end of stage k
 $E^{(k)} : N \rightarrow \mathcal{E}$: tower that $i \in N$ is in at stage k
 $h_i^{(k)}$: height of agent $i \in N$ at stage k
 $\lambda^{(k)}$: length (duration) of stage k
 $N^{(k)} \subseteq N$: agents bouncing off after stage k
 $t^{(k)}$: total time passed at the beginning of stage k

```

1  $\ell^{(0)}(S) = h_i^{(0)} = t^{(0)} = k = 0$  for all  $S \subset A, i \in N$ 
2  $E^{(0)}(i) = i$ 's first equivalence class for all  $i \in N$ 
3 while  $t^{(k)} < 1$  do % get time when first agent bounces off
4    $(\lambda^{(k)}, N^{(k)}) \leftarrow \text{computeLambda}(k)$  % Alg. 2
5    $t^{(k+1)} = t^{(k)} + \lambda^{(k)}$ 
6    $h_i^{(k+1)} \leftarrow h_i^{(k)} + \lambda^{(k)}$  for all  $i \in N$  % climb on
7   for all  $S \in \mathcal{E}$  do % set new ceiling levels
8      $\ell^{(k+1)}(S)$ 
9      $\leftarrow \max\{\ell^{(k)}(S), \max\{h_i^{(k+1)} \mid E^{(k)}(i) = S\}\}$ 
10    for all  $i \in N$  do % move agents that bounced off
11      if  $i \in N^{(k)}$  then
12         $E^{(k+1)}(i) \leftarrow i$ 's next equivalence class
13         $h_i^{(k+1)} \leftarrow 0$ 
14      else
15         $E^{(k+1)}(i) \leftarrow E^{(k)}(i)$ 
16     $k \leftarrow k + 1$ 
17 return lottery  $p$  with  $p(S) \geq \ell^{(k+1)}(S)$  for all  $S \subset A$ 

```

Algorithm 1: The *ESR* algorithm.

Input: k
Output: λ^*, N^*

```

1  $N^* \leftarrow \{\}$ 
2 solve the following Linear Program (LP):
    $\lambda^* \leftarrow \max \lambda$ 
    $\sum_{a \in S} p_a \geq \ell^{(k)}(S) \quad \forall S \subset A$ 
    $\sum_{a \in E^{(k)}(i)} p_a \geq h_i^{(k)} + \lambda \quad \forall i \in N$ 
    $\sum_{a \in A} p_a \leq 1$ 
    $\forall a \in A : p_a \geq 0, \lambda \geq 0$ 
3 for all  $i \in N$  do
4   run the LP again imposing  $\lambda = \lambda^*$  and maximizing
   the slack in the inequality corresponding to  $i$ 
5   if objective function value = 0 then
6      $N^* \leftarrow N \cup \{i\}$ 

```

Algorithm 2: The function `computeLambda`.

class. This can happen at most $|N| \cdot |A|$ times, which means

that the `while`-loop will be executed no more than $|N| \cdot |A|$ times. All the statements inside the loop only take polynomial time. Thus, every iteration and hence the entire algorithm runs in polynomial time. \square

For dichotomous preferences, Bogomolnaia, Moulin, and Stong (2005) defined the egalitarian rule as the SDS that maximizes the leximin ordering (the lexicographic ordering over utility profiles, i.e. probabilities for each agent's approved equivalence class, rearranged increasingly). We point out that *ESR* generalizes their egalitarian rule to the case of general preferences.

Theorem 3. *For dichotomous preferences, ESR is equivalent to the egalitarian rule of (Bogomolnaia, Moulin, and Stong 2005).*

Proof. Let p be a lottery maximizing the leximin ordering. Denote the corresponding utility profile by $u(p) = (\underbrace{u_1, \dots, u_1}_{n_1}, \underbrace{u_2, \dots, u_2}_{n_2}, \dots, \underbrace{u_k, \dots, u_k}_{n_k})$ where $n = \sum_{i=1}^k n_i$ and $u_1 < u_2 < \dots < u_k$. We prove by induction on the number of stages of *ESR* that the guarantees of equivalence classes in *ESR* are fulfilled exactly by all leximin orderings.

Consider the point in time $t_1 := u_1$. At this point, no agent was forced to move away from his most preferred equivalence class (as p certifies that an assignment exists where every agent gets at least u_1 probability for his most preferred equivalence class). By the same argument, for no agent $i \notin \{1, \dots, n_1\}$ can the maximal equivalence class get tight at t_1 .

