An Empirical Comparison of One-Sided Matching Mechanisms

Hadi Hosseini, Kate Larson, and Robin Cohen
Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
{h5hossei,klarson,rcohen}@uwaterloo.ca

ABSTRACT

For one-sided matching problems, two widely studied mechanisms are the Random Serial Dictatorship (RSD) and the Probabilistic Serial Rule (PS). Both mechanisms require only that agents specify ordinal preferences and have a number of desirable economic and computational properties. However, the induced outcomes of the mechanisms are often incomparable and thus there are challenges when it comes to deciding which mechanism to adopt in practice. In this paper, working in the space of general ordinal preferences, we provide empirical results on the (in)comparability of RSD and PS and analyze their respective economic properties. We then instantiate utility functions for agents, consistent with the ordinal preferences, with the goal of gaining insights on the manipulability, efficiency, and envyfreeness of the mechanisms under different risk attitude models.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics, Theory, Experimentation

Keywords

Mechanism Design, Matching, Random Assignment, Probabilistic Serial, Random Serial Dictatorship

1. INTRODUCTION

The problem of assigning a number of indivisible objects to a set of agents, in the absence of monetary transfers, is fundamental in many multiagent resource allocation applications, and has been the center of attention amongst researchers at the interface of artificial intelligence, economics, and mechanism design. Assigning dormitory rooms or offices to students, students to public schools, college courses to students, organs and medical resources to patients, members to subcommittees, etc. are some of the myriad examples of one-sided matching problems [31, 5, 13, 25].

Appears at: 3rd Workshop on Exploring Beyond the Worst Case in Computational Social Choice. Held as part of the 15th International Conference on Autonomous Agents and Multiagent Systems. May 9th-10th, 2016. Singapore.

Two important (randomized) matching mechanisms that only elicit ordinal preferences from agents are *Random Serial Dictatorship* (RSD) [2] and *Probabilistic Serial Rule* (PS) [11]. Both mechanisms have important economic properties and are practical to implement. The RSD mechanism has strong truthful incentives but guarantees neither efficiency nor envyfreeness. PS satisfies efficiency and envyfreeness; however, it is susceptible to manipulation. Therefore, there are subtle points to be considered when deciding which mechanism to use. For example, given a particular preference profile, the mechanisms often produce random assignments which are simply incomparable and thus, without additional knowledge of the underlying utility models of the agents, it is difficult to determine which is the "better" outcome. Furthermore, properties like efficiency, truthfulness, and envyfreeness can depend on whether there is underlying structure in the preferences, and even in general preference models it is valuable to understand under what conditions a mechanism is likely to be efficient, truthful, or envyfree as this can guide designers choices.

We study the comparability of PS and RSD when there is only one copy of each object, and analyze the space of all preference profiles for different combinations of agents and objects. We show that despite the inefficiency of RSD, the fraction of random assignments at which PS stochastically dominates RSD vanishes, especially when the number of agents is less than or equal to the available objects. We then instantiate utility functions for agents to gain insights on the manipulability, social welfare, and envyfreeness of the two mechanisms under different risk attitudes.

Our main result is that under risk aversion, the social welfare of RSD is as good as PS but RSD does create envy among the agents (though the fraction of envious profiles and total envy are small). Moreover, when the number of agents and objects are equal, RSD assignments are less likely to be dominated by PS and overall RSD assignments create negligible envy among agents. We also show that PS is highly susceptible to manipulation in almost all combinations of agents and objects. The fraction of manipulable profiles and the gain from manipulation rapidly increases, particularly when agents become more risk averse.

2. MODEL

In this section, we describe the basic one-sided matching problem and introduce the two mechanisms we study in detail, Random Serial Dictatorship (RSD) [2] and Probabilistic Serial Rule (PS) [11]. We then introduce a number of properties and criteria used to evaluate these mechanisms.

2.1 One Sided Matching

A one-sided matching problem consists of a set of n agents, N, and a set of m indivisible objects, M.[1] Each agent $i \in N$ has a private strict preference ordering, $\succ_i$, over M where $a \succ_i b$ indicates that agent i prefers to receive object a over object b. We represent the preference ordering of agent i by the ordered list of objects $\succ_i = a \succ_i b \succ_i c$ or $\succ_i = (abc)$, for short. We let $\mathcal{P}$ denote the set of all complete and strict preference orderings over M. A *preference profile* $\succ \in \mathcal{P}^n$ specifies a preference ordering for each agent, and we use the standard notation $\succ_{-i} = (\succ_1, \ldots, \succ_{i-1}, \succ_{i+1}, \ldots, \succ_n)$ to denote preferences orderings of all agents except i and thus $\succ = (\succ_i, \succ_{-i})$.

The goal in a one-sided matching problem is to assign the objects in M to the agents in N according to preference profiles, under the constraint that no object can be assigned to more than one agent. If $m = n$ then this means that each agent will receive exactly one object, however if $m < n$ then some agents will receive no object and if $m > n$ then some agents may receive multiple objects. An assignment is represented as a matrix

$$A = \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_n \end{pmatrix} = \begin{pmatrix} A_{1,1} & A_{1,2} & \ldots & A_{1,m} \\ A_{2,1} & A_{2,2} & \ldots & A_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ A_{n,1} & A_{n,2} & \ldots & A_{n,m} \end{pmatrix}$$

where $A_{i,j} \in [0,1]$ is the probability that agent i is assigned object j. We let $\mathcal{A}$ denote the set of all *feasible* assignments where an assignment $A \in \mathcal{A}$ is *feasible* if and only if $\forall j \in M$, $\sum_{i \in N} A_{i,j} = 1$. If $A \in \mathcal{A}$ is such that $A_{i,j} \in \{0,1\}$ then we say that A is a *deterministic* assignment; otherwise, A is a *random* assignment. Every random assignment can be represented as a convex combination of deterministic assignments [37], and thus we view random assignments as a probability distribution over a set of deterministic assignments.

