Ordinal Efficiency, Fairness, and Incentives in Large Multi-Unit-Demand Assignments

Marek Pycia*

July 20, 2011

Abstract

This note extends the single-unit-demand results of Liu and Pycia (2011) to the responsive-preferences and multiple-unit-demand framework of Pycia (2011a).

*Department of Economics, UCLA, 8283 Bunche, Los Angeles, CA 90095; pycia (at) ucla.edu; http://pycia.bol.ucla.edu

1 Introduction

In Liu and Pycia (2011), we studied the canonical model of assignment without transfers, with single-unit demands, and showed that in large markets all sensible, symmetric, and asymptotically strategy-proof ordinal allocation mechanisms coincide asymptotically, and that ordinal efficiency is obtained in the limit. The equivalence was built on a surprising finite-market result: the ordinally-efficient and envy-free allocation is unique and coincides with the outcome of Probabilistic Serial, provided the agents' preference profile has full support. We also gave an easy-to-verify condition for asymptotic ordinal efficiency. In the present note we extend the single-unit demand results to the responsive-preferences and multiple-unit-demand framework of Pycia (2011a); the sequencing of results is the same as in our main paper on the subject, Liu and Pycia (2011).[1]

We study assignment of both divisible and indivisible goods.

2 Model

We study the responsive-preferences model introduced by Pycia (2011a). 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.[2] Each copy $o \in O$ has a uniquely determined type $\theta(o) \in \Theta$. To avoid trivialities, we assume that each object is represented by at least one copy.

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

[1]Please see Liu and Pycia (2011) for a detailed discussion of the background of the results. In this our earlier paper we remarked in Conclusion that our results remain valid in a multiple-unit demand setting; the present note provides the (straightforward) details behind this assertion.

[2]We consider both indivisible and divisible goods; in the divisible case we allow fractional copies.

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 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.[3] 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 $\mathcal{M}$.

In some results we impose a further assumption – which we call *agent-specific capacities* – that 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 consume K_i units, or less.[4] Feasibility structures satisfying agent-specific capacities have been examined in Hatfield (2009), Kojima, 2009, Kasajima (2009), and Heo (2011b).

[3] 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)).

[4] Theorem 3 and half of Theorem 1, as well as Remarks 1,2 and 3 do not rely on this assumption, but Theorems 2 and 4, and half of Theorem 1 rely on it. The agent-specific-capacities assumption fails in many settings of interest, such as course allocation; see, among others, Sönmez and Ünver (2010); Budish and Cantillon (2010); Budish (2010).

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 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

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

A related concept of responsive preferences has been well-studied in matching (c.f. Roth, 1985), but has not been explored in studies of object allocation. 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 (2011b).

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 (an agent's types) 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 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 A Characterization of Ordinal Efficiency and Envy-Freeness

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.[5] 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.

Given preference profile $\succ_N$, an allocation μ is *envy-free* if for any two agents

[5]The conditional form of ordinal efficiency we use has been introduced by Budish, Che, Kojima, and Milgrom (2011). Probabilistic Serial and the definition of envy-freeness was extended to our setting by Pycia (2011a).

$i, j \in N$ 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).

3.1 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:

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

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.[6]

3.2 Main Finite-Market Results

Our goal is to show that ordinal efficiency and envy-freeness fully characterize the allocation of Probabilistic Serial. To do so we restrict attention to preference profiles with full support. A preference profile has *full support* if each agent type $(\succ, F_i) \in \mathcal{P} \times \mathcal{F}$ is represented in the profile.[7] The restriction to full-support preference profiles is strong in small markets, however as the market becomes large the restriction becomes

[6]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 (2011b) extends the definition of probabilistic serial to multiple-demand environments in which demands are determined by additive utility function.

[7]For our results, it is enough to impose this assumption for types that rank all objects as acceptable.

mild.

Theorem 1. *If the preference profile has full support and the feasibility structure satisfies agent-specific capacities, then an allocation is ordinally efficient and envy-free if and only if it is generated by Probabilistic Serial.*

Ordinal-efficiency and envy-freeness of Probabilistic Serial were first proved by Bogomolnaia and Moulin (2001), and extended to our setting by Pycia (2011a).[8] The converse implication has been proved by Liu and Pycia (2011) for single-unit demands, and their proof directly extends to our case.[9] We provide another proof, based on the ideas in Pycia (2011a,b).

Proof. Let t_a be the constants from Proposition 2 in Pycia (2011a). Rename the objects so that

$$0 < t_{a_1} \leq t_{a_2} \leq ... \leq t_{a_{|\Theta|}}.$$

Our assumptions uniquely determine the $0-1$ matrix indexed by i and a where a cell equals 1 iff $\mu(i,a) \neq 0$. Indeed, if i ranks a_1 first then no-envy implies that $\mu(i,a_1) = \min(t_{a_1}, K_i) \neq 0$; if i ranks a second or lower, then ordinal efficiency and full support imply that $\mu(i,a_1) = 0$. The rest of the matrix is determined similarly, by induction.

