Ordinal Efficiency, Fairness, and Incentives in Large Markets*

Qingmin Liu[†] and Marek Pycia[‡]

June 27, 2011

Abstract

We show 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 builds 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 give an easy-to-verify condition for asymptotic ordinal efficiency.

*We thank Andrew Atkeson, Simon Board, Yeon-Koo Che, Moritz Meyer-ter-Vehn, George Mailath, Ichiro Obara, Joseph Ostroy, Mallesh Pai, Andy Postlewaite, Utku Ünver, Kyle Woodward, William Zame, and seminar audiences at Northwestern Matching Workshop, UCLA, and UPenn for valuable comments. Keywords: large market, asymptotic ordinal efficiency, asymptotic strategy-proofness, symmetry, envy-freeness, ordinal efficiency, Random Priority, Probabilistic Serial.

[†]Department of Economics, University of Pennsylvania, 428 McNeil, 3718 Locust Walk, Philadelphia, PA 19104; qingmin (at) econ.upenn.edu

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

1 Introduction

Efficiency and fairness are the twin goals in designing mechanisms to allocate objects. We study them in environments without monetary transfers; examples include assigning school seats to students, and allocating university and public housing.[1] In these examples – and in our model – there are many agents relative to the number of object types (also referred to as objects), each object type is represented by one or more indivisible copies, and each agent consumes at most one object copy.[2] Agents are indifferent among copies of the same object, and have strict preferences among objects. Because object copies are indivisible, fair allocation mechanisms allocate objects randomly. We focus on mechanisms in which the random allocation depends only on agents' reports of their ordinal preferences over objects.[3]

The natural efficiency criterion is ordinal efficiency: an allocation is ordinally efficient if no other allocation first-order stochastically dominates it for all agents. The baseline fairness criterion is symmetric (or equal) treatment of equal agents: an allocation is symmetric if any two agents who reported the same preference ranking are allocated objects according to the same distribution. A more demanding fairness criterion is envy-freeness (or, no envy): an allocation is envy-free if each agent first-order stochastically prefers his distribution over objects to those of other agents.[4]

[1]See Balinski and Sönmez (1999) and Abdulkadiroğlu and Sönmez (2003) for the theory of school seat assignment, and Abdulkadiroğlu, Pathak, Roth, and Sönmez (2005b) and Abdulkadiroğlu, Pathak, and Roth (2005a) for a discussion of practical consideration in seat assignment. See Abdulkadiroğlu and Sönmez (1999) and Chen and Sönmez (2002) for a discussion of house allocation.

[2]In school seat assignment, the number of students is large relative to the number of schools; in allocation of university housing, e.g. at Harvard, MIT, or UCLA, the set of rooms is partitioned into a small number of categories, and rooms in the same category are treated as identical. A related setting was studied by Che and Kojima (2010).

[3]In making this assumption, we follow the prior literature. The standard reasons the literature focused on ordinal mechanisms are (i) learning and reporting one's preference ordering is simpler than learning and reporting one's cardinal utilities, and (ii) ordinal preferences over sure outcomes do not rely on agents' attitude towards risk; we assume instead that an agent prefers one random outcome over another if the former first-order stochastically dominates the latter.

[4]Ordinal efficiency have been introduced by Bogomolnaia and Moulin (2001), and analyzed among others by Abdulkadiroğlu and Sönmez (2003) and McLennan (2002) (see also Postlewaite and Schmeidler (1986) for an early exploration in voting). No envy was introduced by Foley (1967) and Varian (1974), and its variants have been studied by, among others, Schmeidler and Vind (1972),

We make several related contributions to the study of efficient and fair allocations. First, we prove a surprising finite-market equivalence between ordinally efficient and envy-free allocations and the celebrated Probabilistic Serial mechanism of Bogomolnaia and Moulin (2001).[5] We then use the ideas illustrated by this equivalence to derive our main results that – taken together – imply that all asymptotically ordinally efficient, symmetric, and asymptotically strategy-proof ordinal allocation mechanisms that satisfy a mild continuity condition coincide asymptotically.[6]

One important message from these results is that if we care about asymptotic strategy-proofness and treating agents symmetrically then – in large markets – we cannot substantially improve upon the mechanisms we already know and use. In particular, we establish asymptotic ordinal efficiency of a large class of mechanisms – many of those used in practice – and show that the differences among the random allocations generated by these mechanisms vanish in large markets. The allocational equivalence is a strong argument in favor of choosing among these mechanisms primarily on the basis of market-specific implementation considerations.[7]

---

and Alkan, Demange, and Gale (1991).

[5]Bogomolnaia and Moulin (2001) showed that Probabilistic Serial is ordinally efficient and envy-free; the converse implication is new. The mechanism was earlier studied, in a restricted environment, in Crès and Moulin (2001).

[6]A mechanism is strategy-proof if reporting preferences truthfully is a weakly dominant strategy. A mechanism is asymptotically strategy-proof if it is approximately strategy-proof, and the approximation error vanishes as the market becomes large; asymptotic ordinal efficiency and other asymptotic concepts are defined analogously. For the large literature on asymptotic strategy-proofness see footnote 12. We are not aware of prior formal definition of asymptotic ordinal efficiency; Che and Kojima (2010) talk informally about vanishing inefficiency.

[7]In particular, the asymptotic results lend support for the common use of the Random Priority mechanism (Abdulkadiroğlu and Sönmez, 1998) as arguably the simplest and most transparent among these mechanisms. To allocate objects, Random Priority first draws an ordering of agents from a uniform distribution over orderings, and then allocates the first agent a copy of her most preferred object, then allocates the second agent a copy of his most preferred object that still has unallocated copies, etc. Asymptotic ordinal efficiency of Random Priority was established by Che and Kojima (2010); their study left open the question – answered by us – whether there are other asymptotically ordinally efficient and fair allocations preferable to Random Priority on grounds other than efficiency and fairness. This conclusion is qualified by the finite-market result which may be interpreted in favor of using Probabilistic Serial in environments in which agents' report their preferences truthfully, perhaps because of asymptotic strategy-proofness, or because the environment contains enough object copies for Probabilistic Serial to be strategy-proof (see Kojima and Manea, 2010).