2.2 Matching Mechanisms

In general, a *matching mechanism*, $\mathcal{M}$, is a mapping from the set of preference profiles, $\mathcal{P}^n$ to the set of feasible assignments, $\mathcal{A}$. That is, $\mathcal{M} : \mathcal{P}^n \mapsto \mathcal{A}$. In this paper, we focus our attention on two widely studied mechanisms for one-side matching: Random Serial Dictatorship (RSD) [2] and Probabilistic Serial Rule (PS) [10].

RSD relies on the concept of priority orderings over agents. Such an ordering is an ordered list of agents where the first agent gets to select its most preferred object from the set of objects, the second agent then selects its most preferred object from the set of remaining objects and so on, until no objects remain.[2] Given a preference profile $\succ \in \mathcal{P}^n$, RSD returns an assignment $RSD(\succ) \in \mathcal{A}$ which is a uniform distribution over all deterministic assignments induced from all possible priority orderings over the set of agents. RSD has been widely adopted for fair and strategyproof assignment for the school choice problem, course assignment, house allocation, and room assignment [1, 35, 2, 3]

PS treats objects as a set of divisible goods of equal size and simulates a simultaneous eating algorithm. Each agent starts "eating" its most preferred object, all at the same rate. Once an object is gone (eaten away) then the agent starts eating its next preferred object among the remaining objects. This process terminates when all objects have been "eaten". Given a preference profile $\succ \in \mathcal{P}^n$, $PS(\succ) \in \mathcal{A}$ is a random assignment where $A_{i,j}$ is the probability (fraction) that object j is assigned to (or "eaten by") agent i.

2.3 General Properties

In this section, we define key properties for matching mechanisms. In particular, we formally define efficiency, strategyproofness and envyfreeness for (randomized) matching mechanisms under ordinal preferences. To evaluate the quality of a random assignment, we use first-order stochastic dominance [18, 11]. Given a random assignment A_i, the probability that agent i is assigned an object that is at least as good as object ℓ is defined as follows

$$w(\succ_i, \ell, A_i) = \sum_{j \in M : j \succeq_i \ell} A_{i,j} \tag{1}$$

We say an agent always prefers assignment A_i to B_i, if for each object ℓ the probability of assigning an object at least as good as ℓ under A_i is greater or equal that of B_i, and strictly greater for some object.

Definition 1 (Stochastic Dominance). *Given a preference ordering $\succ_i$, random assignment A_i stochastically dominates (sd) assignment $B_i (\neq A_i)$ if*

$$\forall \ell \in M,\ w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, B_i) \tag{2}$$

A matching mechanism is *sd*-efficient if at all preference profiles $\succ \in \mathcal{P}^n$, for all agents $i \in N$, the induced assignment is not stochastically dominated by any other assignment.

Definition 2 (*sd*-Efficiency). *A random assignment is sd-efficient if for all agents, it is not stochastically dominated by any other random assignment.*

An important desirable property in matching mechanisms is strategyproofness, that is the mechanism is designed so that no agent has incentive to misreport its preferences.

Definition 3 (*sd*-Strategyproofness). *Mechanism $\mathcal{M}$ is sd-strategyproof if at all preference profiles $\succ \in \mathcal{P}^n$, for all agents $i \in N$, and for any misreport $\succ'_i \in \mathcal{P}^n$, such that $A = \mathcal{M}(\succ)$ and $A' = \mathcal{M}(\succ'_i, \succ_{-i})$, we have:*

$$\forall \ell \in M,\ w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, A'_i) \tag{3}$$

Sd-strategyproofness is a strict requirement. It implies that under any utility model consistent with the preference orderings, no agent can improve her expected utility by misreporting. We say that a mechanism is *weakly sd-strategyproof* if the inequality in Equation 3 is strict for some $\ell \in M$, but does not hold for all objects. Clearly, *sd*-strategyproofness implies weak *sd*-strategyproofness but the converse does not hold.

An assignment is *manipulable* if it is not *sd*-strategyproof. If there exists some agent who strictly benefits from the

[1]This problem is sometimes called the assignment problem or house allocation problem in the literature.

[2]For $n < m$, RSD requires a careful method for picking sequence at each realized priority ordering based on an arbitrary serial dictatorship quota mechanism, which directly affects the efficiency and envy of the assignments [19, 12]. For simplicity, we use the variant of RSD based on a *quasi-dictatorial* mechanism [29] where the first agent selects its most preferred $(m-n+1)$ objects, and the rest of the agents choose one object each.

	$n \geq m$		$n < m$	
	PS	RSD	PS	RSD
sd-strategyproof	weak	✓	✗	✓
sd-efficiency	✓	✗	✓	✗
sd-envyfree	✓	weak	✓	weak

Table 1: Properties of PS and RSD.

manipulation, (i.e. the mechanism is not even weakly sd-strategyproof) then we say the assignment is sd-manipulable.

Finally, we are interested in whether mechanisms are fair and use the notion of envyfreeness to this end. An assignment is sd-envyfree if each agent strictly prefers her random allocation to any other agent's assignment.

Definition 4 (sd-Envyfreeness). *Given agent i's preference $\succ_i$, assignment A_i is sd-envyfree if for all agents $\forall k \neq i \in N$,*

$$\forall \ell \in M,\ w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, A_k) \quad (4)$$

We say an assignment is weakly sd-envyfree if the inequality in Equation 4 is strict for some $\ell \in M$, but does not hold for all objects. A matching mechanism satisfies sd-envyfreeness if at all preference profiles $\succ \in \mathcal{P}^n$, it induces sd-envyfree assignments for all agents.

2.4 Properties of RSD and PS

The theoretical properties of PS and RSD have been well studied in the economics literature [11], and we summarize the results in Table 1. Both mechanisms are ex post efficient, that is, their realized outcomes cannot be improved without making at least one agent worse off. PS has been shown to be both sd-envyfree and sd-efficient. However, it is not even weakly sd-strategyproof when $n < m$ [22] and is only weakly sd-strategyproof when $n \geq m$. On the other hand, RSD is always sd-strategyproof, but it is only weakly sd-envyfree and is not sd-*efficient*. Example 1 illustrates the inefficiency of RSD.

