Journal of Mathematical Economics 50 (2014) 283–289

Contents lists available at ScienceDirect

Journal of Mathematical Economics

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

A note on object allocation under lexicographic preferences

Daniela Saban[a], Jay Sethuraman[b,*]

[a] Graduate School of Business, Columbia University, New York, NY, United States

[b] IEOR Department, Columbia University, New York, NY, United States

ARTICLE INFO

Article history:
Received 30 January 2013
Received in revised form
25 November 2013
Accepted 3 December 2013
Available online 12 December 2013

Keywords:
Object allocation
Lexicographic preferences
Strategy-proofness
Envy-freeness
Efficiency

ABSTRACT

We consider the problem of allocating m objects to n agents. Each agent has unit demand, and has strict preferences over the objects. There are q_j units of object j available and the problem is balanced in the sense that $\sum_j q_j = n$. An allocation specifies the amount of each object j that is assigned to each agent i, when the objects are divisible; when the objects are indivisible and exactly one unit of each object is available, an allocation is interpreted as the probability that agent i is assigned one unit of object j. In our setting, agent preferences over objects are extended to preferences over allocations using the natural lexicographic order. The goal is to design mechanisms that are efficient, envy-free, and strategy-proof. Schulman and Vazirani show that an adaptation of the probabilistic serial mechanism satisfies all these properties when $q_j \geq 1$ for all objects j. Our first main result is a characterization of problems for which efficiency, envy-freeness, and strategy-proofness are compatible. Furthermore, we show that these three properties *do not* characterize the serial mechanism. Finally, we show that when indifferences between objects are permitted in agent preferences, it is impossible to satisfy all three properties even in the standard setting of "house" allocation in which all object supplies are 1.

1. Introduction

We consider the problem of allocating m objects to n agents. Each agent i has unit demand, and an amount q_j, not necessarily integer, of object j is available. We assume that $\sum_j q_j = n$, so that supply and demand are balanced. An allocation $A = (a_{ij})$ is a non-negative $n \times m$ matrix in which each row i adds up to 1 and each column j adds up to q_j. The ith row of the allocation matrix A is also referred to as agent i's allocation.[1] If the objects are divisible, a_{ij} is simply the amount of object j that agent i receives; if the objects are indivisible and exactly one unit of each object is available, we may interpret a_{ij} as the probability that agent i receives object j.[2]

Each agent i has a strict preference ordering P_i over the set of objects: for objects j and k, we say that $j >_{P_i} k$ (or simply $j >_i k$) if and only if agent i strictly prefers j to k. This strict preference ordering over objects can be extended to a preference ordering over allocations in many different ways. In this paper we shall restrict our attention to the lexicographic extension, first considered in this context by Cho (2012) for the probabilistic assignment of indivisible objects, and, independently, by Schulman and Vazirani (2012) for allocating divisible objects. Given two allocations A and A', agent i prefers the one in which he receives more of his most-preferred object; if he receives the same amount of his most-preferred object in both, he prefers the one in which he receives more of his second most-preferred object, etc. To state this formally, suppose that agent i's preference ordering over the objects is $1 >_i 2 >_i \cdots >_i m$. Given two allocations A and A', agent i lexicographically prefers A to A' if and only if the first non-zero entry in $(a_{i1} - a'_{i1}, a_{i2} - a'_{i2}, \ldots, a_{im} - a'_{im})$ is positive. Of course, agent i is indifferent between A and A' if and only if i's allocation is the same in A and A'. Note that the preference ordering is *complete* in the sense that given two different allocations for agent i, he will always prefer one to the other.

A *problem* is completely specified by the set of agents N along with their preferences, and the set of objects with the corresponding q_j. A *mechanism* is a function that maps each problem to an allocation. Typically, the goal is to design a mechanism that satisfies certain desirable properties. To state these properties formally, it is important to define the notion of *dominance*.

Definition 1 (*Dominance*)**.** An allocation A' dominates an allocation A if each agent i (weakly) prefers his allocation in A' to that in A.

* Corresponding author.
E-mail addresses: dhs2131@columbia.edu (D. Saban), jay@ieor.columbia.edu (J. Sethuraman).

[1] We use the term "allocation" for both the matrix A, which describes the allocation for all agents, and for any of its rows, which describes the allocation for the corresponding agent.

[2] This case corresponds to the well-studied "house allocation" problem (Bogomolnaia and Moulin, 2001; Katta and Sethuraman, 2006; Sonmez and Ünver, 2010).

0304-4068/$ – see front matter
http://dx.doi.org/10.1016/j.jmateco.2013.12.002

Note that each agent compares his allocation in A to that in A' using the lexicographic extension of his preference ordering over the objects. Using this definition of dominance, we next define *efficiency* and *envy-freeness* for allocations.

Definition 2 (*Efficiency*)**.** An allocation A is *efficient* if it is not dominated by any other allocation A'.

Definition 3 (*Envy Freeness*)**.** An allocation A is envy-free if each agent i (weakly) prefers his allocation to the allocation of any other agent.

A mechanism is efficient if it associates an efficient allocation with each problem, and it is envy-free if it associates an envy-free allocation with each problem.[3] In other words, a mechanism is said to be *efficient* if, in any problem, there is no way to strictly improve one agent's allocation without making another agent worse off; it is *envy-free* if each agent prefers his allocation to that of any other agent in any problem. Our final desideratum for a mechanism is strategy-proofness. Informally, a mechanism is strategy-proof if it is not possible for any agent to improve his allocation by falsifying his preferences, regardless of the preferences of the other agents. Formally:

Definition 4 (*Strategy-Proofness*)**.** For each agent i, and for each profile of preferences $P_{-i} = (P_1, P_2, \ldots, P_{i-1}, P_{i+1}, P_{i+2}, \ldots, P_n)$, let A_i and A'_i be the allocations of agent i when the mechanism is applied to the profiles (P_i, P_{-i}) and (P'_i, P_{-i}) respectively. A mechanism is *strategy-proof* if $A_i \geq_{P_i} A'_i$ for all i, P_i, P_{-i}, and P'_i.

