

**PAYOFF EQUIVALENCE OF EFFICIENT MECHANISMS
IN LARGE MATCHING MARKETS**

By

Yeon-Koo Che and Olivier Tercieux

July 2015

COWLES FOUNDATION DISCUSSION PAPER NO. 2015



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.yale.edu/>

Payoff Equivalence of Efficient Mechanisms in Large Matching Markets*

Yeon-Koo Che[†]

Olivier Tercieux[‡]

July 30, 2015

Abstract

We study Pareto efficient mechanisms in matching markets when the number of agents is large and individual preferences are randomly drawn from a class of distributions, allowing for both common and idiosyncratic shocks. We show that, as the market grows large, all Pareto efficient mechanisms—including top trading cycles, serial dictatorship, and their randomized variants—are uniformly asymptotically payoff equivalent “up to the renaming of agents,” yielding the utilitarian upper bound in the limit. This result implies that, when the conditions of our model are met, policy makers need not discriminate among Pareto efficient mechanisms based on the aggregate payoff distribution of participants.

JEL Classification Numbers: C70, D47, D61, D63.

Keywords: Large matching markets, Pareto efficiency, Payoff equivalence

1 Introduction

Assigning indivisible resources without monetary transfers is an important problem in modern market design; applications range from allocation of public housing, public school seats,

*We are grateful to Ludovic Lelièvre, Charles Maurin and Xingye Wu for their research assistance and to Eduardo Azevedo, Itai Ashlagi, Julien Combe, Olivier Compte, Tadashi Hashimoto, Yash Kanoria, Fuhito Kojima, Scott Kominers, SangMok Lee, Bobak Pakzad-Hurson, Debraj Ray, Rajiv Sethi, seminar participants at Columbia, NYU, Wisconsin, Maryland, University College London and Yale, and attendees of the PSE Market Design conference, Warwick Micro theory conference, KAEA Conference, NYC Market Design Workshop, and WCU Market Design conference at Yonsei University for helpful comments.

[†]Department of Economics, Columbia University, USA. Email: yeonkooche@gmail.com.

[‡]Department of Economics, Paris School of Economics, France. Email: tercieux@pse.ens.fr.

employment contracts, and branch postings to the assignment of human organs to transplant patients. A basic desideratum in designing such a market is Pareto efficiency. If a mechanism is not Pareto efficient, a surplus can be generated and distributed in a way that benefits (at least weakly) all participants, suggesting clear room for improvement.

In a centralized matching market, achieving Pareto efficiency is often not difficult. A number of mechanisms are known to produce efficiency, often satisfying additional desirable properties in terms of incentives and (ex ante) fairness.¹ Rather the challenge is often *which mechanism to choose among many Pareto efficient ones*.

This issue is important because alternative Pareto efficient mechanisms often treat individual participants differently (often dramatically so). For instance, in serial dictatorship, individuals are allowed to choose objects, one at a time, from a set according to some serial order; the first dictator (the first individual in the serial order) could very well select the object that is regarded by all as the best, while the last dictator (the final individual in the order) may have to settle for the object regarded by all as the worst. Without monetary transfers to compensate for the loss borne by the latter, the apparent conflict of interests leaves little hope for consensus in terms of selecting from alternative Pareto efficient mechanisms. Ideally, the selection must be based on some measure of aggregate welfare of participants. For instance, if one Pareto efficient mechanism yields a significantly higher utilitarian welfare level or a much more equal payoff distribution than others, that would constitute an important rationale for favoring such a mechanism.

A similar concern arises when the designer has additional policy considerations, such as “affirmative treatment” of some target group (say, identified based on their socio-economic background). Specifically, the designer may select an efficient mechanism based on such additional goals (for instance, by elevating the target agents’ serial orders in serial dictatorship). Any such adjustment will obviously impact the welfare of the participants at the individual level, but do these adjustment impact the total welfare of the agents or their aggregate payoff distribution? If so, how? If accommodating additional social objectives were to entail a significant loss of utilitarian welfare or to produce a significant distributive impact, this would call into question the merit of the policy intervention. This type of policy consideration requires one to evaluate the payoff consequences of alternative Pareto efficient mechanisms.