Example 1. *Suppose there are four agents $N = \{1, 2, 3, 4\}$ and four objects $M = \{a, b, c, d\}$. Consider the following preference profile $\succ = ((abcd), (abcd), (badc), (badc))$. Table 2 shows the outcomes for $PS(\succ)$ and $RSD(\succ)$. In this example, all agents strictly prefer the assignment induced by PS over the RSD assignment. Thus, RSD is inefficient at this preference profile.*

	a	b	c	d
A_1	1/2	0	1/2	0
A_2	1/2	0	1/2	0
A_3	0	1/2	0	1/2
A_4	0	1/2	0	1/2

(a) Assignment under $PS(\succ)$

	a	b	c	d
A_1	5/12	1/12	5/12	1/12
A_2	5/12	1/12	5/12	1/12
A_3	1/12	5/12	1/12	5/12
A_4	1/12	5/12	1/12	5/12

(b) Assignment under $RSD(\succ)$

Table 2: Example showing the inefficiency of RSD

3. INCOMPARABILITY OF RSD AND PS

We argue that the theoretical findings on RSD and PS do not necessarily provide enough guidance to a market designer trying to select the correct mechanism for a specific setting. For example, while we know that PS is sd-efficient and RSD is not, this does not mean that PS always outperforms RSD.

	a	b	c
A_1	1/2	0	1/2
A_2	1/2	1/4	1/4
A_3	0	3/4	1/4

(a) Assignment under $PS(\succ)$

	a	b	c
A_1	1/2	0	1/2
A_2	1/2	1/6	1/3
A_3	0	5/6	1/6

(b) Assignment under $RSD(\succ)$

Table 3: Incomparability of RSD and PS

Example 2. *Suppose there are three agents $N = \{1, 2, 3\}$ and three objects $M = \{a, b, c\}$. Consider the following preference profile $\succ = ((acb), (abc), (bac))$. Table 3 shows $PS(\succ)$ and $RSD(\succ)$. Neither assignment dominates the other since agent 1 is ambivalent between the two assignments while agent 2 prefers $PS(\succ)$ and agent 3 prefers $RSD(\succ)$.*

If we knew the utility functions of the agents, consistent with their ordinal preferences, then we might be able to use the notion of (utilitarian) social welfare to help determine the better assignment.[3] However, it is easy to construct different utility functions for the agents in Example 2 where both RSD and PS maximize social welfare.

Similarly, the envy of RSD and the manipulability of PS both depend on the structure of preference profiles, and thus, a compelling question, that justifies the practical implications of deploying a matching mechanism, is to analyze the space of preference profiles to find the likelihood of inefficient, manipulable, or envious assignments under these mechanisms.

4. GENERAL PREFERENCES

The theoretical properties of PS and RSD only provide limited insight into their practical applications. In particular, when deciding which mechanism to use in different settings, the incomparability of PS and RSD leaves us with an ambiguous choice in terms of efficiency, manipulability, and envyfreeness. Thus, we examine the properties of RSD and PS in the space of all possible preference profiles as well as under lexicographic preferences. Lexicographic preferences are present in various applications and have been extensively studied in artificial intelligence and multiagent systems as a means of assessing allocations based on ordinal preferences [15, 32, 17]. Under lexicographic preferences, we denote the efficiency, strategyproofness, manipulablity, and envyfreeness with ld- (lexicographically dominate) prefix.

The number of all possible preference profiles is super exponential $(m!)^n$. For each combination of n agents and m objects we performed a brute force coverage of all possible preference profiles. Thus, for all subsequent figures each data point shows the fraction of all possible preference profiles. For the cases of $n = 10$ and $m \in \{9, 10\}$, we randomly generated 1,000 instances by sampling from a uniform preference profile distribution. For each preference profile, we ran both PS and RSD mechanisms and compared their outcomes in terms of the stochastic dominance relation. Note that not only is computing RSD probabilities #P-complete (and thus intractable) [6, 33], but checking the desire properties such as envyfreeness, efficiency, and manipulablity of random allocations is shown to be NP-hard for general settings [9, 8]. Thus, for larger settings even if we randomly

[3] Given utility functions for the agents, where $u_i(j)$ is the utility agent i derives from being assigned object j, the (utilitarian) social welfare of an assignment A is $\sum_i \sum_j A_{i,j} u_i(j)$.

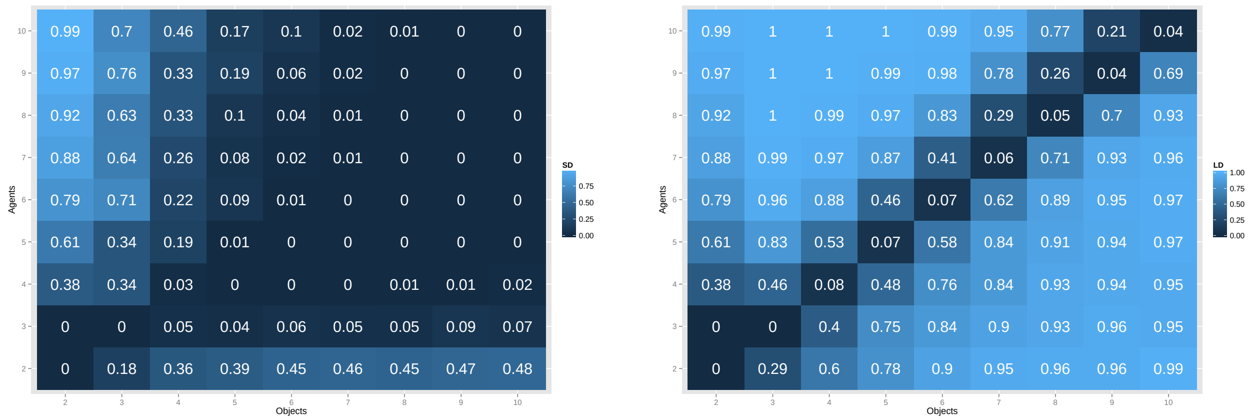

(a) The fraction that PS stochastically dominates RSD. (b) The fraction that PS lexicographically dominates RSD.