On the other hand, for all agents $i \in \{1, \dots, n_1\}$ their maximal equivalence classes must get tight at t_1 . Otherwise there would exist a lottery with at least u_1 probability for all agents in $\{1, \dots, n_1\}$ where some agent $i^* \in \{1, \dots, n_1\}$ (and all agents $> n_1$) gets strictly more than u_1 , this lottery would leximin-dominate p .

We can now remove the agents $1, \dots, n_1$ from consideration (as the guarantees that they accumulate for their last equivalence class don't restrict the set of possible lotteries) and iteratively repeat the argument for all points in time $t_j := u_j$ for $j \leq k$. \square

Axiomatic Properties

We first observe that *ESR* satisfies the minimal requirements of anonymity and neutrality. Another very important property of an SDS is efficiency, i.e. the outcome should be such that no other outcome is preferred by all agents. For strict preferences, the arguably most desirable mechanism *RD* is *DL*-efficient. However, *RSD* and the *maximal recursive rule (MR)* (Aziz 2013) — the two known generalizations of *RD* to non-strict preferences — are not even *SD*-efficient. We show that *ESR* is *DL*-efficient.

Theorem 4. *ESR is DL-efficient.*

Proof. Let p be a lottery obtained by *ESR*. Suppose, p is *DL*-dominated by another lottery q . Denote by t the time when an equivalence class E is frozen to a value smaller

than $q(E)$ for the first time. As E was frozen at t , there is no assignment p' with $p'(E) > \ell_t(E)$ and $p'(T) \geq \ell_t(T)$ for all other equivalence classes T . This means that $q(E') < \ell_t(E') \leq p(E')$ for some equivalence class E' of an agent i that is weakly preferred by i to the equivalence class that he is currently pointing to. As no equivalence class i has previously pointed to gets a higher probability in q (E was the first such equivalence class), q is a DL -disimprovement for i , contradicting our original assumption. \square

The only other SDS for the general preference domain that is known to be SD -efficient is SML . However, this SDS is not monotonic (see Proposition 7 in (Aziz, Brandt, and Brill 2013b)), a property that is fulfilled by ESR . We omit the proof due to space constraints.

Theorem 5. *ESR is monotonic.*

Turning to fairness, a common requirement in many settings is *fair outcome share*⁴: Each agent should get at least $1/n$ of the total utility it can get from any lottery. ESR satisfies the fair outcome share property.

Lemma 1. *If p is an outcome of ESR and E is some agent's most preferred equivalence class, then $p(E) \geq 1/n$.*

Corollary 1. *ESR satisfies the fair outcome share property.*

Finally, we investigate the robustness of ESR to strategic behavior. Here, the results are mixed: ESR satisfies the very basic notion of weak ST -strategyproofness but fails on the stronger (and more common) SD -based notions.⁵

Theorem 6. *For $|N| \leq 2$, ESR is SD -strategyproof. For $|N| \geq 3$, the following holds: (i) ESR is weak SD -strategyproof but not SD -strategyproof for strict preferences (ii) ESR is not weak SD -strategyproof for dichotomous preferences (and thus not for the general domain) (iii) ESR is weak ST -strategyproof.*

Proof. We omit the easy two-agent case. (i) follows easily from Observation 1 and Gibbard's characterization of RD -schemes as ex-post-efficient and SD -strategyproof (Gibbard 1977). Statement (ii) is proven by the following example:

$$1 : \{a, b\}, \{c, d\} \quad 2 : b, \{a, c, d\} \quad 3 : d, \{a, b, c\}$$

The outcome of ESR is $\frac{1}{2}b + \frac{1}{2}d$. However if agent 1 reports, $a, \{b, c, d\}$, then the outcome is $\frac{1}{3}a + \frac{1}{3}b + \frac{1}{3}d$.

Finally, for (iii) let p be an outcome of ESR and i some agent. Denote i 's most preferred equivalence class by E . Then, by Lemma 1, $p(E) \geq 1/n$. But by the characterization of ST in Proposition 2 of (Aziz, Brandt, and Brill 2013b), no lottery assigning positive probability to i 's most preferred equivalence class can be dominated by any other lottery. \square

⁴For dichotomous preferences, fair outcome share was defined by Bogomolnaia, Moulin, and Stong (2005). In the context of cake cutting, fair outcome share is also known as 'proportionality' (Brams and Taylor 1996).