¹Mechanisms such as (deterministic or random) serial dictatorship produce efficiency without regard to existing property rights; top trading cycles mechanisms achieve efficiency by allowing agents to trade pre-existing property rights or priorities (Shapley and Scarf (1974), Abdulkadiroglu and Sonmez (2003)). These mechanisms satisfy strategyproofness and can satisfy ex ante “equal treatment of equals” when the serial order and initial ownership are drawn at random. Efficiency may also be achieved by allowing agents to purchase the objects using “fake money” in an artificial market, as envisioned by Hylland and Zeckhauser (1979).

Unfortunately, little is known about how alternative mechanisms perform in this regard.

The purpose of the current paper is to fill this gap while providing useful insights on practical market design in the process. To make progress, we add some structure to the model. First, we consider markets that are “large” in terms of the number of participants as well as in the number of object types. Large markets are clearly relevant in many settings. For instance, in the US National Resident Matching Program, approximately 20,000 medical applicants participate in filling the positions of 3,000 to 4,000 programs each year. In New York City, approximately 90,000 students apply to over 700 high school programs each year. Second, we assume that the agents’ preferences are randomly generated according to some reasonable distribution. Specifically, we consider a model in which each agent’s utility from an object depends on a common component (i.e., a portion that does not vary across agents) and on an idiosyncratic component that is drawn at random independently (and thus varies across agents). Studying the limit properties of a large market with preferences randomly generated in this way provides a framework for answering our questions.

Our main finding is that all Pareto efficient mechanisms yield aggregate payoffs, or utilitarian welfare, that converge uniformly to the same limit—more precisely, the utilitarian optimum—as the economy grows large (in the sense described above). This result implies that in large economies, alternative efficient mechanisms become virtually indistinguishable in terms of the aggregate payoff distribution of the participants. In other words, agents’ payoffs are asymptotically equivalent across different efficient mechanisms, up to the “renaming” of agents. This result implies that there is no reason to favor one efficient mechanism over another. From a policy perspective, this means that a Pareto efficient allocation favoring or prioritizing a certain group of individuals would not significantly harm utilitarian welfare or significantly alter the distribution of payoffs in a large market.

Importantly, our equivalence holds in terms of the distribution of ordinal ranks enjoyed by the participants, making the result robust to the particular specification of cardinal utilities assumed. The result is also robust to the institutional details that the efficient mechanisms must accommodate (e.g., in terms of the property rights and priorities enjoyed by some agents), which makes the result readily applicable to many practical market design problems. To test the applicability of our results to realistic market settings, we compare alternative efficient mechanisms using simulated data as well as field data from the New York City school choice program. As will be seen, the comparison supports our equivalence result.

The present paper contributes to several strands of literature. First, our equivalence result is closely related to and complements the equivalence result among a class of efficient mechanisms established by [Abdulkadiroglu and Sönmez \(1998\)](#) and its extensions ([Pathak and Sethuraman \(2011\)](#), [Carroll \(2014\)](#), [Lee and Sethuraman \(2011\)](#) and [Bade \(2014\)](#)). As

we shall discuss in detail, this equivalence result holds only in the absence of prior ownership or priority rights, i.e., when participants are treated *ex ante* symmetrically via fair lotteries. By contrast, our result holds for arbitrary priority or ownership structures, as long as the market is sufficiently large. This generality makes our equivalence result applicable to many real-world settings where there are often priority considerations for participants (as was the case with New York City school choice program).

Second, our result contributes to the literature on large matching markets, particularly those with a large number of object types and random preferences; see [Immorlica and Mahdian \(2005\)](#), [Kojima and Pathak \(2008\)](#), [Lee \(2014\)](#), [Knuth \(1997\)](#), [Pittel \(1989\)](#), [Lee and Yariv \(2014\)](#), [Ashlagi, Kanoria, and Leshno \(2013\)](#) and [Che and Tercieux \(2015\)](#).² The first three papers are largely concerned with the incentive issues arising from the deferred acceptance (henceforth, DA) mechanisms of [Gale and Shapley \(1962\)](#). The last five papers are concerned with the ranks of the partners achieved by agents on two sides of a market under DA. We focus on the payoffs enjoyed by agents under *efficient* mechanisms. In a preference environment closer to ours, [Lee and Yariv \(2014\)](#) show that *stable* mechanisms also yield the utilitarian upper bound in a large market limit. As we show below via simulation and data analysis, efficient mechanisms tend to converge much faster than do stable mechanisms, and the magnitude of the difference can be considerable for realistic market sizes. Further, our convergence result is robust, holding even for unbalanced markets, whereas their result does not, as implied by [Ashlagi, Kanoria, and Leshno \(2013\)](#). Most importantly, the uniform equivalence of efficient mechanisms (possibly employing different priority structures) established in the current paper is quite striking and has no analogues in the existing literature.