Figure 1: The fraction of preference profiles under which PS dominates RSD.

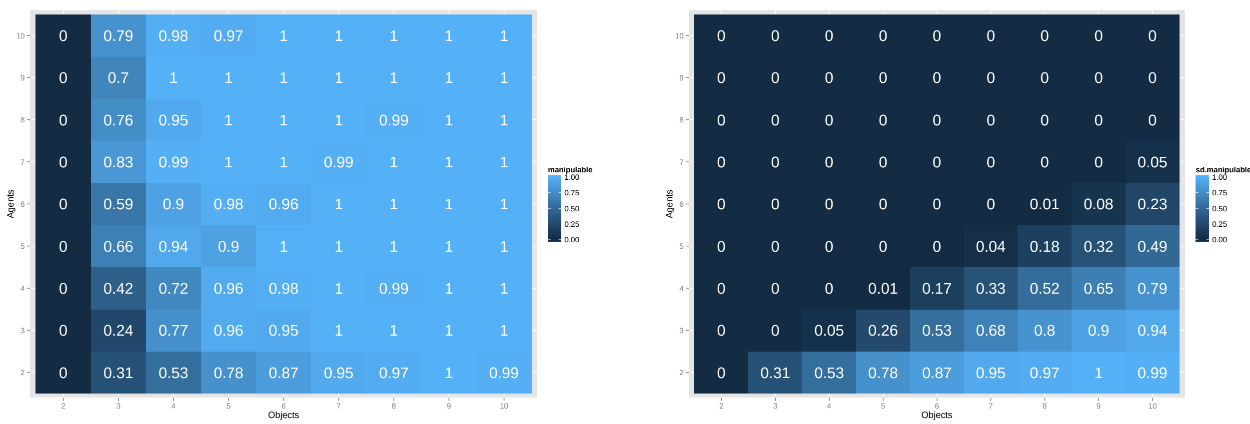

(a) The fraction of manipulable preference profiles under PS. (b) The fraction of *sd*-manipulable profiles under PS.

Figure 2: Heatmaps illustrating the manipulablity of PS.

sample preference profiles it is not easy to verify the aforementioned properties.

A preliminary look at our empirical results illustrates the following: when $m \leq 2, n \leq 3$, PS coincides exactly with RSD, which results in the best of the two mechanisms, i.e., both mechanisms are *sd*-efficient, *sd*-strategyproof, and *sd*-envyfree. Another interesting observation is that when $m = 2$, for all n PS is *sd*-strategyproof (although the PS assignments are not necessarily equivalent to assignments induced by RSD), RSD is *sd*-envyfree, and for most instances PS stochastically dominates RSD, particularly when $n \geq 4$.

4.1 Efficiency

Our first finding [4] is that the fraction of preference profiles at which RSD and PS induce equivalent random assignments goes to 0 when n grows. There are two conclusions that one can draw. First, this result confirms the theoretical results of Manea on asymptotic inefficiency of RSD [24], in that, in most instances, the assignments induced by RSD are not equivalent to the PS assignments. Second, this result suggests that the incomparability of outcomes is significant, that is, the social welfare of the random outcomes is highly dependent on the underlying utility models.

The fraction of preference profiles $\succ \in \mathcal{P}^n$ for which RSD is stochastically dominated by PS at $\succ$ converges to zero as $\frac{n}{m} \to 1$. Figure 1a shows that when n grows beyond $n > 5$, due to incomparability of RSD and PS with regard to the stochastic dominance relation, the RSD assignments are not stochastically dominated by *sd*-efficient assignments induced by PS.

We also see similar results when we restrict ourselves to lexicographic preferences (Figure 1b). The fraction of preference profiles $\succ \in \mathcal{P}^n$ for which RSD is lexicographically dominated by PS at $\succ$ converges to zero as $\frac{n}{m} \to 1$.

For lexicographic preferences, we also observe that the fraction of preference profiles for which PS assignments strictly dominate RSD-induced allocations goes to 1 when the number of agents and objects diverge. The fraction of preference profiles $\succ \in \mathcal{P}^n$ for which RSD is lexicographically dominated by PS at $\succ$ converges to 1 as $|n - m|$ grows.

[4] Periodically, we present results without a figure.

One immediate conclusion is that although RSD does not guarantee either sd-efficiency or ld-efficiency, in most settings when $\frac{n}{m} \to 1$ (and also $n \leq m$ for sd-efficiency), neither of the two mechanisms is preferred in terms of efficiency. Hence, one cannot simply rule out the RSD mechanism.

4.2 Manipulability of PS

One critical issue with deploying PS is that it does not provide incentives for honest reporting of preferences. Although for $n \geq m$ PS is weakly sd-strategyproof [11] and ld-strategyproof [34], when $n < m$ PS no longer satisfies these two properties.[5] The real concern is that, in the absence of strategyproofness, PS allocations are only efficient (or envyfree) with respect to the reported preferences, which in turn may not be truthful. Thus, we are interested in understanding the degree to which PS allocations are manipulable.

Figure 2 shows that the fraction of manipulable profiles goes to 1 as n or m grow. PS is almost 99% manipulable for $n > 5, m > 5$. Another interesting observation is that, for all $n < m$, the fraction of sd-manipulable preference profiles goes to 1 as $m - n$ grows (Figure 2b). These results imply that when agents are entitled to receive more than a single object, agents can strictly benefit from misreporting their preferences. The manipulability of PS under lexicographic preferences has a similar trend and the fraction of ld-manipulable preference profiles converges to 1 even more rapidly.

