Assignment with Multiple-Unit Demands and Responsive Preferences

Marek Pycia[*]

July 20, 2011

Abstract

This note introduces a general framework of assignment with multiple-unit demands and (first-order-stochastic-dominance-) responsive preferences, and provides some preliminary, illustrative, results.

[*]Department of Economics, UCLA, 8283 Bunche, Los Angeles, CA 90095; pycia (at) ucla.edu; http://pycia.bol.ucla.edu. I would like to thank Andrew Atkeson, Qingmin Liu, and William Zame. Keywords: multi-unit demands, responsive preferences, envy-freeness, ordinal efficiency, Random Priority, Probabilistic Serial.

1 Introduction

The present note introduces a general framework of assignment with multiple-unit demands and (first-order-stochastic-dominance-) responsive preferences, and offers some preliminary results on Bogomolnaia and Moulin (2001) Probabilistic Serial mechanism, and on envy-freeness. The results illustrate the usefulness of the proposed framework. We study assignment of both divisible and indivisible goods.

A related concept of responsive preferences has been well-studied in matching, beginning with Roth (1985), but has not been explored in studies of object allocation, except for the special case of responsiveness in which agents preferences are additively separable. Additively separable preferences have been studied by Hatfield (2009), Kojima, 2009, Kasajima (2009), and Heo (2011). In general the literature on multi-unit demands is rich, and the literature review is very preliminary in the current draft of the paper.

2 Model

A finite economy consists of a finite set of agents N, a finite set of object types Θ (or simply objects), and a finite set of object copies O.[1] Each copy $o \in O$ has a uniquely determined type $\theta(o) \in \Theta$.

Agents have multiple-unit demands. For each agent $i \in I$ let F_i be the set of feasible consumptions of agent i. We assume that each F_i is compact. We simultaneously study two variants of multiple-unit assignment:

Divisible goods: fractional consumption is allowed, $F_i \subset \mathbb{R}_+^{|\Theta|}$, and we assume that if a profile $(k_a)_{a\in\Theta}$ of quantities of goods $\theta \in \Theta$ is feasible then so is any profile $(k'_a)_{a\in\Theta}$ such that $0 \leq k'_a \leq k_a$, $a \in \Theta$.

Indivisible goods: only integer consumption is allowed, $F_i \subset \mathbb{N}^{|\Theta|}$, and we assume

[1] We consider both indivisible and divisible goods; in the divisible case we allow fractional copies. Instead of introducing O, we could equivalently talk about the total quantity of each object in the economy.

that if a profile $(k_a)_{a\in\Theta}$ of quantities of goods $\theta\in\Theta$ is feasible then so is any profile $(k'_a)_{a\in\Theta}$ such that $k'_a=0,1,...,k_a$, $a\in\Theta$.

Denote by $\mathcal{F}=\{F_i|i\in N\}$ the class of feasible structures in the economy. In the indivisible case, we denote by $\tilde{F}_i$ lotteries over feasible consumptions of agent i; in the divisible case, denote $\tilde{F}_i=F_i$. A *(random) allocation* μ specifies the expected quantities $\mu(i,a)\geq 0$ of good a assigned to agent i. In the divisible case randomization may be used, but is not needed; μ might entail randomization in the indivisible case.[2] All allocations studied in this paper are assumed to be feasible in that

$$\begin{aligned}\sum_{i\in N}\mu(i,a) &\leq \left|\theta^{-1}(a)\right| \quad \text{for every } a\in\Theta,\\ \mu(i,\cdot) &\in \tilde{F}_i \quad \text{for every } i\in N.\end{aligned}$$

The set of these random allocations is denoted by $\mathcal{M}$.

The setting includes, for instance, course allocation in which an agent can consume at most one unit of each object, and at most some fixed number of objects (see, among others, Sönmez and Ünver (2010); Budish and Cantillon (2010); Budish (2010)). Another special case of interest obtains when each agent $i\in N$ is endowed with capacity $K_i>0$ and an allocation $\mu(i,\cdot)$ is feasible if, and only if, $\sum_{a\in\Theta}\mu(i,a)\leq K_i$. In words, an allocation is feasible if i consumes K_i units, or less; we refer to this latter condition as *agent-specific capacities.* Feasibility structures satisfying agent-specific capacities have been examined in Hatfield (2009), Kojima, 2009, Kasajima (2009), and Heo (2011).