Methodologically, the current paper utilizes a framework developed in random graph theory; see [Dawande, Keskinocak, Swaminathan, and Tayur \(2001\)](#), for instance. In particular, the proof method is similar to the way [Lee \(2014\)](#) exploits the implications of the stability of agents on two sides in a suitably defined random graph; as will be clear, our method exploits

²Another strand of literature studying large matching markets considers a large number of agents matched with a finite number of object types (or firms/schools) on the other side; see [Abdulkadiroglu, Che, and Yasuda \(2015\)](#), [Che and Kojima \(2010\)](#), [Kojima and Manea \(2010\)](#), [Azevedo and Leshno \(2011\)](#), [Azevedo and Hatfield \(2012\)](#) and [Che, Kim, and Kojima \(2013\)](#), among others. The assumption of a finite number of object types enables one to use a continuum economy as a limit benchmark in these models. At the same time, this feature makes the analysis and the resulting insights quite different. The two strands of large matching market models capture issues that are relevant in different real-world settings and are thus complementary. The latter model is more appropriate for situations in which there are a relatively small number of institutions, each with a large number of positions to fill. School choice in some districts, such as Boston Public Schools, could be a suitable application because only a handful of schools enroll hundreds of students each. The former model is descriptive of settings in which there are numerous participants on both sides of the market. Medical matching and school choice in some districts, such as the New York Public Schools, would fit this description.

the implications of Pareto efficiency for an appropriately constructed random graph.

2 Set-up

We consider a model in which a finite set of agents are assigned a finite set of objects, at most one object for each agent. Because our analysis will examine the limit of a sequence of finite economies, it is convenient to index the economy by its size n . An n -**economy** $E^n = (I^n, O^n)$ consists of **agents** I^n and **object types** O^n , where $|I^n| = n$. For much of the analysis, we shall suppress the superscript n for notational simplicity.

The object types can be interpreted as schools or housing units. Each object type o has $q_o \geq 1$ **copies** or **quotas**. Because our model allows for $q_o = 1$ for all $o \in O^n$, one-to-one matching is a special case of our model. We assume that total quantity is $Q^n = \sum_{o \in O^n} q_o = n$. In addition, we assume that the number of copies of each object is uniformly bounded, i.e., there is $\bar{q} \geq 1$ such that $q_o \leq \bar{q}$ for all $o \in O^n$ and all n . The assumption that $Q^n = n$ is made only for convenience—as long as it grows at order n , our results will hold. In particular, as will be clear, our argument will hold even in cases in which the market is unbalanced. Similarly, the assumption that the number of copies of each object is uniformly bounded is not necessary as long as it grows at a sufficiently low rate.³

Throughout, we shall consider a general class of random preferences that allow for a positive correlation among agents on the objects. Specifically, each agent $i \in I^n$ receives the following utility from obtaining object type $o \in O^n$:

$$U_i(o) = U(u_o, \xi_{i,o}),$$

where u_o is a common value, and the *idiosyncratic shock* $\xi_{i,o}$ is a random variable drawn independently and identically from $[0, 1]$ according to a uniform distribution.⁴

We further assume that the function $U(\cdot, \cdot)$ takes values in \mathbb{R}_+ , is strictly increasing in the common values and strictly increasing and continuous in the idiosyncratic shock. The utility of remaining unmatched is assumed to be 0 so that all agents find all objects acceptable.⁵ The symmetry of $U(\cdot, \cdot)$ can be seen as a normalization in the scaling of individual utilities, which also implies a normalization invoked by many authors that the highest possible utility and

³As will be clear from footnote 23, we can allow $\bar{q} = O(n/\log(n))$.