4.3 Envy in RSD

We measured the fraction of agents that are weakly sd-envious of at least one another agent when running RSD. Our results show that the percentage of agents that are weakly envious increases with the number of agents. Moreover, fixing any $n > 3$, the percentage of agents that are (weakly) envious grows with the number of objects; however, there is a sudden drop in the percentage of envious agents when there are equal number of agents and objects.

5. UTILITY MODELS

Given a utility model consistent with an agent's preference ordering, we can find the agent's expected utility for a random assignment. Let u_i denote agent i's Von Neumann-Morgenstern (VNM) utility model consistent with its preference ordering $\succ_i$. That is, $u_i(a) > u_i(b)$ if and only if $a \succ_i b$. Then, agent i's expected utility for random assignment A_i is $\mathbb{E}(u_i|A_i) = \sum_{j \in M} A_{i,j} u_i(j)$.

We say that agent i (strictly) prefers assignment A_i to B_i if and only if $\mathbb{E}(u_i|A_i) > \mathbb{E}(u_i|B_i)$. A mechanism is strategyproof if there exists no agent that can improve its expected utility by misreporting its preference ordering.

Definition 5 (Strategyproof). *Mechanism $\mathcal{M}$ is strategyproof if for all agents $i \in N$, and for any misreport $\succ'_i \in \mathcal{P}^n$, such that $A = \mathcal{M}(\succ)$ and $A' = \mathcal{M}(\succ'_i, \succ_{-i})$, given a utility model u_i consistent with $\succ_i$, we have $\mathbb{E}(u_i|A_i) \geq \mathbb{E}(u_i|A'_i)$.*

[5] A recent experimental study on the incentive properties of PS shows that human subjects are less likely to manipulate the mechanism when misreporting is a Nash equilibrium. However, subjects' tendency for misreporting is still significant even when it does not improve their allocations [20].

A matching mechanism is envyfree if for all preference profiles it prescribes an envyfree assignment.

Definition 6 (Envyfreeness). *Assignment A is envyfree if for all $i, k \in N$, given utility model u_i consistent with $\succ_i$, we have $\mathbb{E}(u_i|A_i) \geq \mathbb{E}(u_i|A_k)$.*

A random assignment A is sd-efficient if and only if there exists a profile of utility values consistent with $\succ$ such that A maximizes the social welfare ex ante [11, 26]. This existence result does not shed light on the social welfare when comparing two random assignments, since an assignment can be sd-efficient but may not have desirable ex ante social welfare. Given utility functions for the agents, the (utilitarian) *social welfare* of an assignment A is $\sum_i \mathbb{E}(u_i|A_i)$. Thus, given a profile of utilities we investigate the (ex ante) social welfare of the assignments under PS and RSD.

5.1 Instantiating Utility Functions

To deepen our understanding as to the performance of the two mechanisms, we investigate different utility models. In particular we look at the performance of the mechanisms when the agents are all risk neutral (i.e. have linear utility functions), when agents are risk seeking and when agents are risk averse.

Our first utility model is the well-studied linear utility model, and we use a variant based on the Borda rule from the social choice literature. Given an agent i's preference ordering $\succ_i$, we let $r(\succ_i, j)$ denote the rank of object j. For example, given preference ordering $a \succ_i b \succ_i c$ then $r(\succ_i, a) = 1$, $r(\succ_i, b) = 2$ and $r(\succ_i, c) = 3$. The utility function for agent i, given object j is $u_i(j) = m - r(\succ_i, j)$.

We use an *exponential* utility model to capture risk attitudes beyond risk-neutrality. An exponential utility has been shown to provide an appropriate translation for individuals' utility models and provides a constant risk aversion rate [4]. In particular,

$$u_i(j) = \begin{cases} (1 - e^{-\alpha(m - r(\succ_i, j))})/\alpha, & \alpha \neq 0 \\ m - r(\succ_i, j), & \alpha = 0 \end{cases} \tag{5}$$

The parameter α represents the agent's risk attitude. If $\alpha > 0$ then the agent is risk averse, while if $\alpha < 0$ then the agent is risk seeking. When $\alpha = 0$ then the agent is risk neutral and we have a linear utility model. The value $|\alpha|$ represents the intensity of the attitude. That is, given two agents with $\alpha_1 < \alpha_2 < 0$, we say that agent 1 is more risk averse than agent 2. Similarly if $\alpha_1 > \alpha_2 > 0$ then agent 1 is more risk seeking than agent 2.

6. RESULTS

For our experiments, we vary three parameters: the number of agents n, the number of objects m, and the risk attitude factor α. Each data point in the graphs shows the average over all possible preference profiles. We study the same settings as in Section 4 when $n \geq m$ and $n < m$. For each utility function, we look at homogeneous populations of agents where agents have the same risk attitudes.

To compare the social welfare, we investigate the percentage change (or improvement) in social welfare of PS compared to RSD under various utility models. That is, $\frac{\sum_i \mathbb{E}(u_i|PS(\succ)) - \sum_i \mathbb{E}(u_i|RSD(\succ))}{\sum_i \mathbb{E}(u_i|RSD(\succ))}$. To measure the manipulability of PS, we are interested in answering two key questions: *i*) In what fraction of profiles PS is manipulable by

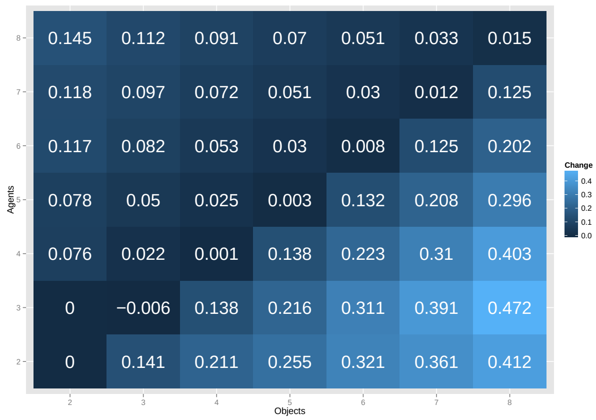

(a) Risk seeking, $\alpha = -0.5$.