Each agent $i\in N$ has preferences $\succ_i^S$ over allocations in $\tilde{F}_i$. We assume that the preferences are responsive – in first order stochastic dominance sense – with respect

[2]When goods are indivisible an allocation needs to be implemented as a lottery over deterministic allocations; a deterministic allocation is a one-to-one mapping from agents to copies of objects from O. Since we cover both assignment of divisible and indivisible goods, we do not take a stance on whether the possibly random assignment can be decomposed. To assure decomposability one can additionally impose the decomposability conditions from Budish, Che, Kojima, and Milgrom (2011) (cf. also Hylland and Zeckhauser (1979) and Bogomolnaia and Moulin (2001)).

to a strict preference relation $\succ_i$ over objects from Θ. To define the responsiveness formally, let us say that $\mu(i,\cdot) \in \tilde{F}_i$ dominates $\mu'(i,\cdot) \in \tilde{F}_i$ (in first order stochastic dominance sense), or $\mu(i,\cdot) \geq^{\text{FOSD}} \mu'(i,\cdot)$ if

$$\sum_{a \succsim b} \mu(i,a) \geq \sum_{a \succsim b} \mu'(i,a) \quad \text{for all } b \in \Theta,$$

and the dominance is strict ($>^{\text{FOSD}}$) if one of the equalities is strict.

Preferences are *responsive* if

$$\begin{aligned} \mu(i,\cdot) \geq^{\text{FOSD}} \mu'(i,\cdot) \quad &\Rightarrow \quad \mu(i,\cdot) \succsim^S \mu'(i,\cdot), \\ \mu(i,\cdot) >^{\text{FOSD}} \mu'(i,\cdot) \quad &\Rightarrow \quad \mu(i,\cdot) \succ^S \mu'(i,\cdot). \end{aligned}$$

An example of such a responsive preference structure is when agents' demands are determined by additively separable von Neumann – Morgenstern utility functions; such additively separable environments have been studied by Hatfield (2009), Kojima, 2009, Kasajima (2009), and Heo (2011).[3]

Agents' preferences over objects define their preferences over copies of objects: agent i prefers object copy o over object copy o' iff she prefers $\theta(o)$ over $\theta(o')$, and the agent is indifferent between two object copies if they are of the same type. We can thus interchangeably talk about preferences over object types and preferences over object copies, or simply about preferences over objects. The indifference also implies that we can interchangeably talk about allocating objects and allocating copies of objects. We refer to the set of preference rankings as $\mathcal{P}$ and to the set of preference profiles as $\mathcal{P}^N$.

We assume that Θ contains the *null object* $\oslash$ ("outside option"), and we assume that it is not scarce: to simplify exposition let us assume that agents have bounded

[3]For studies that allow all possible preference profiles over allocations, not only responsive, see, for instance, Pápai (2001). Responsiveness is a natural structural assumption, and as we illustrate below, it implies that Bogomolnaia and Moulin (2001) Probabilistic Serial is well-behaved in multi-unit demand settings.

demands, and that $|\theta^{-1}(\oslash)|$ is so large that even if all agents are at capacity, $\oslash$ is still available. An object is called acceptable if it is preferred to $\oslash$.

Agent's i feasibility structure F_i and his preference ranking $\succ_i$ over objects are referred to as the *type* of the agent.

A *mechanism* $\phi : \mathcal{P}^N \to \mathcal{M}$ is a mapping from the set of profiles of preferences over objects that agents report to the set of allocations.

3 Ordinal Efficiency, Envy-Freeness, and Weak Strategy-Proofness

In this section we simultaneously characterize the celebrated Probabilistic Serial mechanism of Bogomolnaia and Moulin (2001), and two natural properties of allocations: ordinal efficiency and envy-freeness. Given preference profile $\succ_N$, a random allocation μ *ordinally dominates* another random allocation μ' if for every agent i the distribution $\mu(i,\cdot)$ first order stochastically dominates $\mu'(i,\cdot)$, that is

$$\sum_{b \succsim_i a} \mu(i,b) \geq \sum_{b \succsim_i a} \mu'(i,b), \qquad \forall a \in \Theta.$$

A random allocation is *ordinally efficient* with respect to a preference profile $\succ_N$ if it is not ordinally dominated by any other feasible allocation. Ordinal efficiency is a weak and natural efficiency requirement: if an allocation is not ordinarily efficient, then all agents would ex ante agree there is a better one. Bogomolnaia and Moulin (2001) discuss this requirement in depth. The conditional form of ordinal efficiency has been introduced by Budish, Che, Kojima, and Milgrom (2011).

