Ordinal Efficiency, Fairness, and Incentives in Large Markets*

Qingmin Liu† and Marek Pycia‡

August 2016

Abstract

Efficiency and symmetric treatment of agents are the primary goals of resource allocation in environments without transfers. Focusing on ordinal mechanisms in which no small group of agents can substantially change the allocations of others, we show that all asymptotically efficient, symmetric, and asymptotically strategy-proof mechanisms lead to the same allocations in large markets. In particular, many mechanisms—both well-known and newly developed—are allocationally equivalent. This equivalence is consistent with prior empirical findings that different mechanisms lead to similar allocations in school choice.

*First draft: February 2011. This 2016 draft is currently being revised. We thank Andrew Atkeson, Simon Board, Yeon-Koo Che, Federico Echenique, Eun Jeong Heo, Hugo Hopenhayn, SangMok Lee, Moritz Meyer-ter-Vehn, George Mailath, Kenneth Mirkin, Andrew Newman, Ichiro Obara, Joseph Ostroy, Andy Postlewaite, Al Roth, Tayfun Sonmez, Kyle Woodward, Leeat Yariv, William Zame, Simpson Zhang, and seminar audiences at Northwestern 2011 Matching Workshop, UCLA, Carnegie Mellon, Northwestern, NBER 2012 Market Design Workshop, Gerzensee, SWET, and Caltech-UCLA Theory Conference for valuable comments. Conversations with Kristof Madarasz, Krzysztof Oleszkiewicz, Ron Siegel, and Utku Ünver were especially important. Pycia gratefully acknowledges financial support from the William S. Dietrich II Economic Theory Center at Princeton University. Keywords: ordinal efficiency, sd-efficiency, asymptotic efficiency, asymptotic strategy-proofness, symmetry, envy-freeness, ex-post Pareto efficiency, Random Priority, Probabilistic Serial, Trading Cycles.

†Department of Economics, Columbia University, 1022 IAB, New York, NY 10027; ql2177 (at) columbia.edu

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

1 Introduction

Efficiency and symmetric treatment of participants (fairness) are the primary goals of resource allocation in environments without monetary transfers, such as assigning school seats to students and allocating university and public housing. 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.[1] Agents are indifferent among copies of the same object and have strict ordinal preferences among objects.

A symmetric mechanism allocates the same distribution over objects to each agent who submitted the same preference profile. To achieve symmetry we need to employ a random mechanism because in any deterministic mechanism an object with a single copy can be assigned to only one of the agents who would like to obtain it. For instance, a serial dictatorship—a mechanism in which the agents are ordered and pick objects in turn—is efficient but asymmetric since there is exactly one agent who is always allocated his or her top choice. The standard procedure to symmetrize deterministic mechanisms is to randomize uniformly over them. The often-used Random-Priority mechanism chooses at random an ordering of agents and then computes the allocation of the serial dictatorship in which agents pick objects according to the randomly selected ordering. Randomization introduces symmetry and preserves the incentive properties of mechanisms. In particular, the Random-Priority mechanism is strategy-proof, that is reporting preferences truthfully first-order stochastically dominates any other strategy.[2] Randomization, however, entails an efficiency loss.

[1]See, for instance, Abdulkadiroğlu and Sönmez (2003) for school seat assignment and Chen and Sönmez (2002) for housing allocation. In school seat assignment, the number of students is large relative to the number of schools; in the 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.

[2]Strategy-proofness in ordinal settings has been studied by Gibbard (1978), Roth and Rothblum (1999), and Bogomolnaia and Moulin (2001).

No symmetric strategy-proof mechanism can be efficient as shown by Bogomolnaia and Moulin (2001). For instance, the allocation of the Random-Priority mechanism is inefficient in the simple case in which we allocate two single-copy objects A, B to four agents and agents 1 and 2 rank object A over object B, while agents 3 and 4 rank object B over object A. In this case, Random Priority assigns agent 1 a positive probability of object B while assigning agent 3 a positive probability of object A; were these two agents to ex ante swap these positive probabilities, they would improve their allocations. In general, an allocation is called efficient if no group of agents can advantageously swap probability shares in objects, and no agent can advantageously swap a probability share in an object for a share in an object that is unallocated with positive probability.[3] While symmetry, efficiency, and strategy-proofness are incompatible in finite markets, they become asymptotically compatible in large markets as shown by Kojima and Manea (2010) and Che and Kojima (2010).[4]

The present paper provides the first general examination of symmetric, asymptotically efficient, and asymptotically strategy-proof mechanisms in large markets. Following the standard practice, we focus on mechanisms that are: (i) ordinal, that is the random allocation depends only on agents' reports of their ordinal preferences over objects, and (ii) regular, that is no small group of agents can substantially impact the allocations of other agents.[5]

Our auxiliary result, Theorem 1, shows that randomizations over efficient mecha-

[3]Equivalently, an allocation is efficient if no other allocation first-order stochastically dominates it. This efficiency concept—known as ordinal efficiency or stochastic-dominance efficiency—was introduced by Bogomolnaia and Moulin (2001).

[4]Kojima and Manea (2010) were the first to prove this important point. They showed that a mechanism known as Probabilistic Serial is asymptotically strategy-proof. This mechanism allocates objects through a symmetric "eating" procedure; it was introduced by Bogomolnaia and Moulin (2001) who showed that it is symmetric and efficient but not strategy-proof. Che and Kojima (2010) is the closest paper to ours, and we discuss it below.