Define $A_1(k) = \{i | \, i \text{ ranks } a_1 \text{first and } \mu(i,a_1) \neq 0, \text{ and } k = K_i\}$. Given the above matrix, ordinal efficiency implies that $t_{a_1} = \max_{k \in \mathcal{F}} k$ or is implicitly given by

$$\theta^{-1}(a_1) = \sum_{k \in \mathcal{F}} |A_1(k)| \min(t_{a_1}, k).$$

Proceeding in this way, by induction, we pin down all values t_a, thus uniquely determining the allocation μ. By Proposition 1 in Pycia (2011a), the allocation thus

[8] As shown in Pycia (2011a), this implication does not rely on the assumption of agent-specific capacities.

[9] We just need to replace time 1 by time $K = \max_{i \in N} K_i$, and reinterpret probability as quantity. Heo (2011a) elegantly simplified Liu and Pycia (2011)'s proofs of this result and of related Theorem 2; her simplification remains valid in the multi-unit demand environment.

coincides with Probabilistic Serial. □

The latter implication relies on the preference profile having full support; there are non-full-support preference profiles for which the converse implication fails — see Bogomolnaia and Moulin (2001) (c.f. also Example 2 from Kesten, Kurino, and Ünver (2011)).

4 Allocations in Large Markets

The characterization of efficient and fair allocations given by Theorem 1 holds true in any finite market, including large markets. Ordinal efficiency is a natural requirement, and envy-freeness is an attractive property of allocations; however these requirements are very strong – there are many sensible allocation mechanisms that do not satisfy them. The goal of this section is to show that much weaker requirements – asymptotic ordinal efficiency and asymptotic envy-freeness – are sufficient to determine the allocation as the market becomes large.

To achieve this goal let us fix a sequence of finite economies $\langle N_q, \Theta, O_q \rangle_{q=1,2,\dots}$ in which the set of object types, Θ, and the set of feasibility structures $\mathcal{F}$ are fixed while the set of agents N_q grows in q; we will assume throughout that $|N_q| \to \infty$ as $q \to \infty$. We also assume that the proportion of agents endowed with each feasibility structure $F_i \in \mathcal{F}$ is above some lower bound $\eta > 0$ that does not depend on q. This assumption is trivially satisfied if $\mathcal{F}$ contains only one feasibility structure.

To avoid repetition, in the sequel we refer to $\langle N_q, \Theta, O_q \rangle$ as the q-economy, and maintain a notational assumption that allocations μ_q and mechanisms ϕ_q are defined on q-economies. The q-economy function mapping object copies to their types is denoted θ_q. The set of random allocations in the q-economy is denoted $\mathcal{M}_q$.

Notice that we do not impose any assumptions on the sequence of sets of object copies, O_q, except for some remarks where we explicitly impose the additional assumption that $\left|\theta_q^{-1}(a)\right| \to \infty$. Our main results apply equally well regardless of whether the number of object copies stays bounded, or whether it grows slower than, faster than, or at the same rate as the number of agents in the economy. In particular, replica economies in which the number of agents and the number of object copies grow at the same rate are a special case of our setting, as is the environment studied by Che and Kojima (2010) who assume that the ratio $\left|\theta_q^{-1}(a)\right| / |N_q|$ converges to a positive limit for all non-null objects $a \in \Theta$.

4.1 Asymptotic Ordinal Efficiency

Liu and Pycia (2011) introduced the following concept of asymptotic ordinal efficiency, based on an auxiliary concept of ϵ-ordinal efficiency. For simplicity of exposition let us define it only for environments that satisfy agent-specific capacities.[10]

Given an $\epsilon > 0$, we say that a random allocation μ is *ϵ-ordinally efficient* with respect to a preference profile $\succ$ iff (i) no agent is allocated a higher-than-ϵ probability of an unacceptable object, (ii) if object a is unallocated with probability higher than ϵ and $\mu(i, b) > \epsilon$, then $b \succ_i a$, and (iii) there is no cycle of agents $i_0, i_1, ..., i_n$ and objects $a_0, a_1, ..., a_n$ such that $\mu(i_k, a_k) > \epsilon$ and $a_{k+1} \succ_{i_k} a_k$ (all subscripts modulo $n+1$).