Bogomolnaia and Moulin (2001) proposed an important new mechanism for the house allocation problem that they called the probabilistic serial (PS) mechanism. The specific preference model they assumed was the one based on first order stochastic dominance (hereafter, the *sd*-preference model): an agent i prefers an allocation A to an allocation A' if A stochastically dominates A', or equivalently, agent i's expected utility of A is greater than that of A' for *every* utility function consistent with his preference ordering. Although this notion leads only to a partial ordering over allocations for each agent, they were able to show that the outcome of the PS mechanism was strongly envy-free in the sense that each agent's allocation stochastically dominates the allocation of any other agent. Furthermore, they proved that the PS mechanism was efficient (no other allocation matrix stochastically dominates the PS outcome), but only weakly strategy-proof (by misreporting his preference ordering, an agent cannot obtain an allocation that stochastically dominates his true allocation). Moreover, they show that no mechanism satisfies efficiency, envy-freeness and strategy-proofness in the *sd*-preference model.

Given that no mechanism is efficient, envy-free, and strategy-proof in the *sd*-preference model, one can consider relaxations of these requirements. The model of this paper, due to Cho (2012) and, independently, Schulman and Vazirani (2012), can be seen as a relaxation of the stochastic dominance preference model, in which agents compare allocations based on lexicographic preferences. The key motivation for considering lexicographic preferences is three-fold: first, lexicographic preferences form a complete relation so that *every* pair of allocations can be compared; second, it is a weakening of the stochastic dominance preference model in the sense that stochastic dominance implies lex-dominance; and finally, it can be interpreted as a limiting case of agents with von Neumann–Morgenstern utilities.[4]

Cho (2012) observed that the probabilistic serial mechanism of Bogomolnaia and Moulin (2001) satisfies all three properties when $m = n$, $q_j = 1$ for all j, and the objects are indivisible. Independently, Schulman and Vazirani (2012) studied the lexicographic preference model when agents' demands and objects' supplies are arbitrary[5] and the objects are divisible. Their algorithm – called synchronized greedy (SG) – is efficient in all cases, envy-free when the agents have unit (or equal) demands, and strategy-proof when the maximum amount demanded by an agent is at most the minimum supply of an object. Specialized to our setting then, the SG mechanism is efficient and envy-free, and is strategy-proof whenever $q_j \geq 1$ for all j.

Our first result, proved in Section 2, is that when this last condition is not satisfied *no* mechanism can satisfy all three properties. Specifically, we show that if there is an object j with $q_j < 1$, then there is an instance involving as many agents as objects for which no strategy-proof mechanism can be both efficient and envy-free. In proving this result we assume that we are free to *choose* all the parameters of the problem: the number of agents, the number (and supply) of objects, and the preference ordering of the agents. A natural question is if a similar result is possible for a given number of agents and a fixed set of objects with a given supply vector, with only the agent preferences varying. For this setting, we characterize when it is possible to find a mechanism satisfying all three properties: we show that either the synchronized greedy mechanism is strategy-proof for every preference profile (and hence satisfies all three properties) or *no* mechanism can satisfy all three properties. One way to interpret this result is that whenever the SG mechanism fails strategy-proofness, so does every other mechanism that is envy-free and efficient. Our next two results answer questions left open by Schulman and Vazirani (2012): in Section 3, we show that the synchronized greedy mechanism is not the only mechanism with these properties by constructing a "hybrid" mechanism that incorporates elements of a greedy mechanism and the SG mechanism. This new mechanism satisfies an invariance property called *bounded invariance* (defined formally in Section 3), in addition to being efficient, envy-free and strategy-proof. An important consequence is that envy-freeness, efficiency, and bounded invariance do not characterize the SG mechanism when agents have lexicographic preferences; this is in sharp contrast to the *sd*-preference model, where these properties *do* characterize the serial mechanism (see Bogomolnaia and Heo, 2012). In Section 4 we allow for the agents to be indifferent between objects. The natural extension of the serial mechanism to this setting, developed by Katta and Sethuraman (2006), is efficient and envy-free but not strategy-proof. We show that *no* mechanism can be efficient, envy-free and strategy-proof, even for the special case of unit q_j.

2. Strict preferences: strategy-proofness

Schulman and Vazirani show that the SG mechanism is efficient and envy-free, and is strategy-proof whenever $q_j \geq 1$ for all j. We start this section by showing that, in the absence of this last condition, *no* mechanism can satisfy all three properties.

3 In this paper we shall be concerned with envy-free mechanisms only. To be able to compare allocations across agents, we need the agent demands to be identical, and we normalize this demand to 1; this is really the rationale for insisting a priori that every agent has unit demand. The mechanisms described here admit natural extensions to settings in which the agents have different demands.

4 Lexicographic preferences have a long history in the economics literature, dating back at least to Hausner (1954), who formulated the theory of lexicographic expected utility. The key point of departure from the standard expected utility theory is that continuity is not required; thus, preferences over lotteries that do not satisfy continuity can be represented by a *vector*-valued utility function, and preferences over lotteries translate to lexicographic dominance of the corresponding utility vectors. We refer the reader to Cho (2012) for a more extensive discussion.

5 Earlier, Heo (2010) extended the probabilistic serial mechanism to the more general case of arbitrary object supplies and arbitrary agent demands (not necessarily 1), but focused on the *sd*-preference model. She showed that the serial rule (as she termed it) is efficient, and envy-free (if agents have unit demands).