[5]Regularity is a standard requirement in the study of large markets; see e.g. Champsaur and Laroque (1982). The standard reasons to study ordinal mechanisms are (a) real-life no-transfer mechanisms are typically ordinal, (b) learning and reporting one's rankings of sure outcomes is simpler than learning and reporting one's cardinal utilities, and (c) ordinal preferences over sure outcomes do not rely on agents' attitudes towards risk. For analysis that incorporates cardinal utilities see Hylland and Zeckhauser (1979) and Makowski, Ostroy, and Segal (1999).

nisms are asymptotically efficient, that is the maximum size of advantageous swaps vanishes as the market becomes large. In particular, many known mechanisms are asymptotically efficient, and there are no efficiency reasons to favor one of them over others in large markets.[6] Another consequence of Theorem 1 is an equivalence of efficiency and ex-post Pareto efficiency in large markets. A random allocation is called ex-post Pareto efficient if it can be represented as a randomization over efficient deterministic allocations. The above example of Random Priority shows that ex-post Pareto efficient mechanisms do not need to be efficient. It turns out that in large markets the two concepts become equivalent when we restrict attention to uniform randomizations.[7]

Our main result, Theorem 2, shows that there is a unique way to achieve symmetric and asymptotically efficient allocations in an asymptotically strategy-proof way; the maximum difference between probability shares in an object an agent obtains under any two mechanisms vanishes as the market becomes large.[8] This result establishes the large-market equivalence of mechanisms, both known and unknown. The equivalence of Theorem 2 has several important implications.

First, equivalence results, such as Theorem 2, are relevant for contentious policy decisions such as what mechanism to choose to allocate scarce social resources. For instance, Pathak and Sethuraman (2010) describe the deliberations of policymakers

[6]For instance, all uniform randomizations over any Trading Cycle mechanisms of Pycia and Ünver (2016) (extended to the setting with copies by Pycia and Ünver, 2011) are asymptotically efficient.

[7]Manea (2009) shows that Random Priority may fail to be efficient in environments in which both the number of object types and the number of agents are large; hence the above equivalence fails in his setting. For other negative results on the relation of efficiency and ex post Pareto efficiency see McLennan (2002) and Abdulkadiroğlu and Sönmez (2003).

[8] Following the standard practice in the literature—e.g. in Che and Kojima (2010), Pathak and Sethuraman (2010), and in the empirical papers discussed below—we establish equivalence in terms of agents' marginal distributions over allocated objects. A mechanism is asymptotically strategy-proof if it is approximately strategy-proof, and the approximation error vanishes as the market becomes large. Asymptotic strategy-proofness is a weak incentive compatibility condition that has been intensively studied, see for instance Roberts and Postlewaite (1976), Hammond (1979), Champsaur and Laroque (1982), and Jackson (1992). In our context, the case for restricting attention to strategy-proof mechanisms was made by Abdulkadiroğlu and Sönmez (2003) and Ergin and Sönmez (2006).

from the New York City Department of Education on how to assign school seats to students in the supplementary round of the New York City school choice process. Theorem 2 implies that the mechanisms considered in New York City are asymptotically equivalent.[9]

More generally, Theorem 2 shows that to break away from the large market equivalence one needs to relax one of its assumptions. In particular, in large markets we cannot substantially improve upon the mechanisms we already know and use while requiring asymptotic strategy-proofness. This message is relevant for ongoing efforts to construct new mechanisms. For instance, Erdil (2011) showed how to construct mechanisms that are strategy-proof and dominate Random Priority in efficiency terms; our results imply that his mechanisms are asymptotically equivalent to Random Priority if they are asymptotically symmetric.[10]

Second, Theorem 2 suggests a theoretical explanation for empirical findings that different school choice mechanisms lead to allocationally similar outcomes. For instance, using data from NYC school choice, Abdulkadiroğlu, Pathak, and Roth (2009) estimate that the marginal allocations of different mechanisms differ by at most 1.5%. Using data from Boston school choice, Pathak and Sönmez (2008) estimate that the marginal allocations in different allocation scenarios differ by at most .1%. These surprising empirical findings are consistent with our asymptotic equivalence result.[11]

[9]The supplementary round of the New York City school choice process is well described by our model. In particular, the New York City wanted to run a strategy-proof mechanism and treat all students equally at all schools. Pathak and Sethuraman directly established the equivalence of two mechanisms considered by the New York City, see below.

[10]While we formulate Theorem 2 for symmetric mechanisms, its proof relies only on asymptotic symmetry of the mechanisms. For another interesting new class of mechanisms covered by our results, see Featherstone (2011).

[11]Note that Boston and NYC school districts use coarse priorities in their assignments (and randomize uniformly within the same priority class) while our Theorem 2 studies a canonical assignment model in which we abstract away from the issue of priorities. Our insight here is somewhat related to the literature on stable outcomes in large two-sided matching markets—such as resident and hospital matching—in which each side of the market consists of agents with strict preferences over the other side and hence coarseness of priorities and randomization are not allowed. Under their respective assumptions, Immorlica and Mahdian (2005) and Kojima and Pathak (2008) show that the ratio of agents who obtain the same allocation in all stable matchings asymptotically converges to 1, and Lee (2013) shows that the ratio of agents whose utilities in all stable matchings are close asymptotically

Finally, let us notice that Theorems 1 and 2 taken together imply that all uniform and regular randomizations over subclasses of the strategy-proof and efficient deterministic mechanisms (including all Trading Cycle mechanisms) are asymptotically equivalent because any uniform randomization over such mechanisms is regular, symmetric, and strategy-proof. The allocational equivalence is a strong argument in favor of choosing among these mechanisms primarily on the basis of market-specific considerations, such as tradition or simplicity.

While Theorem 2 is likely the first general result showing that all natural large market mechanisms coincide, the equivalence of some important special mechanisms has been proved before. In finite markets, Abdulkadiroğlu and Sönmez (1998) proved that Random Priority and the Core from Random Endowments (a uniform randomization over Gale's Top Trading Cycles) are equivalent. They focused on environments in which each object has a single copy; Carroll (2010) and Pathak and Sethuraman (2010) proved analogous results for environments with multiple copies.