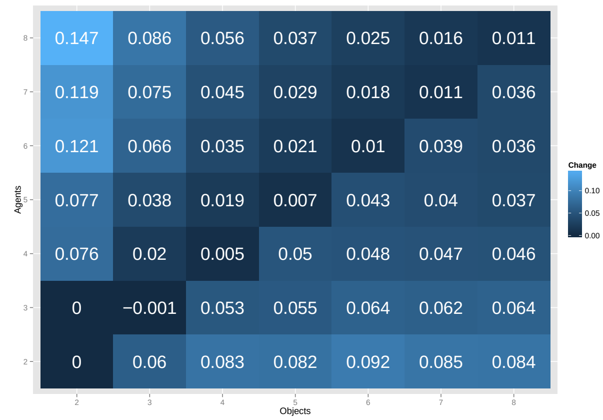

(b) Risk averse, $\alpha = 0.5$.

(c) Risk seeking, $\alpha = -2$.

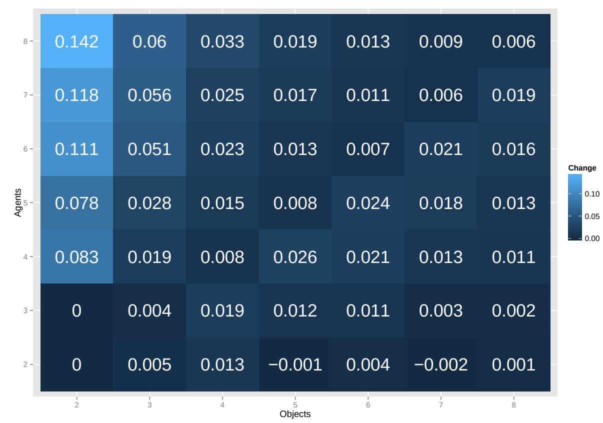

(d) Risk averse, $\alpha = 2$.

Figure 3: The percentage change in social welfare. The negative values show that RSD outperforms PS.

at least one agent? and ii) If manipulation is possible, what is the average percentage of maximum gain, that is $\max_i\{\frac{\mathbb{E}(u_i|PS(\succ_i', \succ_{-i})) - \mathbb{E}(u_i|PS(\succ))}{\mathbb{E}(u_i|PS(\succ))}\}$? To study the envy under the RSD mechanism, we consider two measures: i) the fraction of envious agents, and ii) the total envy felt by all agents.

6.1 Linear Utility Model

We first looked at how RSD and PS perform under the assumption that the utility models are linear. In most cases, the social welfare under PS increases compared to RSD; however, the percentage change from PS to RSD becomes smaller when $n = m$ (less than 0.015 overall improvement in all cases). Interestingly, under RSD the fraction of envious agents is approximately 0 when $n \geq m$. With regards to strategyproofness, PS is manipulable in most combinations of n and m and the fraction of manipulable profiles and the utility gain from manipulation increases as the number of objects compared to agents increases.

6.2 Risk Seeking

Figure 3 presents our results in terms of percentage change in social welfare. Positive numbers show the percentage of improvement in social welfare. Negative values represent those cases where RSD has increased social welfare compared to PS.

Social welfare: Fixing $\alpha < 0$, for $n \geq m$ when $\frac{n}{m}$ grows PS improves the social welfare compared to RSD in most cases and the percentage of improvement also increases. A similar trend holds when varying risk intensity α for fixed n and m where $n \neq m$. For $n < m$, when $\frac{m}{n}$ grows the fraction of profiles at which PS has higher social welfare compared to RSD rapidly increases and the percentage change is also noticeably larger, quickly getting close to 90% improvement (Fig. 3a and 3c). This social welfare gap between PS and RSD grows as the risk intensity $|\alpha|$ increases. Surprisingly, this trend changes for equal number of agents and objects $n = m$: the more risk-seeking agents are (larger $|\alpha|$), RSD becomes more desirable than PS, and in fact, RSD improves the social welfare in more instances.

Envy: For $n \geq m$, the fraction of envious agents under all profiles vanishes and RSD becomes envyfree. This is more evident when agents are more risk-seeking. Intuitively, these observations confirm the theoretical findings about the envyfreeness of RSD under lexicographic preferences [19] since one can consider lexicographic preferences as risk-seeking preferences where an object in a higher ranking is infinitely preferred to all objects that are ranked less

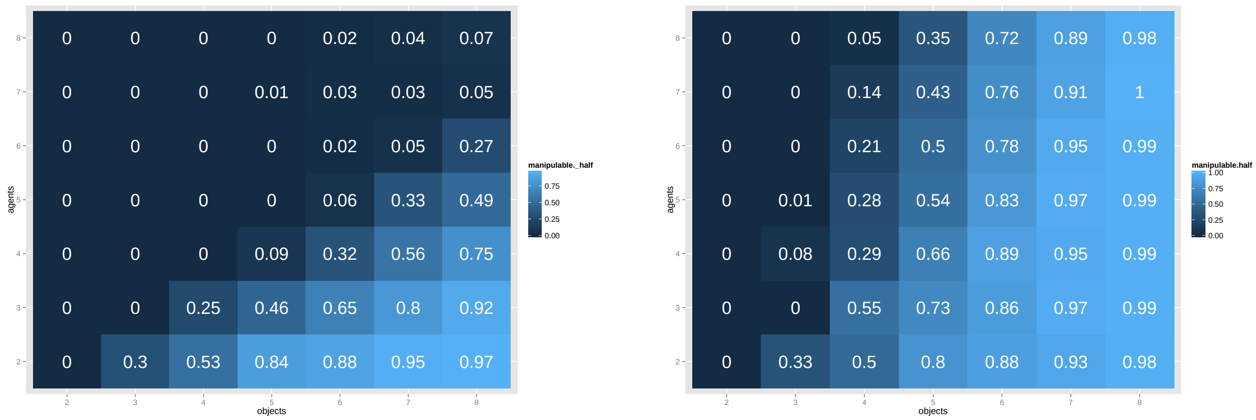

(a) Risk seeking, $\alpha = -0.5$. (b) Risk averse, $\alpha = 0.5$.