⁵However, Aziz, Brandt, and Brill (2013b) suggest that these results are likely to be due to a more fundamental incompatibility of SD -efficiency and SD -strategyproofness if anonymity is required. At least, as the above argument also holds for weak ST -group-strategyproofness, ESR joins SML as the only SDSs known to be both SD -efficient and weak ST -group-strategyproof.

Assignment Domain

Our central result in this section is that for the assignment domain, ESR is equivalent to EPS , the generalization of PS to the domain where agents may express ties between objects. Given a set of objects O with $|O| = |N|$ the set of alternatives is the set of deterministic assignments of objects to agents, or in other words the set of permutations of the set $|O|$. We assume that agents have (complete and transitive) preferences over objects and denote an agent i 's upper contour set by $U_i(o) := \{o' \in O \mid o' \succsim_i o\}$ for all objects $o \in O$. This implies natural preferences over assignments: An agent prefers an assignment in which he gets an object o over an assignment in which he gets o' iff he prefers o over o' . We denote the set of assignments where agent i obtains object o by $A(i, o)$. Note that this yields a 1-to-1-correspondence of equivalence classes of objects to equivalence classes of deterministic assignments.

We denote by $E_i(o)$ the (unique) equivalence class of agent i that contains assignments where he receives object o . Denote by $U_i(E)$ the set of alternatives that is weakly preferred by agent i to all alternatives in E .

Next we prove that ESR is equivalent to EPS for the assignment domain. Instead of dealing directly with the definition of EPS , we will use the result of Heo and Yilmaz (2012) that a *correspondence* is a *subcorrespondence* of EPS if and only if it satisfies *SD-efficiency*, *limited invariance*, and *SD-envy-freeness*. An assignment satisfies *SD-envy-freeness* iff each agent weakly SD -prefers his allocation to any other agent's allocation. A rule satisfies *limited invariance* iff for each agent $i \in N$ and object $o \in O$, the total probability of his receiving the objects in $U_i(o)$ doesn't change if in his preference relation only the preferences among elements not in $U_i(o)$ are changed.

Theorem 7. *For the assignment domain, the ESR correspondence is SD -envy-free and fulfills limited invariance.*

Proof. We deal with both parts of the statement separately.

Limited invariance. Due to limited space, we omit the straightforward argument for limited invariance.

SD -envy-free Let $i, j \in N$ and p a lottery that results from the application of ESR . Suppose i 's allocation does not SD -dominate j 's. Then there exists an object $o \in O$ such that

$$p\left(\bigcup_{o' \succsim_i o} A(j, o')\right) > p\left(\bigcup_{o' \succsim_i o} A(i, o')\right). \quad (1)$$

Denote by t^* the point in time when equivalence class $E_i(o)$ is frozen and note that

$$p\left(\bigcup_{o' \succsim_i o} A(i, o')\right) = \sum_{\substack{E \in \mathcal{E}_i \\ E \succsim_i E_i(o)}} p(E) \geq \sum_{\substack{E \in \mathcal{E}_i \\ E \succsim_i E_i(o)}} \ell_{t^*}(E) = t^*$$

as all equivalence classes E of i with $E \succsim_i E_i(o)$ are already frozen. The sum of lower bounds for all sets from \mathcal{E}_j at time t^* is also equal to t^* , thus (1) implies that, for some object $o^* \succsim_i o$, we can choose $\varepsilon > 0$ such that $p(A(j, o^*)) - \varepsilon \geq \ell_{t^*}(E_j(o^*))$.

Now, consider the fractional assignment that is induced by lottery p . Note that for every agent i' and object o' the probability of receiving o' is exactly $p(A(i', o'))$. We construct

a new fractional assignment where the probability of agent j receiving object o is reduced by ε and this probability is transferred to agent i receiving object o . By the Birkhoff-von-Neumann-Theorem (Lovász and Plummer 2009), we can represent this new fractional assignment as a lottery q over discrete assignments. Note, that for all $(i', o') \in N \times O \setminus \{(i, o^*), (j, o^*)\}$ we have $p(A(i', o')) = q(A(i', o'))$, in particular that means that q , just as p fulfills all guarantees at time t^* except those for $E_{o^*}^i$ and $E_{o^*}^j$. But $q(E_j(o^*)) = p(E_j(o^*)) - \varepsilon \geq p(A(j, o^*)) - \varepsilon \geq \ell_{t^*}(E_j(o^*))$. Finally, we have $q(E_i(o^*)) = p(E_i(o^*)) + \varepsilon \geq \ell_{t^*}(E_i(o^*)) + \varepsilon$ which is a contradiction to the fact that $E_i(o^*)$ was frozen at time t^* (if $o^* \sim_i o$) or before (if $o^* \succ_i o$). \square