In the large market setting of the present paper, Che and Kojima (2010) showed asymptotic equivalence of two mechanisms: Random Priority and Probabilistic Serial.[12] Furthermore, by proving the equivalence of these two mechanisms, they established the asymptotic efficiency of Random Priority. Their work is thus the closest prior work for both our theorems. The asymptotic efficiency of Random Priority is an important special case of our Theorem 1, while our Theorem 2 shows that the equivalence of the two mechanisms studied by Che and Kojima is not a coincidence but rather a fundamental property of allocation in large markets.

converges to 1. In addition, Echenique, Lee, Yenmez, and Shum (2013) examine conditions for the uniqueness of aggregate stable matchings, and Azevedo and Leshno (2013) study the uniqueness of stable matchings in a two-sided model with continuum of agents.

[12]See footnote 4 for a discussion of Probabilistic Serial.

2 Model

A finite economy consists of a finite set of agents N, a finite set of objects Θ (also referred to as object types), and a finite set of object copies O. By $|a|$ we denote the number of copies of object $a \in \Theta$. To avoid trivialities, we assume that each object is represented by at least one copy.

Agents have unit demands and strict preferences $\succ$ over objects.[13] We sometimes refer to an agent's preference ranking as the type of the agent. $\mathcal{P}$ denotes the set of preference rankings, and $\mathcal{P}^N$ is the set of preference profiles. We assume that Θ contains the null object $\oslash$ ("outside option"), and we assume that it is not scarce, $|\oslash| \geq |N|$. An object is called acceptable if it is preferred to $\oslash$.

We can interchangeably talk about allocating objects and allocating copies; one natural interpretation is that object types represent schools, and object copies represent seats in these schools. We study random allocations. An *allocation* μ is given by probabilities $\mu(i,a) \in [0,1]$ that agent i is assigned object a. An allocation is *deterministic* if $\mu(i,a) \in \{0,1\}$ for all $i \in N$, $a \in \Theta$.[14] All allocations studied in this paper are assumed to be feasible in the following sense

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

The set of these random allocations is denoted by $\mathcal{M}$. A *mechanism* $\phi : \mathcal{P}^N \to \mathcal{M}$ is a mapping from the set of profiles of preferences over objects to the set of allocations.

We are primarily concerned with large but finite markets. To study them 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 assume throughout that $|N_q| \to \infty$

[13]By $\succsim$ we denote the weak-preference counterpart of a strict preference ranking $\succ$.

[14]A random allocation needs to be implemented as a lottery over deterministic allocations; Hylland and Zeckhauser (1979) and Bogomolnaia and Moulin (2001) showed how to implement random allocations. The implementation relies on Birkhoff and von Neumann's theorem.

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 we assume that allocations μ_q and mechanisms ϕ_q are defined on q-economies. The number of copies of object a in the q-economy is denoted by $|a|_q$, and the set of random allocations in the q-economy is denoted by $\mathcal{M}_q$. Notice that we do not impose any assumptions on the sequence of sets of object copies, O_q; in particular, we do not need to restrict attention to replica economies.[15]

We study mechanisms ϕ_q in which the effect of reports of any small groups of agents other than j on the allocation of an agent j vanishes as $q \to \infty$. Formally, a sequence of mechanisms ϕ_q is *regular* if for every $\epsilon > 0$ there is $\delta > 0$ such that for any preference profiles $\succ_{N_q}, \succ'_{N_q}$ if

$$\frac{|\{i \in N_q | \succ'_i \neq \succ_i\}|}{|N_q|} < \delta, \tag{1}$$

then for all agents $j \in N_q$ such that $\succ'_j = \succ_j$ we have:

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

We also use a weaker version of this assumption, in which (2) is satisfied for all but a fraction ϵ of agents $j \in N_q$ such that $\succ'_j = \succ_j$; we refer to this weaker assumption as weak regularity.

As the following remark illustrates, many standard mechanisms are regular.

Remark 1. Regularity of many known mechanisms is straightforward to demonstrate. Take, for instance, Random Priority (Abdulkadiroğlu and Sönmez, 1998). To allocate objects, Random Priority first draws an ordering of agents from a uniform distribution over orderings, and then in turn allocates the first agent a copy of her most preferred object, allocates the second agent a copy of his most preferred object that still has

[15]Some of our results rely on additional assumptions on the ratio $|a|_q / |N_q|$, and we explicitly impose these assumptions when needed.

unallocated copies, etc. This mechanism is regular provided for each object $a \in \Theta$ there is $\delta > 0$ such that

$$\liminf_{q\to\infty} |a|_q / N_q > \eta. \tag{3}$$

The proof fixes $\epsilon < \eta$ and has three steps.[16]

Step 1. Fix a priority ordering of agents. Conditional on this ordering, all allocations are deterministic. A change of preferences by δN_q agents can change the allocation of another agent j only if j takes one of the last δN_q copies of an object under at least one of the two preferences rankings submitted by the agents changing their preferences.

Step 2. The probability an agent takes one of the last δN_q copies 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 one of the last δN_q copies of a is bounded above by $\frac{\delta N_q}{|a|_q - \delta N_q}$ (and for large q is bounded above by $\frac{\delta}{\eta-\delta}$). If there are objects ("better objects") that agent j prefers over a then j can take one of the last δN_q copies of a only after these better objects are exhausted; the probability of this happening is bounded above by $\frac{\delta}{\eta}$.

Step 3. We can thus conclude that the probability that a change of preferences by δN_q agents can change the allocation of another agent j is bounded above by $2\,|\Theta|\,\frac{\delta}{\eta-\delta}$ uniformly over j and preference profiles. The regularity claim is thus true.