Given a sequence of preference profiles $\succ_{N_q}$, a sequence of allocations μ_q is *asymptotically ordinally efficient* if for each $q = 1, 2, ...$ there is $\epsilon(q) > 0$ such that $\epsilon(q) \to 0$ when $q \to \infty$ and μ_q is $\epsilon(q)$-ordinally efficient with respect to $\succ_{N^q}$. We say that the asymptotic ordinal efficiency obtains uniformly on a class of sequences of allocations if $\epsilon(q) \to 0$ uniformly on this class.

This definition of asymptotic ordinal efficiency is motivated by the following result from Che and Kojima (2010) (see also Bogomolnaia and Moulin (2001)) — an allocation μ is ordinally efficient iff the following exact analogues of conditions (i)-(iii) hold true: (i') no agent is allocated a positive probability of an unacceptable object,

[10] The extension to our general environment is straightforward but it obscures the thrust of the concept since we need to be more careful in the treatment of feasibility.

(ii') if object a is unallocated with positive probability and $\mu(i,b) > 0$, then $b \succ_i a$, and (iii') there is no cycle of agents $i_0, i_1, ..., i_n$ and objects $a_0, a_1, ..., a_n$ such that $\mu(i_k, a_k) > 0$ and $a_{k+1} \succ_{i_k} a_k$.[11]

4.2 Asymptotic Envy-Freeness

To formulate our main result on allocations in large market, we need to relax envy-freeness to asymptotic envy-freeness. Fix a sequence of preference profiles $\succ_{N_q}$. A sequence of random allocations μ_q is *asymptotically envy-free* if

$$\liminf_q \min_{i,j\in N_q,\, 0\le\hat{\mu}_q(i,\cdot)\le\mu_q(j,\cdot),\, \hat{\mu}_q(i,\cdot)\in\tilde{F}_i,\, a\in\Theta} \left[\sum_{b\succsim_i a} \mu_q(i,b) - \sum_{b\succsim_i a} \hat{\mu}_q(i,b)\right] \ge 0.$$

We say that the asymptotic envy-freeness of allocations obtains uniformly on a class of sequences of allocations if the lim inf convergence obtains uniformly on this class. Of course, any sequence of envy-free allocations is asymptotically envy-free.[12]

4.3 Asymptotic Full Support

We derive our first asymptotic results for sequences of preference profiles that have full support in the limit; in Section 5 we relax this assumption.[13] Following Liu and Pycia (2011), we say that a sequence of preference-profiles $\succ_{N_q}$ has *asymptotically full support* if there exists $\delta > 0$ and $\bar{q}$ such that for any $q > \bar{q}$, and for any agent type $(\succ, F_i) \in \mathcal{P} \times \mathcal{F}$, the proportion of agents whose $\succ_{N_q}$-ranking agrees with $(\succ, F_i)$ is

[11]Condition (i') is known as individual rationality, and condition (ii') as non-wastefulness. Analogues of all of our results remain true if we strengthen the concept of asymptotic ordinal efficiency by substituting the more demanding conditions (i') and (ii') for conditions (i) and (ii). No change in the results and proofs is needed, except for Theorem 4 when we need to impose some additional assumption such as ex post Pareto efficiency (defined in Section 6), or directly conditions (i') and (ii'). All mechanisms we explicitly discuss in this paper, including those listed in Remark 3, satisfy conditions (i') and (ii').

[12]Asymptotic envy-freeness was studied by Jackson and Kremer (2007).

[13]The restriction to asymptotic full-support preference profiles is not needed if the allocations are generated by mechanisms satisfying a mild asymptotic continuity assumption.

above δ. Asymptotic full support holds true uniformly on a class of preference profiles if they have asymptotic full support with the same δ and $\bar{q}$.

Asymptotic full support means that, as q grows, any preference ranking is represented by a non-vanishing fraction of agents.[14] Because a full-support profile can have a single agent of any given type, there are sequences of full-support profiles which are not asymptotically full-support. However, full-support sequences of preference profiles are asymptotically generic; Liu and Pycia (2011) formally defined asymptotic genericity and proved that asymptotically full-support preference profiles are asymptotically generic.

4.4 Main Result on Allocations in Large Markets

The above concepts allow us to state our main result on allocations in large markets:

Theorem 2. *Assume the feasibility structure satisfies agent-specific capacities, and fix a sequence of preference profiles $\succ_{N_q}$ with asymptotically full-support. If two sequences of allocations μ_q and μ'_q are each asymptotically ordinally efficient and asymptotically envy-free then they asymptotically coincide, that is,*

$$\max_{i\in N_q, a\in\Theta} \left|\mu_q(i,a) - \mu'_q(i,a)\right| \to 0 \quad \textit{as} \quad q\to\infty.$$

In the sequel we rely on a slightly stronger version of this result:

