Operations Research Letters 43 (2015) 283–287

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

A note on the assignment problem with uniform preferences

Jay Sethuraman *, Chun Ye

IEOR Department, Columbia University, New York, NY, United States

ARTICLE INFO

Article history:
Received 6 February 2015
Received in revised form
23 February 2015
Accepted 23 February 2015
Available online 3 March 2015

Keywords:
Random assignment
Probabilistic serial mechanism
Fairness
Efficiency

ABSTRACT

Motivated by a problem of scheduling unit-length jobs with weak preferences over time-slots, the random assignment problem is considered on a uniform preference domain. It is shown that the natural extension of the probabilistic serial mechanism to the domain of weak, but uniform, preferences fails strategy-proofness, but so does every other mechanism that is ordinally efficient and treats equals equally. If envy-free assignments are required, any ex-post efficient (probabilistic or deterministic) mechanism must fail even a weak form of strategy-proofness.

1. Introduction

We study the assignment problem, which is concerned with allocating objects to agents, each of whom wishes to receive at most one object. Agents have preferences over the objects, and the goal is to allocate the objects to the agents in a fair and efficient manner. Further, as each agent's preference ordering over the objects is private information, we require the mechanism to be strategyproof: it should be a dominant strategy for the agents to report their preference ordering truthfully. If the objects are divisible, we can think of a fractional assignment in which an object may be allocated in varying amounts to multiple agents so that the total amount allocated of any object is at most 1, and so that each agent receives at most one unit in all. If the objects are indivisible, one can think of a lottery over assignments, which again results in a fractional assignment matrix in which entry (i, a) represents the probability that agent i receives object a. These two views are equivalent for our purposes; while in the rest of the paper we assume that the objects are indivisible, all of our results extend to the case of divisible objects with the obvious change in interpretation. There is now a rich literature on such models with applications to many real-life allocation problems including allocating students to schools in various cities, the design of kidney exchanges, etc. [1,2,7,13] The two prominent mechanisms that have emerged from this literature are the *Random Priority* (RP) mechanism and the *Probabilistic Serial* (PS) mechanism. The PS mechanism is stronger in terms of its efficiency and equity properties, but it is only weakly strategyproof in the strict preference domain and not strategyproof in the full preference domain; whereas the RP mechanism is strategyproof, but satisfies only a weaker version of efficiency and envy-freeness. Furthermore, Bogomolnaia and Moulin [5] show that no strategyproof mechanism can satisfy the stronger form of efficiency and equity that the PS mechanism satisfies.

This paper is inspired by the paper of Bogomolnaia and Moulin [6], which characterizes the PS mechanism on a restricted preference domain. The PS mechanism was introduced in an earlier paper of Crés and Moulin [9] that was motivated by the problem of scheduling unit-length jobs with deadlines. Suppose there are n jobs, each requiring a unit processing time, and all jobs are available at time zero. As the jobs all have unit-length, one could think of the scheduling problem as one of assigning time-slots $1, 2, \ldots, n$ to the jobs, so that slot k represents the interval $(k-1, k]$, and a job assigned to slot k finishes at time k. Jobs have deadlines and earn a non-negative utility if they complete before their deadline. Specifically, if the deadline of job j is d_j, then the utility of assigning j to slot k is monotonically decreasing in k until the deadline, after which it drops to zero. That is, if $u_{j,k}$ denotes the utility of assigning job j to slot k, then

$$u_{j,1} > u_{j,2} \ldots > u_{j,d_j} > 0 = u_{j,d_j+1} = u_{j,d_j+2}, \ldots, u_{j,n}.$$

The goal is to use a mechanism to schedule the jobs in a fair and efficient manner based on their reported utility information without the usage of money. Crés and Moulin [9] proposed the PS mechanism and showed that it finds an ordinally efficient and envy-free allocation (all definitions appear in the next section); furthermore, they showed that the PS mechanism is strategyproof on this domain: in the event each job/agent need only report their

* Corresponding author.
E-mail addresses: jay@ieor.columbia.edu (J. Sethuraman), cy2214@columbia.edu (C. Ye).

http://dx.doi.org/10.1016/j.orl.2015.02.010
0167-6377/

deadline, they show that it is a weakly dominant strategy for each job to report its deadline truthfully. Bogomolnaia and Moulin [6] characterize the PS mechanism on this restricted domain in two different ways: first, they show that ordinal efficiency and envy-freeness characterize the PS outcome on this restricted domain; and second, they show that it is the only strategyproof mechanism that is ordinally efficient and treats equals equally. Taken together, their result shows that the PS mechanism is perhaps the only compelling mechanism on this restricted preference domain. (Crés and Moulin [9] show that the PS mechanism is in fact group strategyproof, although this stronger property is not needed in their characterization result of PS.)

In this paper we consider a slightly more general domain, again inspired by the problem of scheduling unit-length jobs. For simplicity, assume there are n agents and n objects, and suppose the objects are arranged in the order $(1, 2, \ldots, n)$ by all the agents. Each agent's preference ranking, however, is determined by a *weakly* decreasing utility function over the objects, in contrast to a strictly decreasing utility function over the objects till a deadline. (A good way to visualize this preference domain is to have each agent separate the sequence of objects into indifference classes, without disturbing the common order on the objects.) This domain is quite natural in the scheduling context, where completing a job early is always (weakly) better, but jobs may be insensitive to completion times within a certain time interval, and these intervals may change from job to job. The domain considered in the earlier papers is a special case in which, for each agent, all but the final indifference class has a single object. It is then natural to ask if the two characterizations of PS extend to this domain. It turns out that the answer is negative in each case. We show that the PS outcome (actually, a correspondence) is no longer the only outcome that is ordinally efficient and envy-free, nor is the PS mechanism strategyproof on this domain. Somewhat surprisingly, we show that:

- No weakly strategyproof mechanism can satisfy both ex post efficiency and envy freeness on this domain, when there are three or more agents; and
- No strategyproof mechanism can satisfy both ordinal efficiency and equal treatment of equals on this domain, when there are four or more agents.