A similar argument shows that uniform randomizations over fixed-endowment Hierarchical Exchange mechanisms of Pápai (2000) and their counterpart among Trading Cycles of Pycia and Ünver (2016) (extended to the setting with object copies by Pycia and Ünver (2011)) are regular provided the number of object copies satisfies

[16] To see why assumption (3) is needed for regularity, consider an object $a \in \Theta$ and a sequence of economies such that $|a|_q = 1$. If the preference profile $\succ_{N_q}$ is such that only two agents are interested in object a then a change of preferences by one of them has a large impact on the allocation of the other, violating regularity.

condition (3).[17] The regularity of Probabilistic Serial follows from Che and Kojima (2010) and the argument for Random Priority above; the regularity of Probabilistic Serial can also be established directly.

Furthermore, Step 1 of the above argument shows that any deterministic Serial Dictatorships is weakly regular. A similar argument shows that fixed-endowment Hierarchical Exchange mechanisms of Pápai (2000) and their counterpart among Trading Cycles of Pycia and Ünver (2016) and Pycia and Ünver (2011) are also weakly regular.

3 Auxiliary Result: Efficiency

The natural efficiency concept in our ordinal setting is the ordinal efficiency of Bogomolnaia and Moulin (2001): an allocation is efficient if agents cannot trade probability shares, and if no object is wasted.[18] Formally, an allocation μ is *efficient* with respect to preference profile $\succ$ iff (i) for no $n \geq 2$ there is a cycle of agents $i_1, ..., i_n$ and objects $a_1, ..., a_n$ such that $\mu(i_k, a_k) > 0$ and $a_{k+1} \succ_{i_k} a_k$ for all $k = 1, ..., n$ (that is no agents $i_1, ..., i_n$ can trade probability shares), and (ii) if $a \succ_i b$ and $\mu(i, b) > 0$, then all copies of a are allocated with probability 1 (that is no copy of a is wasted).[19] This concept of efficiency naturally extends the concept of efficiency from deterministic allocations to random allocations. A mechanism ϕ is efficient if $\phi(\succ_N)$ is efficient for all preference profiles $\succ_N$.

To define the asymptotic counterpart of this efficiency concept let us say that, for any $\epsilon > 0$, a random allocation μ is *ϵ-efficient* with respect to a preference profile $\succ$ iff (i) for no $n \geq 2$ there is a cycle of agents $i_1, ..., i_n$ and objects $a_1, ..., a_n$ such that $\mu(i_k, a_k) > \epsilon$ and $a_{k+1} \succ_{i_k} a_k$ for all $k = 1, ..., n$ (that is no agents $i_1, ..., i_n$ can

[17] See footnote 25 in Pycia and Ünver (2016) for an extension of Pápai's fixed-endowment idea.

[18] This concept is also known as stochastic-dominance efficiency because the efficiency of an allocation μ is equivalent to the lack of another allocation that first order stochastically dominates μ.

[19] The statement of condition (i) relies on the standard practice of counting subscripts modulo n so that $a_{k+1} = a_1$ when $k = n$.

trade probability shares larger than ϵ), and (ii) if $a \succ_i b$ and $\mu(i,b) > \epsilon$, then all copies of a are allocated with probability at least $1-\epsilon$ (that is no more than ϵ of a is wasted). Given a sequence of preference profiles $\succ_{N_q}$, a sequence of allocations μ_q is *asymptotically efficient* if for each $q = 1, 2, ...$ there are positive $\epsilon(q) \overset{q\to\infty}{\to} 0$ such that μ_q is $\epsilon(q)$-efficient with respect to $\succ_{N^q}$.

Which mechanisms are asymptotically efficient in large economies? We know some asymptotically efficient mechanisms: Bogomolnaia and Moulin (2001) constructed one such mechanism they called Probabilistic Serial. Recently, Che and Kojima (2010) showed that Random Priority is asymptotically equivalent to Probabilistic Serial; it is easy to see that their result implies that Random Priority is asymptotically efficient in the above sense.

The question remained open for other mechanisms. Many of the random mechanisms used in practice are obtained by uniformly randomizing over regular and efficient deterministic mechanisms. This section addresses the efficiency question for this large class of mechanisms. To define uniform randomizations, for any mechanism ψ and permutation $\sigma : N \to N$ let us denote $\psi^\sigma(\succ)(i,a) = \psi\left(\succ_{(\sigma(1),\ldots,\sigma(|N|))}\right)(\sigma(i), a)$. A mechanism $\phi : \mathcal{P}^N \to \mathcal{M}$ is *a uniform randomization* if there exists a mechanism $\psi : \mathcal{P}^N \to \mathcal{M}$ such that

$$\phi\left(\succ_{(1,\ldots,|N|)}\right)(i,a) = \sum_{\sigma : N \overset{1-1}{\to} N} \frac{1}{|N|!} \psi^\sigma(\succ)(i,a)$$

In general, the component mechanism ψ can be random. We say that ψ is *ex-post Pareto efficient* if, for each preference profile $\succ_N$, the allocation $\psi(\succ_N)$ is equal to a mixture over efficient deterministic allocations. For instance, if ψ is a serial dictatorship then ϕ is Random Priority. Other examples include randomizations over fixed-endowment subsets of Papai's Hierarchical Exchange or Pycia and Unver's Trading Cycles.

Theorem 1. *Regular sequences of uniform randomizations over efficient and weakly* *regular deterministic mechanisms are asymptotically efficient on full-support preference profiles.*

The proof is in the appendix. This result shows that all the above-mentioned mechanisms are asymptotically efficient. This result also shows that in our setting efficiency and ex-post Pareto efficiency become equivalent. This is in contrast to small finite markets in which—following the pioneering work of Bogomolnaia and Moulin (2001)—it has been well understood that ex-post Pareto efficiency is a weaker concept than efficiency.[20]

4 Main Result: Equivalence

As argued in the introduction, the natural postulates for allocation are efficiency, symmetry, and strategy-proofness. Unfortunately, we know that in finite markets no mechanism satisfies these three properties, see Bogomolnaia and Moulin (2001). Requiring only that the properties are satisfied in an asymptotic sense, we address the question of what mechanisms satisfy the three properties in large markets.