⁴This assumption entails no loss of generality as long as the distribution of idiosyncratic shocks is atomless and bounded, as one can always focus on the quantile corresponding to the agent's idiosyncratic shock as a normalization and then redefine the payoff function as a function of this normalized shock.

⁵This feature does not play a crucial role in our result, which hold as long as a linear fraction of objects are acceptable to all agents.

the lowest possible utility are identical across all agents. The symmetry assumption serves to normalize individuals' cardinal utilities and thus discipline interpersonal comparison of utilities. Further, as we will discuss in Section 5, our core findings are robust to the rescaling of individual utilities. We assume that the agents' common value for object type $o \in O$, u_o takes an arbitrary value in $[0, 1]$ in an n -economy, and its population distribution is given by a cumulative distribution function (CDF):

$$X^n(u) = \frac{\sum_{o \in O^n: u_o \leq u} q_o}{n},$$

interpreted as the fraction of the *objects* whose common value is less than or equal to u , and by

$$Y^n(u) = \frac{|\{o \in O^n | u_o \leq u\}|}{n},$$

interpreted as the fraction of the *object types* whose common value is no greater than u . Because $q_o \geq 1$ for each $o \in O$, it follows that $X_n(\cdot) \geq Y_n(\cdot)$.

We assume that these CDFs converge to well-defined limits, X and Y , in the Lévy metric. To be precise, for any two distributions F and G , consider their distance measured in the Lévy metric:

$$L(F, G) := \inf \{ \delta > 0 | F(z - \delta) - \delta \leq G(z) \leq F(z + \delta) + \delta, \forall z \in \mathbb{R}_+ \}.$$

According to this measure, any two distributions will be regarded as being close to each other as long as they are uniformly close at all points of continuity.⁶ It follows that the limit distributions X and Y are nondecreasing and satisfy $X(0) = 0$, $X(1) = 1$ and $X(\cdot) - Y(\cdot) \geq 0$. We assume that X has (at most) finite jumps. We allow X and Y to be fairly general, allowing for atoms.

Several special cases of this model are of interest. The first is a **finite-tier model**. In this model, the object types are partitioned into finite tiers, $\{O_1^n, \dots, O_K^n\}$, where $\cup_{k \in K} O_k^n = O^n$ and $O_k^n \cap O_j^n = \emptyset$. (With a slight abuse of notation, the largest cardinality K also denotes the set of indexes.) In this model, the CDFs X^n and Y^n are step functions with finite steps. This model offers a good approximation of situations in which the objects have clear tiers, such as schools classified into different categories or regions or houses existing in clearly distinguishable tiers. A further special case is when $K = 1$ in which the support of the common value is degenerate and agents' ordinal preferences are drawn iid uniformly. [Knuth \(1997\)](#), [Pittel \(1989\)](#), and [Ashlagi, Kanoria, and Leshno \(2013\)](#) employ such a model.

Another special case is the **full-support model** in which the limit distribution Y is strictly increasing in its support. This model is very similar to [Lee \(2014\)](#) and [Lee and Yariv](#)

⁶Here, convergence of CDFs in the Lévy metric is equivalent to weak convergence.

(2014), who also consider random preferences that consist of common and idiosyncratic terms. One difference is that their framework assumes that the common component of the payoff is also drawn randomly from a positive interval. Our model assumes common values to be arbitrary, but with a full-support assumption, the values can be interpreted as realizations of random draws (drawn according to the CDF Y). Viewed in this way, the full-support model is comparable to Lee (2014)'s, except that the current model also allows for atoms in the distribution of Y .

Unless otherwise specified, we are referring to a general model that nests these two as special cases. Fix an n -economy. We shall consider a class of matching mechanisms that are Pareto efficient. A **matching** μ in an n -economy is a mapping $\mu : I \rightarrow O \cup \{\emptyset\}$ such that $|\mu^{-1}(o)| \leq q_o$ for all $o \in O$, with the interpretation that agent i with $\mu(i) = \emptyset$ is unmatched. Let M denote the set of all matchings. All these objects depend on n , although their dependence is suppressed for notational simplicity.