Theorem 2. (Uniform Version) *Assume the feasibility structure satisfies agent-specific capacities. If a class $\mathcal{Q}$ of preference profile sequences has uniformly asymptotic full support, and two classes of allocation sequences*

$$\left\{\phi_q\left(\succ_{N_q}\right) \mid \left(\succ_{N_q}\right)_{q=1,2,\ldots} \in \mathcal{Q}\right\} \quad \textit{and} \quad \left\{\phi'_q\left(\succ_{N_q}\right) \mid \left(\succ_{N_q}\right)_{q=1,2,\ldots} \in \mathcal{Q}\right\},$$

[14] In a continuum economy, the counterpart of asymptotic full support says that every ordering is represented with positive probability; in other words the distribution of orderings has full support. Our results on asymptotically full-support profiles remain valid if the assumption of non-vanishing representation is imposed only for ranking of objects $\succ$ in which all non-null objects are acceptable.

are each uniformly asymptotic ordinally efficient and asymptotic envy-free, then the asymptotic convergence of the allocation sequences is uniform, that is,

$$\max_{\left(\succ_{N_q}\right)_{q=1,2,\ldots}\in\mathcal{Q},\, i\in N_q,\, a\in\Theta} \left|\phi_q\left(\succ_{N_q}\right)(i,a)-\phi_q'\left(\succ_{N_q}\right)(i,a)\right| \to 0 \quad as \quad q\to\infty.$$

The proof below extends the argument from Pycia (2011b) (the original argument in Liu and Pycia, 2011 can also be straightforwardly extended).

Proof. Let ϵ be small relative to the lowest capacity $\min_{k\in\mathcal{F}} k$. Fix an economy far enough in the sequence so that it is ϵ-envy-free and ϵ-ordinally-efficient. Let t_a be the constants defined in Proposition 3 of Pycia (2011a). We can rename the objects so that

$$\frac{5}{2}\epsilon < t_{a_1} \le t_{a_2} \le \ldots \le t_{a_{|\Theta|}}.$$

Our assumptions uniquely determine the $0-1$ matrix indexed by i and a when a cell is 1 iff $\mu(i,a) > \epsilon$. Indeed, if i ranks a_1 first then ϵ-envy-freeness implies that$\mu(i,a_1) \ge \min(k_i, t_{a_1}) - \frac{3}{2}\epsilon > \epsilon$; if i ranks a second or lower, then ϵ-ordinal efficiency and full support imply that $\mu(i,a_1) \le \epsilon$. The rest of the matrix is determined similarly, by induction.

Denote $A_1(k) = \{i|\, i\text{ranks } a_1\text{first and } \mu(i,a_1) > \epsilon \text{ and } K_i = k\}$. Given the above matrix, ϵ-ordinal efficiency implies that either t_{a_1} is ϵ-close to 1 (and the lemma is true), or

$$\sum_{k\in\mathcal{F}} |A_1(k)|\left(t_{a_1} - \frac{3}{2}\epsilon\right) \le \theta^{-1}(a_1) \le \sum_{k\in\mathcal{F}}\left\{|A_1(k)|\left(t_{a_1}+\frac{3}{2}\epsilon\right) + \left(|N_q| - |A_1(k)|\right)\epsilon\right\} + \epsilon.$$

In the latter case,

$$\begin{aligned} t_{a_1} &\in \left(\frac{\theta^{-1}(a_1)}{\sum_{k\in\mathcal{F}}|A_1(k)|} - \frac{5}{2}\epsilon - \frac{|N_q| - \sum_{k\in\mathcal{F}}|A_1(k)|}{\sum_{k\in\mathcal{F}}|A_1(k)|}\epsilon,\ \frac{\theta^{-1}(a_1)}{\sum_{k\in\mathcal{F}}|A_1(k)|} + \frac{3}{2}\epsilon\right) \\ &\subseteq \left(\frac{\theta^{-1}(a_1)}{\sum_{k\in\mathcal{F}}|A_1(k)|} - \frac{5}{2}\epsilon - \frac{1}{\eta}\epsilon,\ \frac{\theta^{-1}(a_1)}{\sum_{k\in\mathcal{F}}|A_1(k)|} + \frac{3}{2}\epsilon\right) \end{aligned}$$

Proceeding similarly through (at most) $|\Theta|$ steps of induction, we pin down all values t_{a_ℓ}, and hence pin down all $\mu(i,a)$ up to some $M(\eta,|\Theta|)\epsilon$. The bound is uniform if all assumptions are imposed uniformly. QED

Using a mild asymptotic continuity assumption, in Section 5 we relax the restriction to asymptotically full-support preference profiles. Note also that an analogue of Theorem 2 holds true in environments in which all objects are acceptable.