Let us first review the definitions of the standard concepts of symmetry and 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, $\mu(i,\cdot) = \mu(j,\cdot)$. A random mechanism ϕ is *strategy-proof* if for any agent $i \in N$ and any profile of preferences $\succ_N$, the allocation agent i obtains by reporting the truth, $\phi\left(\succ_i, \succ_{N-\{i\}}\right)(i,\cdot)$, first-order stochastically dominates the allocation the agent can get by reporting another preference ranking $\succ'_i$, that is

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

[20]See also McLennan (2002), Abdulkadiroğlu and Sönmez (2003), and Che and Kojima (2010).

This is the standard concept of strategy-proofness introduced by Gibbard (1977). We further say that a random mechanism ϕ is *ϵ-strategy-proof* if for any agent $i \in N$ and any profile of preferences $\succ_N$,

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

A sequence of random mechanisms ϕ_q is asymptotically strategy-proof if for each $q = 1, 2, ...$ there are positive $\epsilon(q) \overset{q\to\infty}{\to} 0$ such that μ_q is $\epsilon(q)$-strategy-proof. Footnote 8 discusses the literature on asymptotic strategy-proofness.

Theorem 2. *If two regular sequences of random mechanisms ϕ_q and ϕ_q' are asymptotically efficient, symmetric, and asymptotically strategy-proof, then they are asymptotically equivalent, that is*

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

Theorem 2 has important ramifications for the ongoing work to construct improved allocation mechanisms: in large markets we cannot improve upon the allocations of mechanisms we already know without relaxing the requirement of (approximate) strategy-proofness or of symmetry. Erdil (2011) and Featherstone (2011) provide interesting examples of such an ongoing work.[21]

Theorem 2 also immediately implies the following.

Corollary 1. *All asymptotically strategy-proof sequences of uniform randomizations over Pareto-efficient deterministic mechanisms are asymptotically equivalent.*

[21]The above results allow us to also obtain new insights into the 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 do not rely on assumptions on an agent's utility. The results allow us to relax Che and Kojima's assumption on the number of object copies, as well as to show that no assumption on the number of copies is needed for asymptotic strategy-proofness at full-support preference profile sequences.

As a heuristic for why Theorem 2 is true let us consider a setting with a continuum of agents and assume that the mass of agents of type $\succ$, denoted $|N(\succ)|$, is strictly positive for all $\succ \in \mathcal{P}$. Let $|a|$ be the total mass of object a. One way to understand the continuum setting is as a limit of finite economies. This limit is well-defined provided we assume that (i) the sequence of preference profiles $\succ_{N_q}$ is such that for each preference type $\succ \in \mathcal{P}$, the ratio $\frac{|\{i \in N_q | \succ_i = \succ\}|}{|N_q|}$ converges to a constant $|N(\succ)|$, and that $\frac{|a|_q}{|N_q|}$ converges to a positive constant $|a|$. Because of symmetry of ϕ and ϕ', knowing the total mass of object a allocated to agents of type $\succ$ is equivalent to knowing the probability an agent of type $\succ$ obtains object a. We will refer to this probability as $\mu(i,a)$ or $\mu(\succ, a)$. A key role in our argument is played by the following strong fairness requirement introduced by Foley (1967).[22] 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

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

Since $m(\succ) > 0$ for all types $\succ$, regularity, symmetry and asymptotic strategy-proofness of the sequence of mechanisms ϕ_q allow us to conclude that the limit allocation $(\mu(\succ, a))_{\succ \in \mathcal{P} \times \Theta}$ is envy-free. Thus, proving that there is a unique envy-free and efficient allocation is a step towards proving a weak version of Theorem 2 (weakened by the imposition of the additional restrictive assumptions listed above).[23]

[22]In the proof of Theorem 2 in the appendix we use asymptotic envy-freeness, a concept defined in Jackson and Kremer (2007) who noted its relation to incentive compatibility.

[23]The argument presented here can be used to prove a weak form of Theorem 2 in which we add assumption (b), and we restrict attention to convergence for preference profile sequences such that (a) obtains and $|N(\succ)| > 0$ for all $\succ \in \mathcal{P}$. Indeed, symmetry implies that $\phi_q(\succ_{N_q})(i,a)$ is fully determined by $\succ_{N_q}$, object a, and the preference type $\succ$ of agent i; we can thus write $\phi_q(\succ_{N_q})(\succ_i, a) = \phi_q(\succ_{N_q})(i,a)$. Full-support of $\succ_{N_q}$ ensures that $\phi_q(\succ, a)$ is well-defined for any $\succ$ and large q. At each preference profile along the sequence of profiles $\succ_{N_q}$, the mechanism is given by $(\phi(\succ_{N_q})(\succ, a))_{\succ \in \mathcal{P}, a \in \Theta} \in [0,1]^{\mathcal{P} \times \Theta}$. Given compactness of this latter set, to show that two mechanisms ϕ_q and ϕ'_q coincide asymptotically along $\succ_{N_q}$ it is enough to show that every convergent subsequence of $(\phi(\succ_{N_q})(\succ, a))_{\succ \in \mathcal{P}, a \in \Theta}$ converges to the same limit. We can thus consider a subsequence $(\phi(\succ_{N_q})(\succ, a))_{\succ \in \mathcal{P}, a \in \Theta}$ that converges to a limit $(\mu(\succ, a))_{\succ \in \mathcal{P}, a \in \Theta}$. To prove the