For expositional purposes, let us first describe in detail the finite-market equivalence of ordinal efficiency and envy-freeness with Probabilistic Serial. The mechanism treats copies of an object as a pool of probability shares of the object. Given a preference profile, the random allocation is determined through an "eating" procedure in which, as time passes from 0 to 1, each agent "eats" probability share of the best acceptable object which is still available; an object is available at a point in time if some of its shares have not been eaten before this point in time. We show that there is a unique ordinally efficient and envy-free allocation and that it exactly coincides with the allocation generated by Probabilistic Serial, if the set of reported preferences satisfies a richness condition called "full support," under which any strict ranking of objects is represented by some agent's preference ranking. While true in any finite market, this characterization is particularly relevant when the market is large because in this case full-support preference profiles are asymptotically generic; as the number of agents grows, and the number of object types stays bounded, the ratio of the number of full-support profiles to all profiles goes to 1.[8]

This characterization is surprising because prior literature, starting with Bogomolnaia and Moulin (2001), constructed examples of preference profiles for which there are ordinally efficient and envy-free allocations that differ from the outcome of Probabilistic Serial. Our result implies that, in large markets, such profiles are very rare. Nevertheless, because of their existence, the literature focused on characterizing ordinal efficiency and envy-freeness together with additional conditions – upper invariance in Kesten, Kurino, and Ünver (2011), truncation robustness in Hashimoto and Hirata (2011), and bounded invariance in Bogomolnaia and Heo (2011); these papers show that Probabilistic Serial is the unique mechanism that satisfies ordinal

[8]The study of large markets for allocation of goods has a long tradition; the same or more restrictive concept of large market is at the core of (Debreu and Scarf, 1963) replica economies, Roberts and Postlewaite (1976) examination of incentives, Che and Kojima (2010) study of large markets without transfers, and many others. We differ from this prior literature in that we impose no assumptions on the number of object copies. Manea (2009) examines limits of this concept of large market.

efficiency, no envy, and one of the additional conditions.[9] The remarkable feature of our characterization is that we do not impose any additional conditions on the mechanism beyond the standard assumptions of efficiency and envy-freeness.[10]

We prove three substantive large market results. The first says that in large markets any two mechanisms that are asymptotically ordinally efficient and asymptotically envy-free are asymptotically equivalent for asymptotically generic preference profiles; the equivalence obtains for all preference profiles if we additionally impose a mild asymptotic continuity assumption on the two mechanisms.[11]

The second large-market result says that asymptotic envy-freeness is equivalent to the conjunction of asymptotic symmetry and asymptotic strategy-proofness provided each individual agent's impact on other agents' allocation vanishes as the market becomes large. Asymptotic envy-freeness is thus nothing more than a conjunction of the baseline fairness criterion and a standard incentive compatibility condition.[12] In

[9]Kesten, Kurino, and Ünver (2011) also show that Probabilistic Serial is characterized by non-wastefullness (a weak efficiency criterion), and an additional property they introduce, and Hashimoto and Hirata (2011) also show that Probabilistic Serial is characterized by ordinal efficiency, and two additional properties they introduce. An earlier elegant characterization of Probabilistic Serial was proposed by Bogomolnaia and Moulin (2001) in the case of 3 agents and 3 objects; in this case Probabilistic Serial is the unique mechanism which is ordinally efficient, envy-free, and satisfies an additional incentive compatibility condition (weak strategy-proofness).

[10]A precursor of our characterization was obtained by Bogomolnaia and Moulin (2002); they assume that all agents rank objects in the same way (agents may differ as to which objects are better than receiving no object), and show that there is a unique envy-free and ordinal efficient allocation, and that is achieved by Probabilistic Serial.

[11]Asymptotic envy-freeness is defined in Jackson and Kremer (2007); they also note that it is related to incentive compatibility. Asymptotic continuity assumptions have a long tradition in the studies of large markets, see Debreu and Scarf (1963); Aumann (1964); Hurwicz (1979), and Dubey, Mas-Colell, and Shubik (1980).

[12]Asymptotic strategy-proofness is a weak incentive compatibility condition that has been intensively studied since – building on a seminal analysis of asymptotic incentives by Roberts and Postlewaite (1976) – Hammond (1979); Champsaur and Laroque (1982), and Jackson (1992) proved it for the Walrasian mechanism in large exchange economies. See, for instance, Peleg (1979) (voting), and Gretsky, Ostroy, and Zame (1999) (assignment with transfers), and Roth and Peranson (1999); Immorlica and Mahdian (2005); Kojima and Pathak (2008) (two-sided matching). Azevedo and Budish (2011) review the literature on asymptotic strategy-proofness, and make a general case in favor of imposing this requirement. Strategy-proof mechanisms are hard to manipulate, robust to specification of agents' beliefs, impose minimal costs of searching for and processing strategic information, and do not discriminate among agents based on their access to information and ability to strategize (c.f. Vickrey (1961); Dasgupta, Hammond, and Maskin (1979); Pathak and Sönmez (2008); for some examples of a relaxation of this condition, see Ergin and Sönmez (2006) and Ab-

particular, asymptotic strategy-proofness of asymptotically envy-free mechanisms is reassuring because ordinal efficiency is a property of allocation with respect to the reported preference profile, and only knowing that agents have vanishingly small incentives to misreport allows us to interpret ordinal efficiency with respect to reported preferences as a welfare criterion.

The third large-market result is an easy-to-verify condition for asymptotic ordinal efficiency. Under a mild asymptotic continuity condition, for a symmetric mechanism to be asymptotically ordinally efficient it is enough that its random allocations can be implemented as lotteries over Pareto-efficient deterministic allocations in such a way that random allocations of agents with identical preferences are asymptotically uncorrelated.[13] Many known mechanisms – including Random Priority (Abdulkadiroğlu and Sönmez, 1998), and symmetric randomizations over Hierarchical Exchange of Pápai (2000) and Trading Cycle mechanisms of Pycia and Ünver (2009) (extended to the setting with copies by Pycia and Ünver, 2011) – satisfy the implementation condition, and, if the number of object copies grows unboundedly as the economy becomes large, they also satisfy the asymptotic continuity condition. We can conclude that all these mechanisms are asymptotically ordinally efficient. They are also strategy-proof, and thus they coincide asymptotically with each other, and with Probabilistic Serial.[14]

The closest forerunner of our paper is Che and Kojima (2010). They showed that Random Priority and Probabilistic Serial are asymptotically equivalent, and both are asymptotically ordinally efficient, symmetric, and asymptotically strategy-proof. Our general results allow us to relax their assumptions on the growth rate of the

dulkadiroğlu, Che, and Yasuda (2009)).

[13]Even when the continuity condition fails, the asymptotic ordinal efficiency obtains for asymptotically generic preference profiles.

[14]Abdulkadiroğlu and Sönmez (1998) proved that Random Priority and the Core from Random Endowments (a uniform randomization over Gale's Top Trading Cycles, Shapley and Shubik, 1972) are equivalent in environments in which each object has a single copy, and Pathak and Sethuraman (2010) showed that uniform randomization over Abdulkadiroğlu and Sönmez (2003) Top Trading Cycles for School Choice coincides with Random Priority in environments with copies. The above corollary of our main result provides a large market counterpart of this surprising equivalence for any symmetric randomization (not only uniform) over the substantially larger classes of Hierarchical Exchange and Trading Cycles mechanisms.

number of copies, and thus gain new insights into asymptotic ordinal efficiency of Random Priority and asymptotic strategy-proofness of Probabilistic Serial.[15] More importantly, we show that the equivalence they discovered is not a coincidence but rather a fundamental property of allocation in large markets: in large markets there is effectively only one ordinal allocation mechanism that satisfies the standard postulates of efficiency, fairness, and incentive compatibility.