5 Allocation Mechanisms in Large Markets

In this section we move beyond studying allocations for single preference profiles (and subsets of profiles), and study mechanisms $\phi_q : \mathcal{P}^{N_q} \to \mathcal{M}_q$. We first show that asymptotic envy-freeness is a mild requirement, and then derive analogues of our results for all preference profiles, rather than only asymptotically full-support profiles.

5.1 Symmetric and Asymptotically Strategy-Proof Mechanisms

How strong an assumption is asymptotic envy-freeness? Liu and Pycia (2011) showed that it is surprisingly mild in single-unit demand setting, and we extend this insight to multiple-unit demands.

Symmetry is a basic fairness property of an allocation, and is also known as equal treatment of equals. When there is only one feasibility structure, $|\mathcal{F}| = 1$, we say that a random allocation μ is symmetric, given a preference profile $\succ_N$, if any two agents i and j who submitted the same ranking of objects, $\succ_i = \succ_j$, are allocated the same distributions over objects, $\mu(i,\cdot) = \mu(j,\cdot)$. This motivates our definition in the general case, when there are potentially many feasibility structures. Given a preference profile $\succ_N$, random allocation μ is *symmetric*, if whenever agent i submitted the same ranking of objects as agent j, and they both rank objects $a_1 \succ_{i,j} a_2 \succ_{i,j} ... \succ_{i,j} a_{|\Theta|}$, then i is allocated the distribution $\mu(i,\cdot)$ such that

$|\mu(i,a_1)-\mu(j,a_1)|$ is minimized among $\mu(i,\cdot)\in\tilde{F}_i$, among distributions satisfying this constraint $|\mu(i,a_2)-\mu(j,a_2)|$ is minimized, etc.

Our results will in fact rely only on an asymptotic form of this assumption. Given a preference profile $\succ_N$ an allocation is ϵ-symmetric if whenever agent i submitted the same ranking of objects as agent j, and they both rank objects $a_1 \succ_{i,j} a_2 \succ_{i,j} \ldots \succ_{i,j} a_{|\Theta|}$, then i is allocated the distribution $\mu(i,\cdot)$ such that $|\mu(i,a_1)-\mu(j,a_1)|$ is within ϵ of the minimum among $\mu(i,\cdot)\in\tilde{F}_i$, among distributions satisfying this constraint $|\mu(i,a_2)-\mu(j,a_2)|$ is within ϵ of the minimum, etc. Given a sequence of preference profiles $\succ_{N_q}$, a sequence of allocations is *asymptotically symmetric* if for each $q=1,2,\ldots$ there is $\epsilon(q)>0$ such that $\epsilon(q)\to 0$ when $q\to\infty$ and μ_q is $\epsilon(q)$-symmetric with respect to $\succ_{N^q}$. Asymptotic symmetry obtains uniformly on a class of sequences of allocations if the convergence is uniform on this class. Of course, every sequence of symmetric allocations is asymptotically symmetric.

Before defining asymptotic strategy-proofness, let us review the standard definition of strategy-proofness of random ordinal mechanism (cf. Gibbard 1977). A random mechanism ϕ is *strategy-proof* if for any agent $i\in N$ and any profile of preferences $\succ_{N-\{i\}}$ submitted by other agents, the allocation agent i obtains by reporting the truth, $\phi\left(\succ_i,\succ_{N-\{i\}}\right)(i,\cdot)$, first-order stochastically dominates allocation the agent can get by reporting anther preference ranking $\succ_i'$, that is

$$\sum_{b\succsim_i a}\phi\left(\succ_i,\succ_{N-\{i\}}\right)(i,b)\geq\sum_{b\succsim_i a}\phi\left(\succ_i',\succ_{N-\{i\}}\right)(i,b),\qquad\forall a\in\Theta.$$

A sequence of random mechanisms ϕ_q is *asymptotically strategy-proof on a sequence of preference profiles* $\succ_{N_q}$ if

$$\liminf_q\min_{i\in N_q,\succ_i'\in\mathcal{P},a\in\Theta}\left[\sum_{b\succsim_i a}\phi_q\left(\succ_{N_q}\right)(i,b)-\sum_{b\succsim_i a}\phi_q\left(\succ_i',\succ_{N_q-\{i\}}\right)(i,b)\right]\geq 0.$$

We say that asymptotic strategy-proofness obtains uniformly on a class of sequences of preference profiles if the lim inf convergence obtains uniformly on this class. A sequence of mechanisms is *asymptotically strategy-proof* if the convergence obtains uniformly on the class of all sequences of preference profiles.