(c) Risk seeking, $\alpha = -2$. (d) Risk averse, $\alpha = 2$.

Figure 4: The fraction of manipulable instances under PS.

preferably. When $n < m$, our quasi-dictatorial extension of RSD creates some envy among the agents, but this envy also starts to fade out when the risk intensity $|\alpha|$ increases.

Manipulability: Figure 4 shows the manipulability of the PS assignments when agents are risk seeking. We see that the possibility of manipulation (and any gain) decreases as the risk intensity increases. When $n \geq m$ the fraction of manipulable profiles goes to 0 the more risk seeking agents become. However, when $n < m$ even though the the fraction of manipulable profiles (and manipulation gain) decreases, the fraction of manipulable profiles goes to 1 as $\frac{m}{n}$ grows.

6.3 Risk Aversion

Social welfare: Figures 3b and 3d show that fixing risk factor $\alpha > 0$, when $\frac{n}{m}$ grows PS assignments are superior to that of RSD in more instances and the percentage change in social welfare increases. Fixing risk factor $\alpha > 0$ and when $\frac{m}{n}$ grows, RSD is more likely to have the same social welfare as PS, and in fact in some instances the social welfare under RSD is better than the social welfare under PS. Fixing m and n, when the risk intensity α increases RSD is more likely to have the same social welfare as PS, that is, the welfare gap between PS and RSD closes when agents are more risk averse (α increases). This result is insightful and states that under risk aversion the random allocations prescribed by RSD are either as good as PS or in some cases even are superior to the allocations prescribed by PS. Figure 5 illustrates the percentage change in social welfare based on the difference between available objects and agents ($m-n$) for risk seeking, linear, and risk averse utilities with different risk intensities.

Envy: When $n \geq m$, the fraction of envious agents and total envy grows as $\frac{n}{m} \to 1$. Increasing the risk intensity ($|\alpha|$), the fraction of envious agents increases; however, the total envy among the agents remains considerably low. For $n < m$, the fraction of envious agents and total envy grows as risk intensity increases. Lastly, we noticed that in all instance where RSD creates envy among the agents, around 25% of agents bear more than 50% of envy. That is, few agents feel extremely envious while all other agents are either envyfree or only feel a minimal amount of envy.

Manipulability: Figures 4b and 4d illustrate the manipulability of the PS assignments when agents have risk averse preferences. The fraction of manipulable profiles rapidly goes to 1 as $\frac{m}{n}$ grows. Similarly, as agents become more risk averse (α increases) the fraction of manipulable profiles goes to 1 and the manipulation gain increases.

7. RELATED LITERATURE

Assignment problems with ordinal preferences have attracted interest from many researchers. Svensson showed that serial dictatorship is the only deterministic mechanism that is strategyproof, nonbossy, and neutral [36]. Random Serial Dictatorship (RSD) (uniform randomization over all serial dictatorship assignments) satisfies strategyproofness, proportionality, and ex post efficiency [2]. Bogomolnaia and Moulin noted the inefficiency of RSD from the ex ante perspective, and characterized the matching mechanisms based on first-order stochastic dominance [11]. They proposed the probabilistic serial mechanism as an efficient and envyfree mechanism with regards to ordinal preferences. While PS is not strategyproof, it satisfies weak strategyproofness for problems with equal number of agents and objects. However, PS is strictly manipulable (not weakly strategyproof) when there are more objects than agents [21]. Kojima and Manea, showed that in large assignment problems with sufficiently many copies of each object, truth-telling is a weakly dominant strategy in PS [22]. In fact PS and RSD mechanisms become equivalent [14], that is, the inefficiency of RSD and manipulability of PS vanishes when the number of copies of each object approaches infinity.

The practical implications of deploying RSD and PS have been the center of attention in many one-sided matching problems [1, 27]. In the school choice setting with multi-capacity alternatives, Pathak observed that many students obtained a more desirable random assignment through PS in public schools of New York City [30]; however, the efficiency difference was quite small. These equivalence results and their extensions to all random mechanisms [23], do not hold when the quantities of each object is limited to one.

Other interesting aspects of PS and RSD such as computational complexity and best-responses strategies have also been explored [16, 8, 7]. In this vein, Aziz et al. proved the existence of pure Nash equilibria, but showed that computing an equilibrium is NP-hard [7]. Nevertheless, Mennle et al. [28] showed that agents can easily find near-optimal strategies by simple local and greedy search. In the absence of truthful incentives, the outcome of PS is no longer guaranteed to be efficient or envyfree with respect to agents' true underlying preferences, and this inefficiency may result in outcomes that are worse than RSD, especially in 'small' markets [16].

8. DISCUSSION

We studied the space of general preferences and provided empirical results on the (in)comparability of RSD and PS. It is worth mentioning that at preference profiles where PS and RSD induce identical assignments, RSD is sd-efficient, sd-envyfree, and sd-strategyproof. However, PS is still highly manipulable. We investigated various utility models according to different risk attitudes. Our main results are:

- In terms of efficiency, the fraction of preference profiles $\succ \in \mathcal{P}^n$ for which PS stochastically (or lexicographically) dominates RSD converges to zero as $\frac{n}{m} \to 1$. When instantiating the preferences with actual utility functions, PS allocations are only slightly better than RSD allocations in terms of social welfare when varying n and m, particularly under risk averse utilities. In fact, in some cases RSD allocations are superior in terms of social welfare (see Figure 5).