The literature on random assignment problems focuses on simultaneously satisfying various notions of fairness, efficiency, and strategyproofness, and several impossibility results have been established over the last two decades [3,5,8,10,11,15]. Our two main impossibility results are strengthened versions of similar results in the literature in which preferences are drawn from richer domains. Specifically, versions of the two impossibility results have been obtained by Katta and Sethuraman [11] on the full preference domain (where any weak ordering of the objects is permissible), and by Bogomolnaia and Moulin [5] on the strict preference domain (where any *strict* ordering of the objects is permissible). Thus the surprising element in our result is that these difficulties persist even in domains in which the preferences are severely restricted.

2. Preliminaries

2.1. Model and definitions

An assignment problem is given by a triple $(N, O, \succsim)$, where $N = \{1, \ldots, n\}$ is the set of agents, $O = \{o_1, \ldots, o_n\}$ is the set of objects, and the preference profile $\succsim = (\succsim_1, \ldots, \succsim_n)$ specifies each agent's preference ordering over the objects. If the number of agents is not the same as the number of objects, one can always balance such a problem by adding dummy agents or dummy objects. We will assume that the preference relation of each agent is complete (every pair of objects is comparable) and transitive. By $a \succsim_i b$, we mean that agent i weakly prefers object a to object b. We write $a \succ_i b$ if i strictly prefers a to b, i.e. $a \succsim_i b$ but $b \not\succsim_i a$; and we use $a \sim_i b$ when i is indifferent between a and b, i.e. $a \succsim_i b$ and $b \succsim_i a$. Note that the indifference relation is also transitive. Thus each agent has a most-preferred subset of objects (and the agent is indifferent between all the objects within this set), followed by a most-preferred subset of objects among the remaining ones, etc.

In this paper, we shall consider the *uniform* preference domain in which $o_1 \succsim_i o_2 \succsim_i \ldots \succsim_i o_n$ for every agent $i \in N$. Agents differ in their preference ordering only in their strict preference relation $\succ_i$ (and hence their indifference relation $\sim_i$). In the rest of the paper, we use the following notation for the preference ordering of the agents: all the objects within an indifference class for an agent appear within braces in that agent's preference list, and these maximal indifference classes are separated by a comma; objects are always written in subscript order; and the braces are omitted for singleton indifference classes. Thus, the preference ordering

$o_1 \succ_i o_2 \sim_i o_3 \sim_i o_4 \succ_i o_5$

for agent i is written as

$i : o_1, \{o_2\ o_3\ o_4\}, o_5.$

By a mechanism, we mean a mapping from the set of all preference profiles (within this restricted domain) to a doubly stochastic matrix, which we call the assignment matrix for that profile. The assignment matrix is *deterministic* if its entries are $\{0, 1\}$ (and so the outcome is a *matching* of the agents and objects); otherwise, it is *probabilistic*. If a mechanism maps each preference profile to a deterministic matrix, the mechanism is deterministic; otherwise the mechanism is probabilistic. As a consequence of the Birkhoff–von Neumann theorem [4], the outcome of a probabilistic mechanism can be implemented as a lottery over deterministic assignments.

Given two probabilistic assignments P and Q, we say that agent i prefers P to Q if P_i, the ith row of P stochastically dominates Q_i according to i's preferences. Formally,

$$P_i \succsim_i Q_i \iff \sum_{k:k \succsim_i j} p_{ik} \geq \sum_{k:k \succsim_i j} q_{ik}, \quad \forall j \in O.$$

We say that i strictly prefers P to Q, denoted by $P_i \succ_i Q_i$, if at least one of the inequalities in the above definition is strict. Note that this definition is only a *partial* order, as an agent may not be able to compare two probabilistic allocations. Finally, we say that P *stochastically dominates* Q, denoted by $P \succsim Q$, if $P_i \succsim_i Q_i$ for all $i \in N$, with $P_i \succ_i Q_i$ for some $i \in N$. Again, this notion of stochastic dominance defines a partial order on the set of doubly stochastic matrices.

2.2. Desirable properties

We define some desirable properties of mechanisms that play an important role in the rest of the paper.

Ordinal efficiency. An assignment matrix P is *ordinally efficient* if it is not stochastically dominated by any other random assignment matrix Q such that $Q \succsim P$. It is well known that any ordinally efficient matrix can be implemented as a lottery over deterministic Pareto efficient assignments. Furthermore, checking whether or not a given assignment matrix is ordinally efficient is computationally easy [5,11].

Ex post efficiency. A weaker notion of efficiency that we will consider is ex post efficiency. A bi-stochastic matrix P is *ex post efficient* if it can be written as a convex combination of Pareto efficient assignments.

Envy-freeness. An assignment matrix P is *envy free* if the probabilistic assignment of every agent i stochastically dominates the probabilistic assignment of every other agent with respect to agent i's preference ordering. Let P_i denote the probabilistic assignment of agent i in the matrix P. Then, P is envy-free if $P_i \succsim_i P_{i'}$ for all $i, i' \in N$.