Given preference profile $\succ_N$, an allocation μ is *envy-free* if for any two agents $i, j \in N$ and any allocation $\hat{\mu}(i,\cdot) \in \tilde{F}_i$ such that $0 \leq \hat{\mu}(i,b) \leq \mu(j,b)$, for all $b \in \Theta$, agent i first-order stochastically prefers his allocation $\mu(i,\cdot)$ over $\hat{\mu}(i,\cdot)$, that is

$$\sum_{b \succsim_i a} \mu(i,b) \geq \sum_{b \succsim_i a} \hat{\mu}(i,b), \qquad \forall a \in \Theta.$$

When there is only one feasibility structure, $\mathcal{F}$, this concept is equivalent to the standard concept of first order stochastic dominance envy-freeness. The first order stochastic dominance comparison is well-funded in our context because agents' preferences over allocations are responsive. Envy-freeness (referred to also as no envy) is a strong fairness requirement introduced by Foley (1967).

A mechanism ϕ is weakly strategy-proof if $\phi(\succ_i', \succ_{-i})(i) \succsim_i^S \phi(\succ)(i)$ implies that $\phi(\succ_i', \succ_{-i})(i) = \phi(\succ)(i)$. This concept has been studied by Bogomolnaia and Moulin (2001).

4 Probabilistic Serial

Probabilistic Serial treats copies of an object type as a pool of probability shares of the object type. Given preference profile $\succ_N$, the random allocation produced by Probabilistic Serial can be determined through an "eating" procedure in which each agent "eats" probability share of the best acceptable and available object with speed 1 at every time $t \in [0,1]$; an object a is available at time t if its initial endowment $\theta^{-1}(a)$ is larger than the sum of shares that have been eaten by time t.

Formally, at time $t = 0$, the total quantity of available shares of object type $a \in \Theta$ is $Q_a(0) = |\theta^{-1}(a)|$, and for times $t \in [0,\infty)$ we define the set of available objects $A(t) \subseteq \Theta$ and the available quantity $Q_a(t)$ of probability shares of object $a \in \Theta$ through the following system of integral equations. To formulate the equations we say that *agent i eats from object a at time t* iff $a \in A(t)$ and $\forall b \in A(t)\ a \succsim_i b$, and there exists $\epsilon > 0$ such that $(\psi^t(i,b)|_{b \neq a}, \psi^t(i,a+\epsilon)) \in \tilde{F}_i$. The system of integral equations is:

$$\begin{aligned} A(t) &= \{a \in \Theta | \, Q_a(t) > 0\}, \\ Q_a(t) &= Q_a(0) - \int_0^t |\{i \in N | \, i \text{ eats from } a \text{ at time } \tau\}| \, d\tau, \\ \psi^t(i,a) &= \int_0^t \chi(i \text{ eats from } a \text{ at time } \tau) \, d\tau, \end{aligned}$$

where the Boolean function χ (statement) takes value 1 if the statement is true and 0 otherwise. If stopped at time t, the eating procedure allocates object $a \in \Theta$ to agent $i \in N$ with probability $\psi^t(i,a)$. The allocation $\psi(i,a)$ of *Probabilistic Serial* is given by the eating procedure stopped at the time no agent eats from any object any longer, alternatively $\psi = \psi^{|O|}$.

The continuity of the functions Q_a implies that for any time $T \in [0,1)$ and any $\eta > 0$ sufficiently small, any agent i eats the same object for all $t \in [T, T+\eta)$. In the eating procedure there are some critical times when one or more objects get exhausted. At this time some of the available quantity functions Q_a have kinks; at other times their slope is constant.[4]

5 Some Observations on Probabilistic Serial

To illustrate the usefulness of our setting, this section extends some of the key insights from Bogomolnaia and Moulin (2001) to the present multi-unit demand setting.[5]

Proposition 1. *The allocations of Probabilistic Serial are envy-free and ordinally efficient. If the feasibility structure satisfies agent-specific capacities then Probabilistic Serial is weakly strategy-proof.*

[4]This structure of quantity functions Q_a implies that we can define the allocation of Probabilistic Serial through a system of difference equations; such definitions are given in Bogomolnaia and Moulin (2001), and, for the environment with copies, in Kojima and Manea (2010). Heo (2011) extends the definition of probabilistic serial to multiple-demand environments in which demands are determined by additive utility function.

[5]Cf. also Heo, 2011 who extended the above result to multi-unit demand setting with additively separable preferences.

Ordinal-efficiency and envy-freeness of Probabilistic Serial were proved by Bogomolnaia and Moulin (2001) for single-unit demands, and their proof for envy-freeness directly extends to our case (notice that envy-freeness is satisfied by each ψ^t).