Corollary 2. *The ESR correspondence is equivalent to (i) the EPS correspondence for the assignment setting; (ii) PS to for the assignment problem with strict preferences over objects; (iii) the egalitarian mechanism of Bogomolnaia and Moulin (2004) for the assignment problem with dichotomous preferences over objects.*

Proof. Part (i) follows from the characterization of Heo and Yilmaz (2012) and the fact that *ESR* outputs a complete class of essentially equivalent lotteries. Parts (ii) and (iii) follow from the fact that *EPS* is equivalent to *PS* for strict preferences and to the egalitarian mechanism of Bogomolnaia and Moulin (2004) for dichotomous preferences. \square

Generalizing *ESR*

In this section, we consider the generalization of *ESR* that we call *simultaneous reservation (SR)* in which each agent i has a (piecewise) continuous climbing speed function s_i . The lower bound $\ell(E)$ of a class E is increased at a rate according to the maximal climbing speed among the agents currently touching the ceiling of E .

Characterizing efficiency notions has been an active of area of research (Manea 2008; McLennan 2002; Bogomolnaia and Moulin 2001; Abdulkadiroğlu and Sönmez 2003). We show that outcomes of *SR* characterize all *DL*-efficient lotteries. By the fact that all *serial dictator rules* (Aziz, Brandt, and Brill 2013b; Svensson 1994) are *DL*-efficient, they are thus a subclass of *SR*.

Theorem 8. *A lottery is DL-efficient if and only if it is the outcome of SR.*

Proof. The argument for *DL*-efficiency from Theorem 4 of *ESR* generalizes to *SR* immediately. For the converse, assume that a lottery p is *DL*-efficient. We claim that climbing speed functions s_i for all $i \in N$ exist such that p is the outcome of *SR*. We prove the claim by construction: The functions s_i that we construct will only take values 0 and $|N|$ and be such that for all $t \in [0, 1]$ all but one function will be 0, the agent whose function is non-zero at t will be called *active*.

We will call an agent *selectable* if he either hasn't been selected yet or if he cannot climb any further without bouncing off the ceiling. Whenever a selectable agent is activated, he will thus bounce off the ceiling immediately and move to his next equivalence class E (or to his first, if he hasn't

been active before). As soon as he reaches height $p(E)$, he will *deactivate*. Note that he might be able to climb higher in the tower E just yet, so he will only become selectable again as soon as other lower bounds have risen high enough to make it impossible for him to climb any higher. This procedure makes sure that no agent can ever guarantee more probability than $p(S)$ for any equivalence class S , this also guarantees that all agents can always reach height $p(E)$ as p certifies the existence of a suitable lottery.

We proceed by activating selectable agents until they deactivate, then a new selectable agent is activated. It only remains to be shown that there is always a selectable agent. For contradiction, suppose the opposite: All agents $i \in N$ can climb higher in their current towers E_i without bouncing off the ceiling. This means that there is a lottery q with $q(E) \geq p(E)$ for all equivalence classes E previously visited by the agents and $q(E_i) > p(E_i)$ for all $i \in N$. Thus, q *DL*-dominates p which is a contradiction. \square

Note that for *SR*, instead of selecting a single alternative, we can consider selecting a set (committee) of alternatives. The only change required is to increase the probability guarantees until the total probability weight over A is the size of the committee.

Conclusions

In this paper, we presented the *ESR* correspondence. *ESR* is interesting for a number of reasons: It simultaneously generalizes a number of very prominent mechanisms for restricted preference domains. On the complete preference domain, its properties compare well to other known SDSs: It is essentially single-valued and polynomial-time computable. It satisfies *DL*-efficiency (and thus *SD*-efficiently), monotonicity, fair outcome share and weak *ST*-strategyproofness. While in terms of strategyproofness, both *RSD* and *MR* perform much better, none of these mechanisms is even *SD*-efficient and furthermore, *RSD* is not polynomial-time computable (Aziz, Brandt, and Brill 2013a). *SML* on the other hand, the only other well-known *SD*-efficient SDS, performs no better in terms of strategyproofness (and even worse for strict preferences), is not *DL*-efficient and fails on monotonicity and fair outcome share.

