Investigating the Characteristics of One-Sided Matching Mechanisms

(Extended Abstract)

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). The induced outcomes of these two mechanisms are often incomparable and thus there are challenges when it comes to deciding which mechanism to adopt in practice. Working in the space of general preferences, we provide empirical results on the (in)comparability of RSD and PS and analyze their economic properties.

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. Examples include assigning dormitory rooms or offices to students, courses to students, and medical resources to patients.

Two important (randomized) matching mechanisms that only elicit *ordinal preferences* from agents are *Random Serial Dictatorship* (RSD) [1] and *Probabilistic Serial Rule* (PS) [3]. 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, it is difficult to determine which is the "better" outcome.

We empirically 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 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, the *sd*-inefficiency of RSD does not mean that PS always outperforms RSD. 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 profiles to find the likelihood of inefficient, manipulable, or envious assignments under these mechanisms.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

2. MODEL AND PROPERTIES

Let N denote a set of n agents and M denote a set of m indivisible objects. 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 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 $\succ = (\succ_i, \succ_{-i})$.

In one-sided matching, if $m = n$ then each agent will receive exactly one object, however if $m > n$ then some agents may receive multiple objects. An assignment is represented as a matrix A where each element $A_{i,j} \in [0, 1]$ 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 $\forall j \in M$, $\sum_{i \in N} A_{i,j} = 1$.

A *matching mechanism*, $\mathcal{M}$, is a mapping from the set of preference profiles, $\mathcal{P}^n$ to the set of feasible assignments, $\mathcal{A}$. In this paper, we focus our attention on two widely studied mechanisms for one-side matching: Random Serial Dictatorship (RSD) [1] and Probabilistic Serial Rule (PS) [3].

RSD randomly chooses a priority ordering of agents where the first agent gets to select its most preferred object from the set of objects, the second agent selects its most preferred object from the remaining objects and so on until no objects remain.[1] PS treats objects as a set of divisible goods and simulates a simultaneous eating algorithm. Each agent starts "eating" its most preferred object, all at the same rate. Once an object is exhausted then the agent starts eating its next preferred object among the remaining objects, until all objects have been "eaten".

[1] For $n < m$, we use a variant of RSD based on *quasi-dictatorial* mechanisms [6, 4].

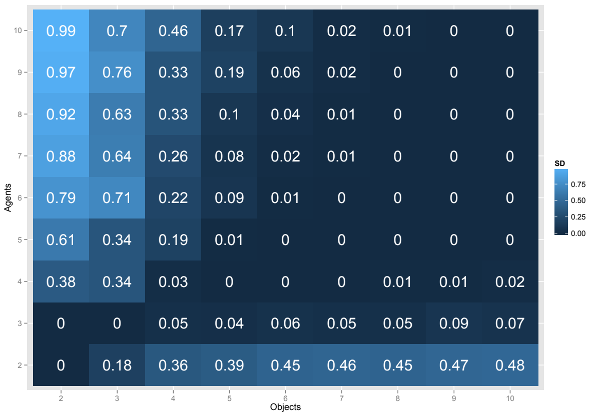

Figure 1: The fraction of preference profiles that PS stochastically dominates RSD.

3. PROPERTIES OF RSD AND PS

To evaluate the quality of a random assignment, we use first-order stochastic dominance [3]. According to our model, 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 $w(\succ_i, \ell, A_i) = \sum_{j \in M: j \succ_i \ell} A_{i,j}$. 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. Formally, A_i *stochastically dominates (sd)* $B_i (\neq A_i)$ if $\forall \ell \in M,\ w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, B_i)$.

A matching mechanism is sd-efficient if at all preference profiles $\succ \in \mathcal{P}^n$, for all agents $i \in N$, the induced random assignment is not stochastically dominated by any other assignment. An important desirable property in matching mechanisms is strategyproofness. A 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 $\forall \ell \in M,\ w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, A'_i)$. An assignment is *manipulable* if it is not sd-strategyproof. If there exists some agent who strictly benefits from the manipulation, then we say the assignment is sd-manipulable. Similarly, an assignment is sd-envyfree if each agent strictly prefers her random allocation to any other agent's assignment.

The theoretical properties of PS and RSD have been well studied in the literature [3, 5, 2]. PS has been shown to be both sd-envyfree and sd-efficient. However, it is sd-manipulable when $n < m$ [5], and is not sd-strategyproof when $n \geq m$. On the other hand, RSD is sd-strategyproof, but it does not guarantee *sd-efficiency*.

4. RESULTS

We studied the space of preference profiles and provided empirical results on the (in)comparability of RSD and PS under general and lexicographic preferences. Under lexicographic preferences, any fraction of a more preferred object is preferred to other less preferred objects. The number of all possible preference profiles is super exponential $(m!)^n$. For each combination of n agents and m objects we did a brute force coverage 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. Our main results are as follows:

General preferences: In terms of efficiency, the fraction of preference profiles $\succ \in \mathcal{P}^n$ for which PS stochastically dominates RSD converges to zero as $\frac{n}{m} \to 1$. PS is almost 99% manipulable when $n \leq m$ and the fraction of sd-manipulable profiles rapidly goes to 1 as $\frac{m}{n}$ grows (Figure 1).

Lexicographic preferences: 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.

5. DISCUSSION

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 can receive more than one object, and designers need to take this into consideration. We have also instantiated utility functions for agents under variety of risk attitude models to gain deeper insights on the manipulability, social welfare, and envyfreeness of PS and RSD. Our results show that while RSD does not theoretically guarantee sd-efficiency, 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.

REFERENCES

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

[2] 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.

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

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

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

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