The proof of ordinal efficiency extends directly if we additionally assume that the feasibility structure satisfies agent-specific capacities. Without this additional assumption the proof requires slight modifications. First, notice that after the eating procedure stops, either the feasibility constraint does not allow agent i to have more of an object $a \in \Theta$, or a is no longer available, or a is unacceptable for i. Given this it is enough to notice that at each time t, the allocation ψ^t is undominated by any feasible allocation that uses only copies eaten till time t. The proof of this last claim relies on the observation that the eating time is naturally divided into intervals such that during the interval, no agent changes the object eaten (or stops eating). We may conduct the argument by induction on these intervals. The claim is true in the first interval, as each agent eats from their most preferred object (among objects the agent can receive in positive quantity). The responsive structure of preferences now implies that if the claim is true until n-th interval, then it cannot be violated in the n-th interval.

Bogomolnaia and Moulin (2001) also proved that Probabilistic Serial is weakly strategy-proof, and this insight also extends to our setting, with no substantive change in its proof, provided the feasibility structure satisfies agent-specific capacities.

6 On Envy-Freeness and ϵ-Envy-Freeness

This section extends some of insights from Pycia (2011a). Given preference profile $\succ_N$, an allocation μ is *envy-free* if for any two agents $i, j \in N$ and any allocation $\hat{\mu}(i, \cdot) \in \tilde{F}_i$ such that $0 \le \hat{\mu}(i,b) \le \mu(j,b)$, for all $b \in \Theta$, agent i first-order stochastically prefers his allocation $\mu(i,\cdot)$ over $\hat{\mu}(i,\cdot)$, that is

$$\sum_{b \succsim_i a} \mu(i,b) \geq \sum_{b \succsim_i a} \hat{\mu}(i,b), \qquad \forall a \in \Theta.$$

When $|\mathcal{F}| = 1$, this general definition reduces to the standard definition: given preference profile $\succ_N$, an allocation μ is envy-free if any agent i first-order stochastically prefers his allocation over the allocation of any other agent j, that is

$$\sum_{b \succsim_i a} \mu(i,b) \geq \sum_{b \succsim_i a} \mu(j,b), \qquad \forall a \in \Theta.$$

The standard form of the condition also obtain if $F_i = F_j$.[6]

A preference profile has *full support* if each agent type is represented in the profile. The restriction to full-support preference profiles is strong in small markets, however as the market becomes large the restriction becomes mild.

Proposition 2. *If $\succ_N$ has full support and an allocation μ is envy-free and ordinally efficient then there are constants $t_{a,F} > 0$, $a \in \Theta$, $F \in \mathcal{F}$, such that $\mu(i,a) \neq 0$ and $F_i = F$ imply $t_{a,F} = \sum_{b \succsim_i a} \mu(i,b)$. If, additionally, the feasibility structure satisfies agent-specific capacities, then there are constants $t_a > 0$ such that $t_{a,F} = \min(K_i, t_a)$.*

Proof. Let $\mu(i,a) \neq 0$ and $\mu(j,a) \neq 0$, and $F_i = F_j$. Take agent i' such that $F_{i'} = F_i$ and who ranks a first and otherwise ranks objects as i does; no envy implies $\mu(i',a) = \sum_{b \succsim_i a} \mu(i,b)$. Similarly, for agent j' such that $F_{i'} = F_i$ and who ranks a first and otherwise ranks objects as j does; $\mu(j',a) = \sum_{b \succsim_i a} \mu(j,b)$. Finally, no envy further implies that $\mu(i',a) = \mu(j',a)$ because $F_{i'} = F_i = F_j = F_{j'}$. If agent specific capacities are satisfied, then this argument goes through without the assumption $F_i = F_j$ as long as the capacity constraint is not binding. $\square$

Given $\epsilon > 0$ and preference profile $\succ_N$, an allocation μ is *ϵ-envy-free* if for any

[6] Envy-freeness (referred to also as no envy) is a strong fairness requirement introduced by Foley (1967), see Bogomolnaia and Moulin (2001) for the $|\mathcal{F}| = 1$ formulation.

two agents $i, j \in N$ and any allocation $\hat{\mu}(i, \cdot) \in \tilde{F}_i$ such that $0 \leq \hat{\mu}(i, b) \leq \mu(j, b)$, for all $b \in \Theta$,

$$\sum_{b \succsim_i a} \mu(i, b) \geq \sum_{b \succsim_i a} \hat{\mu}(i, b) - \epsilon, \qquad \forall a \in \Theta.$$

When $F_i = F_j$ this condition simply means that

$$\sum_{b \succsim_i a} \mu(i, b) \geq \sum_{b \succsim_i a} \mu(j, b) - \epsilon, \qquad \forall a \in \Theta.$$