Figure 5: The percentage change in social welfare between RSD and PS for $\alpha \in (-2, -1, -0.5, 0, 0.5, 1, 2)$ and different combinations of $m - n$. Positive α indicates risk averse and negative α risk taking profiles. Linear utility is indicated by $\alpha = 0$. As agents become more risk averse the social welfare gap between RSD and PS closes.

- PS is almost 99% manipulable when $n \leq m$ and the fraction of sd- and ld- manipulable profiles rapidly goes to 1 as $\frac{m}{n}$ grows. When instantiating the preferences with utility functions, the manipulability of PS increases as agents become more risk averse. Moreover, an agent's utility gain from manipulation also grows when the risk intensity increases.

- For risk seeking utilities, when $n \geq m$ the fraction of envious agents under all profiles vanishes and RSD becomes envyfree. For risk averse utilities, the fraction of envious agents increases as agents become more risk averse. However, the total amount of envy just slightly grows, and surprisingly, only few agents feel extremely envious while all other agents are either envyfree or only feel a minimal amount of envy.

Our work in this paper can be used to help guide designers of multiagent systems who need to solve allocation problems. If a designer strongly requires sd-efficiency then the theoretical results of PS indicate that it is better than RSD. However, our results show that PS is highly prone to manipulation for various combinations of agents and objects. This manipulation and the possible gain from manipulation become more severe particularly when agents are risk averse, and designers need to take this into consideration. On the other hand, while RSD does not theoretically guarantee sd-efficiency, our results show that it tends to do quite well – sometimes even outperforming PS in terms of social welfare. RSD also has the added advantage of being sd-strategyproof and thus is not prone to the manipulation problems of PS.

An interesting future direction is to study egalitarian social welfare of the matching mechanisms in single and multi unit assignment problems as well as in the full preference domain. Another open direction is to provide a parametric analysis of the matching mechanisms according to the risk aversion factor.

REFERENCES

[1] A. Abdulkadiroglu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifferences: redesigning the new york city high school match. Technical report, National Bureau of Economic Research, 2009.

[2] A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

[3] A. Abdulkadiroğlu and T. Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999.

[4] K. J. Arrow. Essays in the theory of risk-bearing. 1974.

[5] I. Ashlagi, F. Fischer, I. A. Kash, and A. D. Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, 2013.

[6] H. Aziz, F. Brandt, and M. Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013.

[7] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Equilibria under the probabilistic serial rule. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pages 1105–1112. AAAI Press, 2015.

[8] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, N. Narodytska, and T. Walsh. Manipulating the probabilistic serial rule. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, pages 1451–1459, 2015.

[9] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.

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

[11] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

[12] S. Bouveret and J. Lang. Manipulating picking sequences. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*, Prague, Czech Republic, August 2014. IOS Press.

[13] E. Budish and E. Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–71, 2012.

[14] Y.-K. Che and F. Kojima. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica*, 78(5):1625–1672, 2010.

[15] C. Domshlak, E. Hüllermeier, S. Kaci, and H. Prade. Preferences in AI: An overview. *Artificial Intelligence*, 175(7):1037–1052, 2011.

[16] Ö. Ekici and O. Kesten. An equilibrium analysis of the probabilistic serial mechanism. *International Journal of Game Theory*, pages 1–20, 2015.

[17] P. C. Fishburn. Lexicographic orders, utilities and decision rules: A survey. *Management Science*, pages 1442–1471, 1974.

[18] J. Hadar and W. R. Russell. Rules for ordering uncertain prospects. *The American Economic Review*, pages 25–34, 1969.

[19] H. Hosseini and K. Larson. Strategyproof quota mechanisms for multiple assignment problems. *arXiv preprint arXiv:1507.07064*, 2015.

[20] D. Hugh-Jones, M. Kurino, and C. Vanberg. An experimental study on the incentives of the probabilistic serial mechanism. Technical report, Discussion Paper, Social Science Research Center Berlin (WZB), Research Area Markets and Politics, Research Unit Market Behavior, 2013.

[21] F. Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009.

[22] F. Kojima and M. Manea. Incentives in the probabilistic serial mechanism. *Journal of Economic Theory*, 145(1):106–123, 2010.

[23] Q. Liu and M. Pycia. Ordinal efficiency, fairness, and incentives in large markets. *Unpublished mimeo*, 2013.

[24] M. Manea. Asymptotic ordinal inefficiency of random serial dictatorship. *Theoretical Economics*, 4(2):165–197, 2009.

[25] D. Manlove. *Algorithmics of matching under preferences*. World Scientific Publishing, 2013.

[26] A. McLennan. Ordinal efficiency and the polyhedral separating hyperplane theorem. *Journal of Economic Theory*, 105(2):435–449, 2002.

[27] T. Mennle and S. Seuken. Hybrid mechanisms: Trading off strategyproofness and efficiency in one-sided matching. Technical report, Working Paper, 2014.

[28] T. Mennle, M. Weiss, B. Philipp, and S. Seuken. The power of local manipulation strategies in assignment mechanisms. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pages 82–89. AAAI Press, 2015.

[29] S. Pápai. Strategyproof multiple assignment using quotas. *Review of Economic Design*, 5(1):91–105, 2000.

[30] P. A. Pathak. Lotteries in student assignment. *Unpublished mimeo, Harvard University*, 2006.

[31] A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.

[32] D. Saban and J. Sethuraman. A note on object allocation under lexicographic preferences. *Journal of Mathematical Economics*, 2013.

[33] D. Saban and J. Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005–1014, 2015.

[34] L. J. Schulman and V. V. Vazirani. Allocation of divisible goods under lexicographic preferences. *arXiv preprint arXiv:1206.4366*, 2012.

[35] T. Sönmez and M. U. Ünver. Course bidding at business schools. *International Economic Review*, 51(1):99–123, 2010.

[36] L.-G. Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.

[37] J. Von Neumann. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games*, 2:5–12, 1953.