To show that asymptotic envy-freeness is implied by symmetry and asymptotic strategy-proofness, we restrict attention to mechanism satisfying a regularity condition known as asymptotic non-atomicity. A sequence of random mechanisms $\phi_q : N_q \to \mathcal{M}_q$ is *asymptotically non-atomic on a sequence of preference profiles* $\succ_{N_q}$ if

$$\max_{i,j\in N_q,\, i\neq j,\, \succ'_i\in\mathcal{P},\, a\in\Theta} \left|\phi_q\left(\succ_i, \succ_{N_q-\{i\}}\right)(j,a) - \phi_q\left(\succ'_i, \succ_{N_q-\{i\}}\right)(j,a)\right| \to 0 \quad \text{as} \quad q\to\infty.$$

We say that asymptotic non-atomicity obtains uniformly on a class of sequences of preference profiles if the convergence obtains uniformly on this class. A sequence of mechanisms is *asymptotically non-atomic* if the convergence obtains uniformly on the class of all sequences of preference profiles.

In words, a sequence of random mechanisms is asymptotically non-atomic if the impact on allocations of other agents from a preference change by one agent vanishes as the economy grows. Asymptotic non-atomicity is a natural regularity condition – as markets grow we expect individuals' impact on allocations of other agents to become arbitrarily small; see Debreu and Scarf (1963) and Aumann (1964).

Remark 1. Asymptotic non-atomicity of Probabilistic Serial is straightforward. Random Priority is asymptotically non-atomic for asymptotically full-support preference profiles. 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 seminal work on Random Priority). The proof that Random Priority is asymptotically non-atomic is the same as in Liu and Pycia (2011).

It is straightforward to observe that in large asymptotically non-atomic markets, symmetry and strategy-proofness are equivalent to asymptotic envy-freeness.[15]

Theorem 3. *For any asymptotically non-atomic sequence of random mechanisms ϕ_q, the mechanisms are asymptotically symmetric and asymptotically strategy-proof if and only if they are asymptotically envy-free.*

This result and the above discussion allow us to conclude that asymptotic envy-freeness is a mild assumption. Theorems 2 and 3 furthermore imply

Corollary 1. *Assume the feasibility structure satisfies agent-specific capacities and that two sequences of random mechanisms ϕ_q and ϕ'_q are each (i) asymptotically non-atomic, (ii) asymptotically ordinally efficient, and (iii) either asymptotically envy-free, or asymptotically symmetric and asymptotically strategy-proof. If a sequence of preference profiles $\succ_{N_q}$ has asymptotically full-support, then the sequences of allocations $\phi_q\left(\succ_{N_q}\right)$ and $\phi'_q\left(\succ_{N_q}\right)$ asymptotically coincide, that is*

$$\max_{i\in N_q, a\in\Theta}\left|\phi_q\left(\succ_{N_q}\right)(i,a)-\phi'_q\left(\succ_{N_q}\right)(i,a)\right|\to 0 \quad \text{as} \quad q\to\infty.$$

As in Liu and Pycia (2011), we will later see that – in addition to Probabilistic Serial – Random Priority, and many other mechanisms satisfy the conditions of this equivalence result.

Analogues of the above two results are true when formulated uniformly on any class of sequences of preference profiles $\succ_{N_q}$, and resulting sequences of allocations $\phi_q\left(\succ_{N_q}\right)$.

[15] We apply the efficiency and no envy terms directly to mechanisms: a sequence of mechanisms is asymptotically ordinally efficient if the mechanisms generate asymptotically ordinally efficient allocations for every sequence of preference profiles; similarly, a sequence of mechanisms is asymptotically envy-free if the mechanisms generate asymptotically envy-free allocations for every sequence of preference profiles.

5.2 Main Results on Allocation Mechanisms in Large Markets

Results of Section 4 are derived for asymptotically full-support sequences of preference profiles. The analogues of these results are true for all preference profiles if we impose mild continuity assumptions on the mechanisms.

A sequence of mechanisms ϕ_q is *asymptotically equicontinuous* if for every $\epsilon > 0$, and every q large enough, there is $\delta > 0$ such that for every agent $j \in N_q$ the inequality

$$\max_{a \in \Theta} \left| \phi_q \left(\succ_{N_q} \right) (j, a) - \phi_q \left(\succ'_{N_q} \right) (j, a) \right| < \epsilon, \tag{1}$$

is satisfied for all $\succ_{N_q}, \succ'_{N_q} \in \mathcal{P}^{N_q}$ such that $\succ'_j = \succ_j$ and

$$\frac{\left| \{ i \in N_q | \succ'_i \neq \succ_i \} \right|}{|N_q|} < \delta. \tag{2}$$

Asymptotic equicontinuity is stronger than asymptotic non-atomicity. It is an asymptotic and ordinal counterpart of the uniform equicontinuity of Kalai (2004).[16] Continuity of large market allocation has been studied by Hurwicz (1979) and Dubey, Mas-Colell, and Shubik (1980). Champsaur and Laroque (1982) directly address the need for such an assumption.