To prove that efficiency and envy-freeness fully determine an allocation when $|N(\succ)| > 0$ for all $\succ \in \mathcal{P}$, first note that there are constants $t_a > 0$, $a \in \Theta$, such that $\mu(i,a) > 0$ imply $t_a = \sum_{b \succsim a} \mu(i,b)$. Indeed, take i, j such that $\mu(i,a), \mu(j,a) > 0$. Consider agent i' who ranks a first and otherwise ranks objects as agent i does; no envy and efficiency imply $\mu(i',a) = \sum_{b \succsim_i a} \mu(i,b)$. Similarly, for agent j' who ranks a first and otherwise ranks objects as agent j does, $\mu(j',a) = \sum_{b \succsim_j a} \mu(j,b)$. Finally, no envy further implies that $\mu(i',a) = \mu(j',a)$, proving the claim. We refer to indices t_a as time and say that object a is not exhausted at all times $t < t_a$ and exhausted at all times $t \geq t_a$.[24]

Our assumptions uniquely determine agent-object pairs (i,a) such that $\mu(i,a) > 0$. To see this first note that efficiency and envy-freeness imply that agents i ranking a first get it with positive probability, $\mu(i,a) = t_a > 0$. No envy then implies that $\mu(i,a) > 0$ for all agents who rank a above all objects b that are not exhausted at time t_a. Let us denote the set of such types by $N(a)$. The definition of times t_{a_k} then implies that for $i \in N(a)$ we have

$$\mu(i,a) = t_a - t'_a \tag{4}$$

where t'_a equals 0 if agent i ranks a first and otherwise equals t_b for the object b that is $\succ_i$-worst among objects $\succ_i$ ranks above a. Furthermore, only agents from $N(a)$ get a with positive probability. Indeed, if agent i ranks a below some object b non-exhausted at t_a, and $\mu(i,a) > 0$ then: (i) agent i gets object $b \succ a$ with probability below 1; (ii) consider agent j that ranks a first and b second and observe that agent j gets a with positive probability because $t_b > t_a > 0$. Equation (4) implies that j gets b with positive probability. Now (i) and (ii) imply that agents i and j have a profitable trade, a contradiction showing that $\mu(i,a) > 0$ iff $i \in N(a)$.

weak version of the theorem it is then enough to show that $\mu(\succ, a)$ for all $\succ \in \mathcal{P}$ and $a \in \Theta$ are uniquely determined by efficiency and envy-freeness.

[24] This choice of terminology is motivated by parallels to the "eating-in-time" terminology of the Probabilistic Serial mechanism of Bogomolnaia and Moulin (2001).

To conclude the proof let us rename the objects so that

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

The efficiency then implies that

$$t_{a_1} = \min\left\{\frac{|a_1|}{|N(a_1)|}, 1\right\}, \quad t_{a_2} = \min\left\{t_{a_1} + \frac{|a_2| - t_{a_1}\left|\{i \in N(a_2) \mid \text{ i ranks } a_2 \text{ first}\}\right|}{|N(a_2)|}, 1\right\},$$

and, proceeding in this way by induction, we can pin down all values t_a, thus uniquely determining allocation μ.

5 Comments on assumptions

We study regular mechanisms.[25] However, the gist of our results remain valid for all mechanisms, not only regular ones. If we drop the regularity assumption, the analogue of Theorems 1 and 2 remain true for full-support sequences of preference profiles defined as follows: a sequence of preference-profiles $\succ_{N_q}$ has *full support* if there 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. Full support holds true uniformly on a class of sequences if they have full support with the same δ and $\bar{q}$. Full support implies that, as q grows, any preference ranking is represented by a non-vanishing fraction of agents.[26]

The proofs of Theorems 1 and 2 show that they are true for full-support sequences of profiles whether or not the mechanism is regular, and we then employ regularity

[25] For a discussion of regularity in the large market literature see, for instance, Champsaur and Laroque (1982).

[26] In a continuum economy, the counterpart of full support says that every ordering is represented with positive probability; in other words the distribution of orderings has full support. Our observations on full-support profiles remain valid if the assumption of non-vanishing representation is imposed only for rankings of objects $\succ$ in which all non-null objects are acceptable. Note also that in the case of Theorem 1 we are relaxing the assumption that the uniform randomizations are regular, while maintaining the weak regularity of the base mechanism over which we randomize.

to extend the results to any preference profiles. Regularity is thus a dispensable assumption because full-support sequences are typical. Indeed, full-support sequences have asymptotically full measure in the following sense: a set $\mathcal{S}$ of sequences of preference profiles is *asymptotically full-measure* if for every $\epsilon > 0$ there exists a sequence of sets $\mathcal{S}_q \subset \mathcal{P}_q$ of preference profiles 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 have asymptotically full-measure.*

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 $\mathcal{S}_q = \mathcal{S}_q^{\delta(\epsilon)}$. □

Another assumption we can relax is symmetry. While the heuristic argument for Theorem 2 made substantial use of symmetry, the formal argument from Appendix B does not, and this argument remains valid when Theorem 2 is strengthened by relaxing the symmetry assumption to its asymptotic counterpart.[27]

6 Conclusion

Theorem 1 shows that – in ordinal settings – many known mechanisms are asymptotically efficient, and Theorem 2 establishes that effectively there is only one way to allocate objects in an efficient, symmetric, and strategy-proof way. Thus, in large markets, the choice among ordinal mechanisms needs to be based on criteria other

[27] Given a sequence of preference profiles $\succ_{N_q}$, a sequence of random allocations μ_q is asymptotically symmetric if

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

A mechanism is *asymptotically symmetric* if the convergence is uniform across preference profiles. In the auxiliary Theorem 1, we can also relax the assumption that the randomization is uniform and replace it by its asymptotic counterpart.

than efficiency or fairness. This has important implications for ongoing efforts to construct new ordinal mechanisms.

A Proof of Theorem 1

Consider a sequence of preference profiles $(\succ_N)_{N=\{1\},\{1,2\},\ldots}$. The key step in our proof is the following.