One may point out that in the context of voting, *ESR* is too egalitarian: it gives as much weight to any minority as to the majority. In this sense, it constitutes the other extreme to *SML* which is oblivious to minority opinions. It will be interesting to identify subdomains other than assignment where *ESR* is an especially desirable rule. Finally, characterizing *ESR* is another direction for future work.

Acknowledgments

NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

- Abdulkadiroğlu, A., and Sönmez, T. 2003. Ordinal efficiency and dominated sets of assignments. *Journal of Economic Theory* 112(1):157–172.
- Aziz, H.; Brandt, F.; and Brill, M. 2013a. The computational complexity of random serial dictatorship. *Economics Letters* 121(3):341–345.
- Aziz, H.; Brandt, F.; and Brill, M. 2013b. On the tradeoff between economic efficiency and strategyproofness in randomized social choice. In *Proc. of 12th AAMAS Conference*, 455–462. IFAAMAS.
- Aziz, H.; Brandt, F.; and Stursberg, P. 2013. On popular random assignments. In Vöcking, B., ed., *Proc. of 6th International Symposium on Algorithmic Game Theory (SAGT)*, volume 8146 of LNCS, 183–194. Springer.
- Aziz, H. 2013. Maximal Recursive Rule: A New Social Decision Scheme. In *Proc. of 22nd IJCAI*, 34–40.
- Bogomolnaia, A., and Moulin, H. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100(2):295–328.
- Bogomolnaia, A., and Moulin, H. 2004. Random matching under dichotomous preferences. *Econometrica* 72(1):257–279.
- Bogomolnaia, A.; Moulin, H.; and Stong, R. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122(2):165–184.
- Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Brandt, F.; Conitzer, V.; and Endriss, U. 2013. Computational social choice. In Weiß, G., ed., *Multiagent Systems*. MIT Press, 2nd edition. chapter 6, 213–283.
- Chambers, C. 2004. Consistency in the probabilistic assignment model. *Journal of Mathematical Economics* 40:953–962.
- Cho, W. J. 2012. Probabilistic assignment: A two-fold axiomatic approach. Unpublished manuscript.
- Conitzer, V., and Sandholm, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. of 21st AAAI Conference*, 627–634. AAAI Press.
- Conitzer, V. 2010. Making decisions based on the preferences of multiple agents. *Communications of the ACM* 53(3):84–94.
- Fishburn, P. C. 1982. Monotonicity paradoxes in the theory of elections. *Discrete Applied Mathematics* 4(2):119–134.
- Gibbard, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica* 45(3):665–681.
- Heo, E. J., and Yilmaz, O. 2012. A characterization of the extended serial correspondence.
- Katta, A.-K., and Sethuraman, J. 2006. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory* 131(1):231–250.
- Kavitha, T.; Mestre, J.; and Nasre, M. 2011. Popular mixed matchings. *Theoretical Computer Science* 412(24):2679–2690.
- Kojima, F., and Manea, M. 2010. Incentives in the probabilistic serial mechanism. *Journal of Economic Theory* 145(1):106–123.
- Lovász, L., and Plummer, M. D. 2009. *Matching Theory*. AMS Chelsea Publishing.
- Manea, M. 2008. Random serial dictatorship and ordinally efficient contracts. *International Journal of Game Theory* 36:489–496.
- Manea, M. 2009. Asymptotic ordinal inefficiency of random serial dictatorship. *Theoretical Economics* 165–197.
- McLennan, A. 2002. Ordinal efficiency and the polyhedral separating hyperplane theorem. *Journal of Economic Theory* 105(2):435–449.
- Procaccia, A. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proc. of 24th AAAI Conference*, 836–841. AAAI Press.
- Saban, D., and Sethuraman, J. 2013. A note on object allocation under lexicographic preferences. *Journal of Mathematical Economics*.
- Schulman, L. J., and Vazirani, V. V. 2012. Allocation of divisible goods under lexicographic preferences. Technical Report arXiv:1206.4366, arXiv.org.
- Service, T. C., and Adams, J. A. 2012. Strategyproof approximations of distance rationalizable voting rules. In *Proc. of 11th AAMAS Conference*, 569–576.
- Svensson, L.-G. 1994. Queue allocation of indivisible goods. *Social Choice and Welfare* 11:323–330.