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. 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 unit demands and strict preferences over objects from Θ. Agent's preference ranking is also referred to as the type of the agent. 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. One natural interpretation is that object types represent schools, and object copies represent seats in these schools. 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, $|\theta^{-1}(\oslash)| \geq |N|$. An object is called acceptable if it is preferred to $\oslash$.

[15]Kojima and Manea (2010) show that agents' have incentives to report preferences truthfully in Probabilistic Serial if the number of copies is large enough relative to a measure of variability of agents' utility; Che and Kojima (2010) and our results on asymptotic strategy-proofness of Probabilistic Serial does not rely on assumptions on agents' utility.

A *random allocation* μ is determined by probabilities $\mu(i,a) \in [0,1]$ that agent i is assigned object type a.[16] All random allocations studied in this paper are assumed to be feasible in the following sense

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

The set of these random allocations is denoted by $\mathcal{M}$. A *random 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 random 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. 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 allocation. Ordinal efficiency is a weak

[16] A random 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. Hylland and Zeckhauser (1979) and Bogomolnaia and Moulin (2001) showed how to implement random allocations. The implementation relies on the Birkhoff and von Neumann's theorem. For recent work on decomposition of random allocations, see Budish, Che, Kojima, and Milgrom (2011).

and natural efficiency requirement: if an allocation is not ordinarily efficient, then all agents would ex ante agree there is a better one. P Bogomolnaia and Moulin (2001) discuss this requirement in depth. 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

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

Envy-freeness (referred to also as no envy) is a strong fairness requirement introduced by Foley (1967). Footnote 4 gives more historic details on the two concepts.

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 $t \in [0,1)$ 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

$$\begin{aligned} A(t) &= \{a \in \Theta |\, Q_a(t) > 0\}, \\ Q_a(t) &= Q_a(0) - \int_0^t |\{i \in N|\ a \in A(\tau) \text{ and } \forall b \in A(\tau)\ a \succsim_i b\}|\, d\tau. \end{aligned}$$

We say that *agent i eats from object a at time t* iff $a \in A(t)$ and $\forall b \in A(t)\ a \succsim_i b$. If stopped at time t, the eating procedure allocates object $a \in \Theta$ to agent $i \in N$ with probability

$$\psi^t(i,a) = \int_0^t \chi(i \text{ eats from } a \text{ at time } \tau)\, d\tau,$$

where the Boolean function χ (statement) takes value 1 if the statement is true and 0 otherwise. The allocation $\psi(i,a)$ of *Probabilistic Serial* is given by the eating procedure stopped at time 1; that is $\psi = \psi^1$.

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.[17]

3.2 Main Finite-Market Result

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 for each ranking of objects, there exists an agent whose preferences over objects agree with this ranking.[18] The restriction to full-support preference profiles is strong in small markets, however as the market becomes large the restriction becomes mild: as the number of agents grow while the number of object types stays constant, the proportion of the number of full-support profiles to the number of all preference profiles goes to 1.

Theorem 1. *For every full-support preference profile, 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 proved by Bogomolnaia and Moulin (2001). The converse implication is new and relies on the

[17]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).

[18]For our results, it is enough to assume for each ranking of objects such that all objects are acceptable, there exists an agent whose preferences over objects agrees with this ranking.

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)).

Proof. Fix any full-support preference profile and random allocation μ that is envy-free and ordinally efficient. To prove the theorem it is enough to show that

$$\sum_{a' \succsim_i a} \mu(i, a') \geq \sum_{a' \succsim_i a} \psi^t(i, a') \tag{1}$$

for all $t \in [0, 1]$, agents $i \in N$, and objects $a \in \Theta$. Indeed, this set of inequalities for $t = 1$, together with ordinal efficiency of Probabilistic Serial ψ^1 imply that $\mu(i, a) = \psi^1(i, a)$ for all $i \in N, a \in \Theta$.

By way of contradiction, assume the above inequality fails for some time, agent, and object. Let T be the infimum of $t \in [0, 1]$ such that there exist $i \in N$ and $b \in \Theta$ such that $\sum_{a \succsim_i b} \mu(i, a) < \sum_{a \succsim_i b} \psi^t(i, a)$. Since there are finite number of agents and objects, there is an agent and object for which the infimum is realized; let us fix such an agent and such an object, and call them i and b, respectively. Let us assume that, among objects for which the infimum is realized, b is the highest ranked in i's preferences .

We proceed in several steps.

Step 1. Inequalities (1) are satisfied for all $t \in [0, T]$. In particular, the cutoff time T belongs to $[0, 1)$. Indeed, by definition, inequalities (1) are satisfied for all $t \in [0, T)$. Because the inequalities are satisfied when $t = 0$, and the mapping $t \mapsto \psi^t(i, a)$ is continuous, inequalities (1) are also satisfied for $t = T$.

Step 2. In the eating procedure, agent i must be eating from b at time T. Indeed, if i is eating from an object $a \succ_i b$ at T, then $\psi^T(a') = 0$ for all objects $a' \prec_i a$, and hence if (1) is violated for agent i and object b then it is violated for agent i and object a. This would contradict the assumption i ranks b above all other objects for which the infimum T is realized. If i is eating from an object $a \prec_i b$ at time T then

$\sum_{a' \succsim_i a} \mu(i, a') \geq \sum_{a' \succsim_i b} \psi^T(i, a') = \sum_{a' \succsim_i b} \psi^t(i, a')$ for t just above T, again contrary to T being the infimum of t at which (1) is violated for i and b.

Step 3. Agent i gets object b or better with probability T, that is $\sum_{a' \succsim_i b} \mu(i, a') = T$. Indeed, by Step 2, agent i is eating from b at time T in the eating procedure, and thus $\sum_{a' \succsim_i b} \psi^T(i, a') = T$. Because (1) is satisfied for $t = T$, we get $\sum_{a' \succsim_i b} \mu(i, a') \geq \sum_{a' \succsim_i b} \psi^T(i, a') = T$. The inequality is binding because $t \mapsto \psi^t(i, a')$ are continuous in t and T is the infimum of times at which (1) is violated.

Step 4. If b is the favorite object of an agent $j \in N$, then $\mu(j, b) = T$. Indeed, by Step 1, the top choice object b is still available at time t, and thus $\psi^T(j, b) = T$. Because (1) is satisfied at time T we thus get $\mu(j, b) \geq T$. Furthermore, envy-freeness of μ implies that $\mu(j, b) \leq T$ as otherwise agent i outcome would not $\succeq_i$-first-order stochastically dominate that of j.

Step 5. If b is the favorite object of an agent $j \in N$, then $\psi^1(j, b) > T$. Indeed, if not, then in the eating procedure b would be exhausted at time T contrary to i eating out of b at time T and thus at some times $t > T$.