Remark 2. Asymptotic equicontinuity of Probabilistic Serial is straightforward to demonstrate. Random Priority is asymptotically equicontinuous provided

$$\left| \theta_q^{-1} (a) \right| \to \infty \quad \text{as } q \to \infty. \tag{3}$$

The proof is the same as in Liu and Pycia (2011).

Imposing asymptotic equicontinuity allows us to extend the claim of Theorem 2 to all sequences of preference profiles.

[16] Kalai imposes the continuity assumption uniformly on all games (mechanisms) rather than only in an asymptotic limit, and he requires agents' utilities rather than their allocations to be ϵ-close. This assumption is at the core of his analysis of a general class of large market games.

Imposing the asymptotic equicontinuity assumption allows us to derive our main equivalence results for allocation mechanisms in large markets.[17]

Corollary 2. *Assume the feasibility structure satisfies agent-specific capacities, and that the sequences of random mechanisms ϕ_q and ϕ'_q are (i) asymptotically equicontinuous, (ii) asymptotically ordinally efficient, and (iii) either asymptotically envy-free, or asymptotically symmetric and asymptotically strategy-proof. Then, the sequences of mechanisms coincide asymptotically and uniformly across all preference profiles, that is*

$$\max_{\succ_{N_q}\in\mathcal{P}_q, i\in N_q, a\in\Theta}\left|\phi_q\left(\succ_{N_q}\right)(i,a)-\phi'_q\left(\succ_{N_q}\right)(i,a)\right|\rightarrow 0 \quad as \quad q\rightarrow\infty.$$

The proof is the same as in Liu and Pycia (2011).

To be able to apply this equivalence result we need to know which mechanisms – other than Probabilistic Serial – are asymptotically ordinally efficient. We explore this question in the next section.

6 Ex-Post Pareto Efficiency And Asymptotic Ordinal Efficiency

Che and Kojima (2010) and Liu and Pycia (2011) showed that – with single-unit demands – Random Priority is ordinally efficient when the number of copies of each object grows to infinity. In this section we extend the general criterion for asymptotic ordinal efficiency to multi-unit demands.

As in Liu and Pycia (2011) let us say that a sequence of allocations can be implemented in a Pareto-efficient and asymptotically-uncorrelated way if the allocations can be implemented as lotteries over Pareto-efficient deterministic allocations in such

[17] An analogue of this corollary, with the same proof, holds true holds true uniformly on any class of sequences of preference profiles, $\mathcal{Q}_q\subseteq\mathcal{P}^{N_q}$. We may then relax the equicontinuity assumption by restricting it to $\succ_{N_q}\in\mathcal{Q}_q$ (rather than all $\succ_{N_q}\in\mathcal{P}^{N_q}$).

a way that random allocations of agents with identical preferences are asymptotically uncorrelated. Formally, a sequence of allocations μ_q can be *implemented in a Pareto-efficient and asymptotically-uncorrelated way* if there exists a probability space Ω such that conditional allocations $\mu_q(\cdot,\cdot|\omega)$ for $\omega \in \Omega$ are deterministic and Pareto efficient, and for any $a \in \Theta$ the maximum over $i, j \in N_q$ with the same preference type of the covariance of random variables $X_{i:q} : \Omega \ni \omega \mapsto \mu_q(i, a|\omega)$ and $X_{j:q} : \Omega \ni \omega \mapsto \mu_q(j, a|\omega)$ goes to 0 as $q \to \infty$. The first part of this assumption – that $\mu_q(\cdot,\cdot|\omega)$ are deterministic and Pareto efficient – is known as *ex-post Pareto efficiency.*

Remark 3. Any sequence of allocations generated by Random Priority on a full-support preference profile has Pareto-efficient and asymptotically-uncorrelated implementation. The argument is the same as in Liu and Pycia (2011).[18]

Theorem 4. *Assume the feasibility structure satisfies agent-specific capacities. If a sequence of symmetric mechanisms ϕ_q is asymptotically equicontinuous, and random allocations $\phi_q\left(\succ_{N_q}\right)$ can be implemented in Pareto-efficient and asymptotically uncorrelated way for any $\succ_{N_q} \in \mathcal{P}^{N_q}$ with asymptotically full-support, then mechanisms ϕ_q are asymptotically ordinally efficient.*

The proof of this theorem relies on the following lemma.

Lemma 1. *Assume the feasibility structure satisfies agent-specific capacities, and fix a sequence of preference profiles $\succ_{N_q}$ with asymptotically full-support. If a sequence of symmetric random allocations μ_q can be implemented in Pareto-efficient and asymptotically uncorrelated way, then it is asymptotically ordinally efficient.*

This lemma is of independent interest as it shows that asymptotic ordinal efficiency obtains on asymptotically full-support profiles (an asymptotically generic class of

[18]The convergence is uniform on any class of uniformly asymptotically full-support profiles. Note also that this result relies on agents having global capacity constraints, as opposed to capacity constraints on each object separately; in the latter case Random Priority is not ordinally efficient as shown by Budish and Cantillon (2010).

profiles) even if the asymptotic equicontinuity assumption is violated.[19]

The proofs of the lemma and of the theorem are the same as in Liu and Pycia (2011).