Proposition 3. *If $\succ_N$ has full support, $\epsilon > 0$, and an allocation μ is ϵ-envy free and ordinally efficient, then there are constants $t_a(F) > \frac{5}{2}\epsilon$, $a \in \Theta$, $F \in \mathcal{F}$, such that $\mu(i, a) > \epsilon$ and $F = F_i$ imply $\left| t_{a,F} - \sum_{b \succsim_i a} \mu(i, b) \right| \leq \frac{3}{2}\epsilon$. If, additionally, the feasibility structure satisfies agent-specific capacities, then there are constants $t_a > 0$ such that $t_{a,F} = \min(K_i, t_a)$.*

Proof. As in the proof of Proposition 1, ϵ-envy-freeness implies that if $\mu(i, a), \mu(j, a) > \epsilon$ and $F_i = F_j$ then $\left| \sum_{b \succsim_i a} \mu(i, b) - \sum_{b \succsim_i a} \mu(i, b) \right| \leq 3\epsilon$. Thus there exists t_a such that $\sum_{b \succsim_i a} \mu(i, b)$ is within $\frac{3}{2}\epsilon$ of t_a for all agents obtaining a with probability ϵ or higher. We can assume $t_{a_1} > \frac{5}{2}\epsilon$. If agent specific capacities are satisfied, then this argument goes through without the assumption $F_i = F_j$ as long as the capacity constraint is not binding. □

7 A Comment on Random Priority

To allocate objects, Random Priority first draws an ordering of agents from a uniform distribution over orderings, and then allocates the first agent her most preferred feasible allocation, then allocates the second agent his most preferred feasible allocation (drawing on objects that that still has unallocated copies), etc. (see Abdulkadiroğlu and Sönmez, 1998 for a seminal study of Random Priority). In our setting Random Priority remains ex post efficient, treats agents with identical feasibility constraints symmetrically, and is strategy-proof.

8 Conclusion

The present note introduces the general framework of assignment with multiple-unit demands and responsive preferences. This framework is used in Pycia (2011b) to extend the single-unit demand results of Liu and Pycia (2011) to multi-unit demand environments.

References

Abdulkadiroğlu, A. and T. Sönmez (1998): "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems," *Econometrica*, 66, 689–701.

Bogomolnaia, A. and H. Moulin (2001): "A New Solution to the Random Assignment Problem," *Journal of Economic Theory*, 100, 295–328.

Budish, E. (2010): "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes," Working Paper.

Budish, E. and E. Cantillon (2010): "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *American Economic Review*, forthcoming.

Budish, E., Y.-K. Che, F. Kojima, and P. Milgrom (2011): "Implementing Random Assignments: A Generalization of the Birkhoff - von Neumann Theorem," University of Chicago, Columbia University, Stanford University, unpublished mimeo.

Foley, D. K. (1967): "Resource Allocation and the Public Sector," *Yale Economic Essays*, 7, 45–98.

Hatfield, J. W. (2009): "Strategy-Proof, Efficient, and Nonbossy Quota Allocations," *Social Choice and Welfare*, 33 No. 3, 505–515.

Heo, E. J. (2011): “Probabilistic Assignment with Quotas: A Generalization and a Characterization of the Serial Rule,” .

Hylland, A. and R. Zeckhauser (1979): “The Efficient Allocation of Individuals to Positions,” *Journal of Political Economy*, 87, 293–314.

Kasajima, Y. (2009): “More on the probabilistic assignment of indivisible goods when agents receive several,” .

Kojima, F. (2009): “Random Assignment of Multiple Indivisible Objects,” *Mathematical Social Sciences*, 57, 134–142.

Kojima, F. and M. Manea (2010): “Incentives in the Probabilistic Serial Mechanism,” *Journal of Economic Theory*, 144, 106–123.

Liu, Q. and M. Pycia (2011): “Ordinal Efficiency, Fairness, and Incentives in Large Markets,” .

Pápai, S. (2001): “Strategyproof and Nonbossy Multiple Assignments,” *Journal of Public Economic Theory*, 3, 257–271.

Pycia, M. (2011a): “On Envy-Freeness,” .

——— (2011b): “Ordinal Efficiency, Fairness, and Incentives in Large Multi-Unit-Demand Assignments,” .

Roth, A. E. (1985): “The college admission problem is not equivalent to the marriage problem,” *Journal of Economic Theory*, 36, 277–288.

Sönmez, T. and M. U. Ünver (2010): “Course Bidding at Business Schools,” *International Economic Review*, 51, 99–123.