Step 6. There is an agent $k \in N$ such that $\mu(k, b) > \psi^1(k, b)$. Indeed, by the full-support assumption there is an agent $j \in N$ who ranks b as his first choice. Steps 3 and 4 imply that this agent j gets less b under μ than under ψ. Ordinal efficiency of μ implies that there must be another agent k who gets more b under μ than under ψ^1.

Step 7. There is an object $c \neq b$ that agent k from Step 6 ranks just above b. Indeed, the claim follows from Steps 4, 5, and 6.

Let us fix agent k and object c satisfying Steps 6 and 7.

Step 8. Under μ, agent k gets object b or better with probability strictly higher than T. Indeed, Step 1 and the availability of object b at time T in the eating procedure imply that $\sum_{a \succsim_k c} \mu(k, a) \geq \sum_{a \succsim_k c} \psi^T(k, a) = T - \psi^T(k, b) \geq T - \psi^1(k, b)$. The claim then follows from Step 6.

To conclude the proof, notice that by the full support assumption, there exist an agent j who ranks objects the same way as k except that he puts b first. By Step 6, $\mu(k,b) > 0$, and thus ordinal efficiency of μ implies that $\mu(j,a) = 0$ for all objects $a \succ_k b$. Step 4 thus implies that under μ the probability j gets object c or better equals T, and, by Step 8, it is smaller than the probability k gets these objects. This contradicts envy-freeness of μ. The contradiction proves (1), and the theorem. □

The analogue of Theorem 1 holds true in environment in which all objects are acceptable. The above argument is valid in such environment because in the argument b is a proper object.

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 $\langle N_q, \Theta, O_q \rangle_{q=1,2,\ldots}$ in which the set of object types, Θ, is fixed while the set of agents N_q grows in q; we will assume throughout that $|N_q| \to \infty$ as $q \to \infty$. As discussed in the introduction, similar or more restrictive assumptions are standard in the study of large markets.

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

There are mechanisms, such as Random Priority (Abdulkadiroğlu and Sönmez, 1998), that are not ordinally efficient, but which in large markets — and under some additional assumptions — have only small inefficiencies as demonstrated by Che and Kojima (2010). To formally capture the efficiency properties of such mechanisms, we now introduce a concept of asymptotic ordinal efficiency.

Let us first define an auxiliary concept of ϵ-ordinal efficiency. 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, (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 $a_{k+1} \succ_{i_k} a_k$.[19] In particular, their result implies that any sequence of ordinally efficient allocations is asymptotically ordinally efficient.

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

$$\liminf_q \min_{i,j \in N_q, a \in \Theta} \left[\sum_{b \succsim_i a} \mu_q(i,b) - \sum_{b \succsim_i a} \mu_q(j,b) \right] \geq 0.$$

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.[20]

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.[21] Formally, we say that a sequence of preference-profiles $\succ_{N_q}$ has *asymptotically full support* if there

[19]Condition (i') is known as individual rationality, and condition (ii') as non-wastefullness. 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').

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

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

exists $\delta > 0$ and $\bar{q}$ such that for any $q > \bar{q}$, and for any ranking of objects $\succ \in \mathcal{P}$, the proportion of agents whose $\succ_{N_q}$-ranking agrees with $\succ$ is 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.[22] 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; we formally define asymptotic genericity and show that asymptotically full-support preference profiles are asymptotically generic in Appendix B.

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. *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,*

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

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

Theorem 2. (Uniform Version) *If a class $\mathcal{Q}$ of preference profile sequences has uniformly asymptotic full support, and two classes of allocation sequences*

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

[22] 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 starts with the observation that it is enough to show that μ_q asymptotically coincides with the allocation of Probabilistic Serial. The rest of the proof follows roughly the same outline as the proof of Theorem 1, except that substantive additional care must be taken to handle the approximations. The proof is in Appendix A.

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? It is surprisingly mild — it is implied by two standard postulates: symmetry and asymptotic strategy-proofness.

Symmetry is a basic fairness property of an allocation, and is also known as equal treatment of equals. Given preference profile $\succ_N$, a random allocation μ is *symmetric* 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, $\mu(i,\cdot)=\mu(j,\cdot)$. Our results will in fact rely only on a weak form of this assumption: given a sequence of preference profiles $\succ_{N_q}$, a sequence of random allocations μ_q is *asymptotically symmetric* if

$$\max_{i,j\in N_q \text{ such that} \succ_i=\succ_j,\, a\in\Theta} \left|\mu_q(i,a)-\mu_q(j,a)\right| \to 0 \quad \text{as} \quad q\to\infty.$$

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. Footnote 12 discusses the literature on asymptotic strategy-proofness.

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 a copy of her most preferred object, then allocates the second agent a copy of his most preferred object that still has unallocated copies, etc. (see Abdulkadiroğlu and Sönmez, 1998). To see that Random Priority is asymptotically non-atomic for asymptotically full-support preference profiles note that (i) a change of preferences by agent i from $\succ_i$ to $\succ_i'$ can change the allocation of another agent j only in Random Priority orderings in which agent j takes the last copy of an object under at least one of the two preference rankings submitted by i; (ii) the probability of agent j taking the last copy of an object a under a preference profile is bounded above by $\frac{1}{n}$ where n is the number of agents with preference rankings identical to j's, and hence the probability of agent j taking the last copy of an object under one of two profiles is bounded above by $\frac{2|\Theta|}{n-1}$, and (iii) this probability converges to 0 along any sequence of asymptotically full-support preference profiles because along such sequences $n \to \infty$. A similar argument shows that uniform randomizations over Hierarchical Exchange of Pápai (2000) or Trading Cycles of Pycia and Ünver (2009) (extended to the setting with object copies by Pycia and Ünver (2011)) are asymptotically non-atomic for asymptotically full-support preference profiles.

It is straightforward to observe that in large asymptotically non-atomic markets, symmetry and strategy-proofness are equivalent to asymptotic envy-freeness.[23]

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. *Suppose 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*

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

We will later see that – in addition to Probabilistic Serial – Random Priority, and many other mechanisms satisfy the conditions of this equivalence result.

[23]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.

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 $\phi_q\left(\succ_{N_q}\right)$.

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

$$\max_{a\in\Theta}\left|\phi_q\left(\succ_{N_q}\right)(j,a) - \phi_q\left(\succ'_{N_q}\right)(j,a)\right| < \epsilon, \tag{2}$$

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

$$\frac{\left|\{i \in N_q \mid \succ'_i \neq \succ_i\}\right|}{\left|N_q\right|} < \delta. \tag{3}$$

Asymptotic equicontinuity is stronger than asymptotic non-atomicity. It is an asymptotic and ordinal counterpart of the uniform equicontinuity of Kalai (2004).[24] 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

$$\left|\theta_q^{-1}(a)\right| \to \infty \quad \text{as } q \to \infty. \tag{4}$$

[24]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.

The proof relies on the definition of Random Priority provided in Remark 1, and has three steps.