Theorem 4 and Remarks 2 and 3 allow us to answer the questions posed at the beginning of this section: all the mechanisms listed are asymptotically ordinally efficient as long as the number of copies of objects is unbounded as the economy grows (with no further assumption on the rate of growth).

Let us finish by noting the following direct corollary of Theorem 4 and 2:

Corollary 3. *Assume the feasibility structure satisfies agent-specific capacities. All asymptotically equicontinuous sequences of symmetric and asymptotically strategy-proof mechanisms* ϕ_q *such that random allocations* $\phi_q\left(\succ_{N_q}\right)$ *can be implemented in Pareto-efficient and asymptotically uncorrelated way for any* $\succ_{N_q}\in\mathcal{P}^{N_q}$ *with asymptotically full-support, coincide asymptotically.*

This corollary phrases the conditions for convergence in easy-to-verify terms.

7 Conclusion

In Liu and Pycia (2011), we established asymptotic equivalence of a broad class of mechanisms that include Probabilistic Serial and Random Priority.[20] We have shown that all these mechanisms are symmetric, asymptotically ordinally efficient, and asymptotically strategy-proof (and also asymptotically envy-free). In large markets, the choice among these mechanisms need to be based on criteria other than efficiency or fairness. We remarked in Conclusion that the results remain valid in the multiple-unit assignment setting. The present note's role has been to explicate this comment.

[19] A uniform analogue of the lemma holds true, see Liu and Pycia (2011).

[20] The equivalence class established by Liu and Pycia (2011) includes also such mechanisms as symmetric randomization over Pápai (2000) Hierarchical Exchange, and Pycia and Ünver (2011a) Trading Cycles, extended to the setting with copies by Pycia and Ünver (2011b).

References

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

Aumann, R. J. (1964): "Markets with A Continuum of Traders," *Econometrica*, 39–50.

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.

Champsaur, P. and G. Laroque (1982): "A Note on Incentives in Large Economies," *The Review of Economic Studies*, 49, 627–635.

Che, Y.-K. and F. Kojima (2010): "Asymptotic Equivalence of Random Priority and Probabilistic Serial Mechanisms," *Econometrica*, 78, 1625–1672.

Debreu, G. and H. Scarf (1963): "A Limit Theorem on the Core of an Economy," *International Economic Review*, 4, 235–246.

DUBEY, P., A. MAS-COLELL, AND M. SHUBIK (1980): "Efficiency properties of strategies market games: An axiomatic approach," *Journal of Economic Theory*, 22, 339–362.

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

GIBBARD, A. (1977): "Manipulation of Schemes That Mix Voting with Chance," *Econometrica*, 45, 665–681.

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

HEO, E. J. (2011a): "Consumption Schedules and the Serial Rule: Shorter Proofs of Liu and Pycia (2011)," .

——— (2011b): "Probabilistic Assignment with Quotas: A Generalization and a Characterization of the Serial Rule," .

HURWICZ, L. (1979): "On allocations attainable through Nash equilibria," *Journal of Economic Theory*, 21, 140–165.

HYLLAND, A. AND R. ZECKHAUSER (1979): "The Efficient Allocation of Individuals to Positions," *Journal of Political Economy*, 87, 293–314.

JACKSON, M. O. AND I. KREMER (2007): "Envy-freeness and Implementation in Large Economies," *Review of Economic Design*, 11, 185–198.

KALAI, E. (2004): "Large Robust Games," *Econometrica*, 72, 1631–1665.

KASAJIMA, Y. (2009): "More on the probabilistic assignment of indivisible goods when agents receive several," .

KESTEN, O., M. KURINO, AND M. U. ÜNVER (2011): "Fair and Efficient Assignment via The Probabilistic Serial Mechanism," Working paper.

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. (2000): "Strategyproof Assignment by Hierarchical Exchange," *Econometrica*, 68, 1403–1433.

Pycia, M. (2011a): "Assignment with Multiple-Unit Demand and Responsive Preferences," .

——— (2011b): "On Envy-Freeness," .

Pycia, M. and M. U. Ünver (2011a): "Incentive Compatible Allocation and Exchange of Discrete Resources," UCLA and Boston College, unpublished mimeo.

Pycia, M. and U. Ünver (2011b): "Trading Cycles for School Choice," UCLA and Boston College, unpublished mimeo.

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.