Lemma 1. *Let θ be a preference ranking, let a be an object, and let ϕ be a uniform randomization over weakly regular deterministic mechanism ψ. Let $(\succ_N)_{N=\{1\},\{1,2\},\ldots}$ be a sequence of preference profiles such that every preference ranking is represented by at least a fraction δ of all agents at all preference profiles in the sequence. If there is an agent with preference ranking θ who obtains a under ϕ with probability at least ϵ at all $\succ_N$, then the share of permutations σ at which some agent of type θ is allocated object a under ψ^σ tends to 1 as $|N| \to \infty$.*

Proof of the lemma. Let $A_N \subseteq \Sigma_N$ be the set of permutations σ such that no agent of type θ obtains a at ψ^σ. Let $B_N(\delta) \subseteq \Sigma_N$ be the set of permutations $\sigma \in \Sigma_N$ such that at least the fraction δ of agents in N obtains a under ψ^σ. For $\delta < \frac{\epsilon}{2}$ sufficiently small, we can assume that $B_N(\delta)$ contains at least the fraction δ^2 of all permutations. Indeed, suppose not. Then, for some large set of agents N, at most for fraction δ^2 of the permutations of ψ contains more than fraction δ of agents obtaining a, and at all other permutations at most fraction δ of agents obtain a. Thus the average probability an agent of type θ_i obtains a is bounded above by $\frac{\delta^2}{\delta} + \delta(1-\delta^2) \leq 2\delta$. By symmetry of the randomization, all these agents have the same probability of obtaining a, and this probability is at least ϵ, a contradiction as $\epsilon > 2\delta$.

The reminder of the proof builds on the following result due to Maurey:[28]

[28] In addition to Maurey's original paper, one may find an exposition of his result in many surveys of concentration inequalities.

Lemma 2. *(Maurey, 1979) If $f : \Sigma_N \to \mathbb{R}$ is Lipschitz with constant L with respect to the distance $d(\sigma, \sigma') = \frac{1}{|N|} |\{i \in N : \sigma(i) \neq \sigma'(i)\}|$, then:*

$$\left| \left\{ \sigma \in \Sigma_N : \left| f(\sigma) - \frac{\sum_{\tau \in \Sigma_N} f(\tau)}{|\Sigma_N|} \right| \geq t \right\} \right| \leq 2e^{-\frac{t^2 n}{16L^2}} |\Sigma_N| .$$

Let now $f(\sigma)$ be $\frac{1}{|N|}$ times the smallest number of transpositions needed to obtain $\sigma \in \Sigma_N$ from some $\tau \in A_N$; the function f is Lipschitz with respect to the distance d.

Since ψ_N is regular, so is each ψ_N^σ. By regularity, there is some fixed $\tilde{\delta} > 0$ such that for a fraction δ of agents to obtain a different allocation, at least a fraction $\tilde{\epsilon}$ of agents needs to change their allocation. Thus, in any $\psi_N^\sigma(\succ)$ for the fraction δ of agents (all of them of of type θ_i) to change their allocations from something different than a to a, we need at least fraction $\tilde{\epsilon}$ of agents to change their reports. Since every transposition in σ can be implemented by a simultaneous change of reports by two agents, we conclude that we need at least $\frac{\tilde{\delta}}{2}|N|$ transpositions to move from any permutation in A to any permutation in B. We can conclude that $\mathbb{E}f = \frac{\sum_{\tau \in \Sigma_N} f(\tau)}{|\Sigma_N|} \geq \frac{\tilde{\delta}}{2}\delta > 0$. Since f equals zero on A_N, we conclude that all $\sigma \in A_N$ belong to $\left\{ \sigma \in \Sigma_N : \left| f(\sigma) - \frac{\sum_{\tau \in \Sigma_N} f(\tau)}{|\Sigma_N|} \right| \geq t \right\}$ for $t = \frac{\tilde{\delta}}{2}\delta$. By Maurey's lemma, A_N can contain at most $2e^{-\frac{t^2 n}{16L^2}} |\Sigma_N|$ permutations. Thus, as $n = |N| \to \infty$, the fraction of permutations in A_N vanishes, contradicting the assumption that this ratio is at least δ, and hence proving the lemma. QED

To finish the proof of Theorem 1, suppose that ϵ-efficiency fails for the sequence of preference profiles $\succ_N$. Thus, either condition (i) or (ii) fails on an infinite subsequence, and we may assume that it fails at all N considered. The analyses of these two cases are similar; let us consider the slightly more complex case when (i) fails. Then, at each N, there is a cycle of c agents $i_1, ..., i_c$ that have a profitable swap of size ϵ, that is $\mu(i_k, a_k) > \epsilon$ and $a_{k+1} \succ_{i_k} a_k$ for all $k = 1, ..., c$. Notice that the uniformity of randomization implies that the agents with the same preference type obtain the same probability shares. The resulting symmetry of allocations allows us to assume that $c \leq |\Theta|$ as otherwise we could shorten the cycle.

First, consider a sequence of preference profiles such that each preference type is represented by at least some fraction δ of agents. Because there are at most $|\Theta|$ agents in the cycle, the lemma we just proved implies that the share of permutations σ at which all types θ_k have a representative who is allocated a_k tends to 1 , and in particular, there exists a permutation σ and agents $j_k \in \theta_k$ such that each j_k is allocated a_k under ψ^{σ}. But, this implies that ψ^{σ}, and hence ψ, are not Pareto efficient, a contradiction.