Step 1. A change of preferences by a fraction δ of agents can change the (deterministic) allocation of another agent j only in Random Priority orderings in which agent j takes the last copy of an object under at least one of the two preferences rankings submitted by the fraction of agents changing their preferences.

Step 2. The probability an agent takes the last copy of an object a under a preference profile vanishes as $q \to \infty$. Indeed, fix q and an ordering of agents other than $j \in N_q$; and consider probabilities conditional on such an ordering. If a is the favorite object for j then j would take it as long as it is available and the conditional probability j takes the last copy of a is bounded above by $\left|\theta_q^{-1}(a)\right|$. If there are objects ("better objects") that agent j prefers over a then consider three cases. If one of the better objects is still available when the last copy of a goes, then the probability j takes the last copy of a is 0. If the better objects are exhausted by some agent with order number t in the sequence, then if object a is still available when order number $t + \left\lfloor \sqrt{\left|\theta_q^{-1}(a)\right|} \right\rfloor$ comes up, then the probability j takes the last copy of a is bounded above by $\frac{1}{\sqrt{\left|\theta_q^{-1}(a)\right|}}$; if a is not available when order number $t + \left\lfloor \sqrt{\left|\theta_q^{-1}(a)\right|} \right\rfloor$ comes up, then the probability j takes the last copy of a is bounded above by $\frac{1}{t}$ and $t \geq \left|\theta_q^{-1}(a)\right| - \left\lfloor \sqrt{\left|\theta_q^{-1}(a)\right|} \right\rfloor$. In either case, the conditional probability of j taking the last object is bounded above by $\max\left\{\frac{1}{\sqrt{\left|\theta_q^{-1}(a)\right|}}, \frac{1}{\left|\theta_q^{-1}(a)\right| - \sqrt{\left|\theta_q^{-1}(a)\right|}}, \right\}$, and so does the unconditional probability.

Step 3. By Step 2, the probability agent j takes the last copy of an object under one of two profiles of Step 1 is bounded above by

$$2\left|\Theta\right| \max\left\{\frac{1}{\sqrt{\left|\theta_q^{-1}(a)\right|}}, \frac{1}{\left|\theta_q^{-1}(a)\right| - \sqrt{\left|\theta_q^{-1}(a)\right|}}, \right\}.$$

By condition (4), this bound goes to zero uniformly over agents and preference profiles, and the claim is true.

An analogous argument can show that mechanisms obtained by uniform randomization over Hierarchical Exchange or Trading Cycles are asymptotically equicontinuous provided the number of object copies satisfies condition (4).

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

Imposing the asymptotic equicontinuity assumption allows us to derive our main equivalence results for allocation mechanisms in large markets.[25]

Corollary 2. *Suppose 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*

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

Proof. First notice that Theorem 2 (Uniform Version) provides a uniform convergence for all sequences of preference profiles$\succ_{N_q}$ such that for some $\delta>0$ and positive integer $\bar{q}$, for all $q>\bar{q}$ each ranking $\succ\in\mathcal{P}$ is represented in $\succ_{N_q}$ by at least fraction δ of agents. A uniform counterpart of Theorem 3 is also true for such sequences of profiles. Because every asymptotically equicontinuous mechanism is asymptotically non-atomic, Theorems 2 and 3 yield the result for asymptotically full-support preference profiles. We can then fix any sequence of profiles $\succ_{N_q}$, and use the convergence for asymptotically full-support sequences of profiles $\succ'_{N_q}$ such that (3) is satisfied, and the asymptotic equicontinuity of ϕ_q, to derive the convergence for $\succ_{N_q}$. □

To be able to apply this equivalence result we need to know which mechanisms

[25]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 $\succ_{N_q}\in\mathcal{P}^{N_q}$).

– 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

Which mechanisms are asymptotically ordinally efficient, besides Probabilistic Serial? Che and Kojima (2010) demonstrate ordinal efficiency of Random Priority in the case in which the number of copies of each object grows at asymptotically the same rate as the number of agents, and offer a counterexample to asymptotic ordinal efficiency in settings with few copies. How far can we relax the rate of growth restriction? What can be said about asymptotic ordinal efficiency of symmetric randomizations over Hierarchical Exchange or over Trading Cycles?

To address these questions, 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 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. We give an argument for Random Priority; the arguments for the other mechanisms are analogous. Take Ω to be the space of sequences of orderings of agents in q-economies; conditional on the ordering of agents Random Priority is deterministic and Pareto efficient. Fix q, a preference profile in $\mathcal{P}^{N_q}$, and an ordering $\succ \in \mathcal{P}$. Denote by n_a the number of agents of type $\succ$ getting a under the fixed preference profile, and let $n = \sum_{a\in\Theta} n_a$. The correlation between two agents i and j of type $\succ$ getting a is the average of such correlations conditional on the profile of numbers n_a, $a \in \Theta$. The symmetry among agents of the same preference type implies that conditional on a profile of n_a, $a \in \Theta$, the covariance is

$$\frac{n_a}{n}\frac{n_a-1}{n-1}\left(1-\frac{n_a}{n}\right)^2+2\frac{n_a}{n}\left(\frac{n-n_a}{n-1}\right)\left(1-\frac{n_a}{n}\right)\left(0-\frac{n_a}{n}\right)\quad + \\ +\left(\frac{n-n_a}{n}\right)\left(\frac{n-1-n_a}{n-1}\right)\left(0-\frac{n_a}{n}\right)^2=\frac{n_a\left(n_a-n\right)}{n^2\left(n-1\right)}\quad\in\quad\left[-\frac{1}{4\left(n-1\right)},0\right].$$

The correlation thus converges to 0 as $q\to\infty$ as along asymptotically full-support profiles, n, the number of agents of type $\succ$ grows to infinity.[26]

Theorem 4. *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. *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 is of independent interest as it shows that asymptotic ordinal efficiency obtains on asymptotically full-support profiles (an asymptotically generic class of profiles) even if the asymptotic equicontinuity assumption is violated.[27]

[26] The convergence is uniform on any class of uniformly asymptotically full-support profiles.
[27] A uniform analogue of the lemma holds true; see the proof of Theorem 4 for an argument.

Proof. To prove asymptotic ordinal efficiency we need to prove that the allocations are $\epsilon(q)$-ordinally efficient for some $\epsilon(q) \to 0$ as $q \to \infty$. By way of contradiction, assume that there is $\epsilon > 0$, and a sequence of $q_n \to \infty$, such that the allocations in the q_n-economies are not ϵ-ordinally efficient. Because there is a finite number of agent types $\succ \in \mathcal{P}$, the compactness of $[0,1]$ allows us to subsample the sequence q_n and assume that the proportion of each type converges to a constant. The limit proportions then add up to 1, and – by asymptotic full support of the preference profiles – are positive.

For any $a \in \Theta$, symmetry of allocations implies that the probability $\mu_{q_n}(i,a)$ is the same for all agents i of a particular type $\succ \in \mathcal{P}$. Because there is a finite number of agent and object types, the compactness of $[0,1]$ allows us also to subsample the sequence q_n and assume that $\mu_{q_n}(i,a)$ converges to a constant $\mu_\infty(\succ, a) \in [0,1]$.

Because ex post Pareto efficiency implies conditions (i) and (ii) of ϵ-ordinal efficiency, it must be condition (iii) that is violated for each q_n. The violation of (iii) would mean that for each q_n there is a cycle of agents $i_0, ..., i_m$ and objects $a_0, ..., a_m$ such that i_k gets a higher-than-ϵ probability of a_k, and $a_{k+1} \succ_{i_k} a_k$ (subscripts modulo $m+1$). Denoting by $\succ_k$ the preference ranking of agent i_k, we get $\mu_\infty(\succ_k, a_k) \geq \epsilon$.

Consider now a Pareto-efficient and asymptotically uncorrelated implementation of μ_{q_n}. Let Ω be the associated probability space. Applying the weak law of large numbers to random variables $X_{i:q} : \Omega \ni \omega \mapsto \mu_q(i, a|\omega)$, we conclude that for any $\tilde{\epsilon} > 0$ and q large enough, the proportion of agents of type $\succ_k$ is within $\tilde{\epsilon}$ of $\mu_\infty(\succ_k, a_k)$ with probability at least $1 - \tilde{\epsilon}$. This implies that there are some agents $i'_0, ..., i'_m$ of types $\succ_0, ..., \succ_m$ (respectively) who are allocated objects $a_0, ..., a_m$ (respectively) at some state of nature ω. This, however, contradicts Pareto efficiency of the allocation $\mu_q(\cdot, \cdot|\omega)$. □

We are now ready to prove Theorem 4:

Proof. Fix a class $\mathcal{S}$ of asymptotically full-support sequences of preference profiles $\succ_{N_q}^{AFS}$ such that for some $\delta > 0$ and positive integer $\bar{q}$, for all $q > \bar{q}$, each ranking

$\succ \in \mathcal{P}$ is represented in $\succ_{N_q}^{AFS}$ by at least fraction δ of agents.

First notice that if the asymptotic ordinal efficiency of $\phi_q\left(\succ_{N_q}^{AFS}\right)$ does not obtain uniformly on $\mathcal{S}$ then there is $\epsilon > 0$ and a sequence of sequences of preference profiles $\left(\succ_{N_q}^{AFS_k}\right)_{q=1,2,...}$, $k = 1, 2, ...$ such that $\phi_q\left(\succ_{N_q}^{AFS_q}\right)$ is not ϵ-ordinally efficient. While the sequence of profiles $\left(\succ_{N_q}^{AFS_q}\right)_{q=1,2,...}$ does not need to belong to $\mathcal{S}$, we can conclude that it has asymptotically full-support. By Lemma 1, $\phi_q\left(\succ_{N_q}^{AFS}\right)$ is asymptotically ordinally efficient, and the contradiction allows us to conclude that the asymptotic ordinal efficiency of $\phi_q\left(\succ_{N_q}^{AFS}\right)$ obtains uniformly on $\mathcal{S}$.

To finish the proof, take any sequence of profiles $\succ_{N_q} \in \mathcal{P}^{N_q}$. There is a sequence of profiles $\left(\succ_{N_q}^{AFS}\right)_{q=1,2,...} \in \mathcal{S}$ such that (3) is satisfied. The above conclusion and inequality (2) imply that $\phi_q\left(\succ_{N_q}\right)$ satisfy conditions (i)-(iii) of ϵ-ordinal efficiency, where $\epsilon \to 0$ as $q \to \infty$, uniformly on $\succ_{N_q} \in \mathcal{P}^{N_q}$. Allocations $\phi_q\left(\succ_{N_q}\right)$ are thus asymptotically ordinally efficient, uniformly on $\succ_{N_q} \in \mathcal{P}^{N_q}$. □

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. *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

Theorem 4, Corollary 2, and Remarks 2 and 3, establish asymptotic equivalence of a broad class of mechanisms that include Probabilistic Serial, Random Priority, and symmetric randomizations over Hierarchical Exchange, and Trading Cycles. 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.

With the exception of the equivalence between Random Priority and Probabilistic Serial discovered in the seminal paper by Che and Kojima (2010), the asymptotic equivalence of this mechanisms, and their ordinal efficiency properties, are new. Furthermore, our general results show that the surprising asymptotic equivalence of Probabilistic Serial and Random Priority is not a coincidence but a fundamental property of allocation in large markets.[28]

While we study single-unit assignment, Theorems 1, 2, and 3 remain valid for multi-unit assignment problems when agents have lexicographic-responsive preferences over bundles of objects (c.f. Roth, 1985); each agent has a strict ranking of objects and chooses between two bundles A and B by comparing the best object from A not available in B with the best object from B not available in A. Our proofs remain valid. In general, this part of our analysis builds on the ingenious construction of Bogomolnaia and Moulin (2001) Probabilistic Serial mechanism. This central mechanism has been extended to many settings and analogues of some of our results

[28]The above results allow us also to obtain new insights into strategy-proofness properties of Probabilistic Serial. Kojima and Manea (2010) showed that agents have incentives to report preferences truthfully in Probabilistic Serial if the number of copies is large enough relative to a measure of variability of an agent's utility, and Che and Kojima (2010) showed asymptotic strategy-proofness of Probabilistic Serial provided the number of copies has asymptotically the same rate of growth as $|N_q|$. Like Che and Kojima (2010), our results does not rely on assumptions on an agent's utility. The results allow us to relax Che and Kojima (2010) assumption on the number of object copies, as well as show that no assumption on the number of copies is needed for asymptotic strategy-proofness at asymptotically full-support preference profiles.

are likely true in some of these more general settings; exploring them is a topic for future work.[29]

A Proof of Theorem 2

It is enough to prove the theorem under the additional assumption that one of the allocation sequences, μ'_q, is generated by Probabilistic Serial. To prove the first part of the theorem, fix any sequence of full-support preference profiles $\succ_{N_q}$, and random allocations μ_q that are envy-free and ordinally efficient with respect to the preference profiles. Asymptotic ordinal efficiency implies that for any small $\epsilon > 0$ and large $M > 0$ there is $\bar{q}$ such that for $q \geq \bar{q}$, the allocations are $\frac{\epsilon}{M^2}$-ordinally efficient, and, in particular, there are no two agents who could swap probability shares of size $\frac{\epsilon}{M^2}$ in some two objects. By asymptotic envy-freeness we can assume that for $q \geq \bar{q}$ each agent's i allocation $\frac{\epsilon}{M^2}$-first order stochastically dominates allocations of any other agent j in agent's i preferences, that is

$$\sum_{b \succsim_i a} \mu_q(i,b) - \sum_{b \succsim_i a} \mu_q(j,b) \geq -\frac{\epsilon}{M^2} \qquad \text{for all } a \in \Theta.$$

We fix $q \geq \bar{q}$ and, to economize on notation, we drop the q-subscript when referring to this fixed economy $N = N_q$ and its allocation $\mu = \mu_q$. As in Section 3, we do not explicitly mention the preference argument when referring to the allocation of Probabilistic Serial ψ.