Second, consider the general case. The regularity of sequence of mechanisms ψ_N implies that there is some δ such changing the preferences of less than fraction δ of agents changes the distributions over objects the remaining agents obtain by at most $\frac{\epsilon}{2}$. By changing the preferences of at most fraction δ of agents, none of them in the cycle $i_1, ..., i_c$, we thus obtain a sequence of preference profiles $\succ'_N$ at which cycles of agents $i_1, ..., i_c$ have a beneficial swaps of size $\frac{\epsilon}{2}$. However, we just proved the asymptotic efficiency of preference profile sequences such as $\succ'_N$, a contradiction completing the proof of the Theorem. QED

B Proof of Theorem 2

In the proof we will use the ingenious construction of the Probabilistic Serial mechanism of Bogomolnaia and Moulin (2001); this mechanism is both efficient and envy-free. 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 shares 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.[29]

Before proving Theorem 2, let us also define asymptotic envy-freeness. First, given an $\epsilon > 0$ and preference profile $\succ_{N_q}$, let us say that an allocation μ is ϵ-envy free if

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

Given a sequence of preference profiles $\succ_{N_q}$, we say that a sequence of allocations

[29] 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).

μ_q is asymptotically envy free if there are positive $\epsilon_q \to 0$ such that μ_q is ϵ_q-envy free. We say that a sequence of mechanisms ϕ_q is asymptotically envy-free if there are positive $\epsilon(q) \overset{q\to\infty}{\to} 0$ such that allocations $\phi_q\left(\succ_{N_q}\right)$ are ϵ_q-envy free for all $\succ_{N_q}$.

Lemma 3. *Fix a full-support sequence of preference profiles. If a sequence of random allocations μ_q is asymptotically efficient and asymptotically envy-free, then it coincides asymptotically with the allocations of Probabilistic Serial. Moreover, if a sequence of mechanisms ϕ_q is asymptotically efficient, symmetric, and asymptotically envy-free, and a class of profile sequences $\mathcal{S}$ has uniformly full support, then ϕ_q converges to Probabilistic Serial uniformly on this class, that is*

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

Proof. Fix any sequence of full-support preference profiles $\succ_{N_q}$ and an asymptotically efficient and asymptotically envy-free sequence of allocations μ_q. Asymptotic 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 $\frac{\epsilon}{M^2}$-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 i's allocation $\frac{\epsilon}{M^2}$-first order stochastically dominates allocations of any other agent j in agent i's 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 $\mu = \mu_q$. We also drop the preference argument when referring to the allocation of Probabilistic Serial ψ.

To prove the lemma it is enough to show that

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

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

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 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 a 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.

Step 1. Inequalities (5) are satisfied at $t = T$ because the mapping $t \mapsto \psi^t(i,a)$ is continuous and the inequalities are satisfied for $t \in [0,T)$. In particular, the cutoff time T belongs to $[0,1)$.

Step 2. At time T of the eating procedure, agent i must be eating from b. 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 that 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 $\sum_{a' \succsim_i b} \mu(i,a') \geq \sum_{a' \succsim_i b} \psi^T(i,a') - \frac{\epsilon}{M} = \sum_{a' \succsim_i b} \psi^t(i,a') - \frac{\epsilon}{M}$ 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 $\frac{\epsilon}{M^2}$-efficiency of μ implies that there must be an 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 the lack of swaps of size $\frac{\epsilon}{M^2}$ (the consequence of $\frac{\epsilon}{M^2}$-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 lemma.

An examination of the above argument shows that the choice of ϵ, M, and $\bar{q}$ can be made uniformly for $\mu_q = \phi_q \left(\succ_{N_q}\right)$ on a class of preference profile sequences with uniformly full-support, proving the second part of the lemma.

Proof of Theorem 2. First note that it is enough to prove the result assuming that ϕ'_q are Probabilistic Serial. Second, notice that regularity, symmetry, and asymptotic strategy-proofness imply that $\phi_q \left(\succ_{N_q}\right)$ is asymptotically envy free for all $\succ_{N_q}$. The above lemma shows a uniform convergence of $\phi \left(\succ_{N_q}\right)$ to Probabilistic Serial 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. Let us then fix any sequence of profiles $\succ_{N_q}$, and derive the convergence for $\succ_{N_q}$ from the regularity of ϕ_q and the convergence for asymptotically full-support sequences of profiles $\succ'_{N_q}$ such that (1) is satisfied. The convergence is uniform across preference profiles, proving Theorem 2. QED

References

Abdulkadiroğlu, A., P. A. Pathak, and A. E. Roth (2009): "Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *American Economic Review*, 99, 1954–1978.

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

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.

Azevedo, E. and J. Leshno (2013): "A Supply and Demand Framework for Two-sided Matching Markets," .

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

Carroll, G. (2010): "A General Equivalence Theorem for Allocation of Indivisible Objects," MIT, Working paper.

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.

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

Echenique, F., S. Lee, B. Yenmez, and M. Shum (2013): "The Revealed Preference Theory of Stable and Extremal Stable Matchings," *Econometrica*, 81, 153–171.

Erdil, A. (2011): "Strategy-Proof Stochastic Assignment," .

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

Featherstone, C. R. (2011): "Rank Efficiency: Investigating a Widespread Ordinal Welfare Criterion," .

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.

——— (1978): "Straightforwardness of Game Forms with Lotteries as Outcomes," *Econometrica*, 46, 595–614.

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

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.

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.

Lee, S. (2013): "Incentive Compatibility of Large Centralized Matching Markets," .

Makowski, L., J. M. Ostroy, and U. Segal (1999): "Efficient Incentive Compatible Economies Are Perfectly Competitive," *Journal of Economic Theory*, 85, 169–225.

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

Maurey, B. (1979): "Construction de Suites Symetriques," *Comptes Rendus de l'Académie des Sciences*, 288, 679–681.

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.

Pycia, M. and M. U. Ünver (2016): "Incentive Compatible Allocation and Exchange of Discrete Resources," *Theoretical Economicse*, (forthcoming).

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

Roberts, J. D. and A. Postlewaite (1976): "The incentives for price-taking behavior in large exchange economies," *Econometrica*, 115–127.

Roth, A. E. and U. Rothblum (1999): "Truncation Strategies in Matching Markets: In Search of Advice for Participants," *Econometrica*, 67, 21–43.