To prove the theorem it is enough to show that

$$\sum_{a' \succsim_i a} \mu(i,a') \geq \sum_{a' \succsim_i a} \psi^t(i,a') - \frac{\epsilon}{M} \tag{5}$$

[29]For extensions of Probabilistic Serial see Katta and Sethuraman (2006) (non-strict preferences), Yilmaz (2010) (existing property rights), and Kojima (2009) and Budish, Che, Kojima, and Milgrom (2011) (multiple unit assignment). Theorem 4 is unlikely to generalize to multiple-unit assignment; see Budish and Cantillon (2010) for examples of inefficiency of Random Priority in multi-unit assignment.

for all $t \in [0,1]$, agents i, and objects $a \in \Theta$. Indeed, this set of inequalities for $t = 1$, together with $\frac{\epsilon}{M}$-ordinal efficiency of Probabilistic Serial ψ^1 imply that $|\mu(i,a) - \psi^1(i,a)| < \epsilon$ for all i and a, provided M is high enough.

By way of contradiction, assume the above inequality fails for some time, agent, and object. Let T be the infimum of $t \in [0,1]$ such that there exists $i \in N$ and $b \in \Theta$ such that $\sum_{a \succsim_i b} \mu(i,a) < \sum_{a \succsim_i b} \psi^t(i,a) - \frac{\epsilon}{M}$. Since there are finite number of agents and objects, there is an agent and object for which the infimum is realized; let us fix such an agent and such an object, and call them i and b, respectively. Let us assume that b is the highest ranked object in i's preferences for which the infimum is realized. We structure the rest of the proof as to highlight the parallels to the proof of Theorem 1.

Step 1. Inequalities (5) are satisfied for all $t \in [0,T]$. In particular, the cutoff time T belongs to $[0,1)$. Indeed, by definition, inequalities (5) are satisfied for all $t \in [0,T)$. Because the inequalities are satisfied when $t = 0$, and the mapping $t \mapsto \psi^t(i,a)$ is continuous, inequalities (5) are also satisfied for $t = T$.

Step 2. In the eating procedure, agent i must be eating from b at time T. Indeed, if i is eating from an object $a \succ_i b$ at T, then $\psi^T(a') = 0$ for all objects $a' \prec_i a$, and hence if (5) is violated for agent i and object b then it is violated for agent i and object a. This would contradict the assumption i ranks b above all other objects for which the infimum T is realized. If i is eating from an object $a \prec_i b$ at time T then $\sum_{a' \succsim_i a} \mu(i,a') \geq \sum_{a' \succsim_i b} \psi^T(i,a') - \frac{\epsilon}{M} = \sum_{a' \succsim_i b} \psi^t(i,a') - \frac{\epsilon}{M}$ for t just above T, again contrary to T being the infimum of t at which (5) is violated for i and b.

Step 3. Agent i gets object b or better with probability $T - \frac{\epsilon}{M}$, that is $\sum_{a' \succsim_i b} \mu(i,a') = T - \frac{\epsilon}{M}$. Indeed, by Step 2, agent i is eating from b at time T in the eating procedure, and thus $\sum_{a' \succsim_i b} \psi^T(i,a') = T$. Because (5) is satisfied for $t = T$, we get $\sum_{a' \succsim_i b} \mu(i,a') \geq \sum_{a' \succsim_i b} \psi^T(i,a') - \frac{\epsilon}{M} = T - \frac{\epsilon}{M}$. The inequality is binding because functions $t \mapsto \psi^t(i,a')$ are continuous in t and T is the infimum of times at which (5) is violated.

Step 4. If b is the favorite object of agent $j \in N$, then $\mu(j,b) \in \left[T - \frac{\epsilon}{M}, T - \frac{\epsilon}{M} + \frac{\epsilon}{M^2}\right]$. Indeed, by Step 1, the top choice object b is still available at time t in the eating procedure, and thus $\psi^T(j,b) = T$. Because (5) is satisfied at time T we thus get $\mu(j,b) \geq T - \frac{\epsilon}{M}$. Furthermore, envy-freeness of μ implies that $\mu(j,b) \leq T - \frac{\epsilon}{M} + \frac{\epsilon}{M^2}$ as otherwise the outcome of agent i would not $\frac{\epsilon}{M^2}$-first-order stochastically dominate for agent i the outcome of agent j .

Step 5. If b is the favorite object of agent $j \in N$, then $\psi^1(j,b) > T$. Indeed, if not, then in the eating procedure b would be exhausted at time T, contrary to i eating b at time T and thus at some times $t > T$.

Step 6. There is an agent $k \in N$ such that $\mu(k,b) > \psi^1(k,b) + \frac{3\epsilon}{M^2}$. Indeed, by the asymptotic full-support assumption at least a fraction δ of agents ranks b as their first choice. Steps 3 and 4 imply that under μ these agents get at least $\frac{(M-1)\epsilon}{M^2}$ less b than they get under ψ. Because $\delta > 0$ and is independent of M, for M large enough, the $\frac{\epsilon}{M^2}$-ordinal efficiency of μ implies that there must be another agent k who gets $\frac{3\epsilon}{M^2}$ more b under μ than under ψ^1.

Step 7. There is an object $c \neq b$ that agent k from Step 6 ranks just above b. Indeed, the claim follows from Steps 4, 5, and 6.

Let us fix agent k and object c satisfying Steps 6 and 7.

Step 8. Under μ, agent k gets object b or better with probability strictly higher than $T - \frac{\epsilon}{M} + \frac{3\epsilon}{M^2}$. Indeed, Step 1 and the availability of object b at time T in the eating procedure imply that $\sum_{a \succsim_k c} \mu(k,a) \geq \sum_{a \succsim_k c} \psi^T(k,a) - \frac{\epsilon}{M} = T - \psi^T(k,b) - \frac{\epsilon}{M} \geq T - \psi^1(k,b) - \frac{\epsilon}{M}$. The claim then follows from Step 6.

To conclude the proof, notice that by the asymptotic full support assumption, there exists an agent j who ranks objects in the same way as agent k except that i puts b first. By Step 6, $\mu(k,b) > \frac{\epsilon}{M^2}$, and thus lack of swaps of size $\frac{\epsilon}{M^2}$ (the consequence of ordinal efficiency of μ) implies that $\sum_{a \succ_k b} \mu(j,a) < \frac{\epsilon}{M^2}$. Step 4 thus implies that under μ the probability j gets object c or better is between $T - \frac{\epsilon}{M}$ and $T - \frac{\epsilon}{M} + \frac{2\epsilon}{M^2}$, and, by Step 8, it is smaller than the probability k gets these objects.

This contradicts envy-freeness of μ. The contradiction proves (5), and the first part of the theorem.

An examination of the above argument shows that the choice of ϵ, M, and $\bar{q}$ can be made uniformly on a class of preference profile sequences with uniformly asymptotic full-support, proving the second part of the theorem.

B Asymptotic Genericity of Asymptotically Full-Support Sequences of Preference Profiles

A set $\mathcal{S}$ of sequences of preference profiles is *asymptotically generic* if for every $\epsilon > 0$ there exists a sequence of sets $\mathcal{S}_q \subset \mathcal{P}_q$ of preference profiles in q-economies such that for q large enough the ratio $\frac{|\mathcal{S}_q|}{|\mathcal{P}_q|} > 1 - \epsilon$, and all sequences of profiles from $\mathcal{S}_q$ are in the set of sequences $\mathcal{S}$.

Proposition 1. *Asymptotically full-support profiles are asymptotically generic.*

Proof. Let $\mathcal{S}_q^{\delta} \subset \mathcal{P}_q$ be the set of preference profiles such that, for any ranking of objects $\succ$ in the q-economy, the proportion of agents whose ranking agrees with $\succ$ to $|N_q|$ is above δ. Take $\epsilon > 0$ and notice that for q large enough there exists $\delta(\epsilon) > 0$ such that $\frac{\left|\mathcal{S}_q^{\delta(\epsilon)}\right|}{|\mathcal{P}_q|} > 1 - \epsilon$. To complete the proof it is enough to set $\mathcal{S}_q = \mathcal{S}_q^{\delta(\epsilon)}$. □

References

Abdulkadiroğlu, A., Y. Che, and Y. Yasuda (2009): "Resolving Conflicting Preferences in School Choice: the Boston Mechanism Reconsidered," *American Economic Review*, forthcoming.

Abdulkadiroğlu, A., P. A. Pathak, and A. E. Roth (2005a): "The New York City High School Match," *American Economic Review Papers and Proceedings*, 95, 364–367.

Abdulkadiroğlu, A., P. A. Pathak, A. E. Roth, and T. Sönmez (2005b): "The Boston Public School Match," *American Economic Review Papers and Proceedings*, 95, 368–372.

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

——— (1999): "House Allocation with Existing Tenants," *Journal of Economic Theory*, 88, 233–260.

Abdulkadiroğlu, A. and T. Sönmez (2003): "Ordinal Efficiency and Dominated Sets of Assignments," *Journal of Economic Theory*, 112, 157–172.

Abdulkadiroğlu, A. and T. Sönmez (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729–747.

Alkan, A., G. Demange, and D. Gale (1991): "Fair Allocation of Indivisible Goods and Criteria of Justice," *Econometrica*, 59, 1023–1059.

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

Azevedo, E. and E. Budish (2011): "Strategyproofness for "Price-Takers" as a Desideratum for Market Design," Unpublished manuscript.

Balinski, M. and T. Sönmez (1999): "A Tale of Two Mechanisms: Student Placement," *Journal of Economic Theory*, 84, 73–94.

Bogomolnaia, A. and E. J. Heo (2011): "Probabilistic Assignment of Objects: Characterizing the Serial Rule," .

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

——— (2002): "A Simple Random Assignment Problem with a Unique Solution," *Economic Theory*, 19, 623–635.

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 Mechainsms," *Econometrica*, 78, 1625–1672.

Chen, Y. and T. Sönmez (2002): "Improving Efficiency of On-campus Housing: An Experimental Study," *American Economic Review*, 92, 1669–1686.

Crès, H. and H. Moulin (2001): "Scheduling with Opting Out: Improving upon Random Priority," *Operations Research*, 49, 565–577.

Dasgupta, P., P. Hammond, and E. Maskin (1979): "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *Review of Economic Studies*, 46, 185–216.

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.

Ergin, H. and T. Sönmez (2006): "Games of School Choice under the Boston Mechanism," *Journal of Public Economics*, 90, 215–237.

Foley, D. (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.

Gretsky, N. E., J. M. Ostroy, and W. R. Zame (1999): "Perfect Competition in the Continuous Assignment Model," *Journal of Economic Theory*, 88, 60–118.

Hammond, P. J. (1979): "Straightforward Individual Incentive Compatibility in Large Economies," *The Review of Economic Studies*, 46, 263–282.

Hashimoto, T. and D. Hirata (2011): "Characterizations of the Probabilistic Serial Mechanism," Working paper.

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.

Immorlica, N. and M. Mahdian (2005): "Marriage, Honesty, and Stability," *SODA 2005*, 53–62.

Jackson, M. O. (1992): "Incentive Compatibility and Economic Allocation," *Economic Letters*, 299–302.

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.

Katta, A.-K. and J. Sethuraman (2006): “A Solution to The Random Assignment Problem on The Full Preference Domain,” *Journal of Economic Theory*, 131, 231–250.

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.

Kojima, F. and P. A. Pathak (2008): “Incentives and stability in large two-sided markets,” *American Economic Review*, 99, 608–627.

Manea, M. (2009): “Asymptotic Ordinal Inefficiency of Random Serial Dictatorship,” *Theoretical Economics*, 4, 165–197.

McLennan, A. (2002): “Ordinal Efficiency and The Polyhedral Separating Hyperplane Theorem,” *Journal of Economic Theory*, 105, 435–449.

Pápai, S. (2000): “Strategyproof Assignment by Hierarchical Exchange,” *Econometrica*, 68, 1403–1433.

Pathak, P. A. and J. Sethuraman (2010): “Lotteries in Student Assignment: An Equivalence Result,” *Theoretical Economics*, forthcoming.

Pathak, P. A. and T. Sönmez (2008): “Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism,” *American Economic Review*, 98, 1636–1652.

Peleg, B. (1979): “A Note on Manipulability of Large Voting Schemes,” *Theory and Decision*, 11, 401–412.

Postlewaite, A. and D. Schmeidler (1986): "Strategic behaviour and a notion of ex ante efficiency in a voting model," *Social Choice and Welfare*, 3, 37–49.

Pycia, M. and M. U. Ünver (2009): "Incentive Compatible Allocation and Exchange of Discrete Resources," UCLA and Boston College, unpublished mimeo.

Pycia, M. and U. Ünver (2011): "Trading Cycles for School Choice," .

Roberts, D. J. and A. Postlewaite (1976): "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica*, 44, 115–127.

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

Roth, A. E. and E. Peranson (1999): "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 89, 748–780.

Schmeidler, D. and K. Vind (1972): "Fair Net Trades," *Econometrica*, 637–642.

Shapley, L. and M. Shubik (1972): "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111–130.

Varian, H. R. (1974): "Equity, Envy, and Efficiency," *Journal of Economic Theory*, 9, 63–91.

Vickrey, W. (1961): "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.

Yilmaz, Özgür. (2010): "The Probabilistic Serial Mechanism with Private Endowments," *Games and Economic Behavior*, 69, 475–491.