In practice, the matching chosen by the designer will depend on the realized preferences of the agents as well as on other features of the economy. For instance, if the objects O are institutions or individuals, their preferences over their matching partners will typically impact which matching will arise. Alternatively, one may wish the matching to respect the existing rights that individuals may have over the objects; for instance, if the objects are housing, some units may be occupied by existing tenants who have priority over these units. Likewise, a school choice matching may favor students whose siblings already attend the school or those living nearby. Some of these factors may depend on the features not captured by their idiosyncratic component. Our model is completely general in this regard.

Specifically, we collect all assignment-relevant variables, call its generic realization a “state,” and denote it by $\omega = (\{\xi_{i,o}\}_{i \in I, o \in O}, \theta)$, where $\{\xi_{i,o}\}_{i \in I, o \in O}$ is the realized profile of the idiosyncratic component of payoffs, and θ is the realization of all other variables that influence the matching, and let Ω denote the set of all possible states. We make no assumption on how θ is drawn and how its realized value affects the outcome. The generality on the θ contrasts the current model with the others, many of which tend to impose a particular random structure on the agents’ priorities (or objects’ preferences). For instance, Ashlagi, Kanoria, and Leshno (2013) assumes that the preferences are iid, and Lee (2014) and Lee and Yariv (2014) assume that the preferences consist of random common shocks with full support and iid idiosyncratic shocks. Our results do not require any such assumptions on the object side.

A **matching mechanism** is a function that maps from a state in Ω to a matching in M . With a slight abuse of notation, we shall use $\mu = \{\mu_\omega(i)\}_{\omega \in \Omega, i \in I}$ to denote a matching mechanism, which selects a matching $\mu_\omega(\cdot)$ in state ω . Let \mathcal{M} denote the set of all match-

ing mechanisms. For convenience, we shall often suppress the dependence of the matching mechanism on ω .

A matching $\mu \in M$ is Pareto efficient if there is no other matching $\mu' \in M$ such that $U_i(\mu'(i)) \geq U_i(\mu(i))$ for all $i \in I$ and $U_i(\mu'(i)) > U_i(\mu(i))$ for some $i \in I$. A matching mechanism $\mu \in \mathcal{M}$ is Pareto efficient if, for each state $\omega \in \Omega$, the matching it induces, i.e., $\mu_\omega(\cdot)$, is Pareto efficient. Let \mathcal{M}_n^* denote the set of all Pareto efficient mechanisms in the n -economy.

3 Payoff Equivalence of Pareto Efficient Mechanisms

We first define an upper bound for the utilitarian welfare—the highest possible level of total surplus that can be realized under any matching mechanism. To this end, suppose that every agent is assigned an object and realizes the highest possible idiosyncratic payoff. Because the common values of the objects are distributed according to X^n , the resulting (normalized) utilitarian welfare is $\int_0^1 U(u, 1) dX^n(u)$. This obviously yields the upper bound for the utilitarian welfare in the n -economy. We consider its limit, the **limit utilitarian upper bound**:

$$U^* := \int_0^1 U(u, 1) dX(u).$$

The payoff distribution of an economy, whether a finite n -economy or its limit, can be represented by a distribution function, i.e., a nondecreasing right-continuous function F mapping from $[0, U(1, 1)]$ to $[0, 1]$. $F(z)$ is interpreted as the fraction of agents who realize payoffs no greater than z . We let F^μ denote the payoff distribution induced by mechanism μ .

We are now in a position to state our main theorem.

THEOREM 1. *Let F^* be the distribution of payoffs attaining the limit utilitarian upper bound U^* , and recall \mathcal{M}_n^* is the set of Pareto efficient mechanisms in the n -economy. Then,*

$$\sup_{\mu^n \in \mathcal{M}_n^*} L(F^{\mu^n}, F^*) \xrightarrow{p} 0.^7$$

In words, the theorem states that the distance (in the Lévy metric) between a payoff distribution resulting from *every* Pareto efficient mechanism and that of the utilitarian upper

⁷We say $Z_n \xrightarrow{p} z$, or Z_n **converges in probability** to z , where both Z_n and z are real-value random variables, if for any $\epsilon > 0, \delta > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$, we have

$$\Pr\{|Z_n - z| > \epsilon\} < \delta.$$

bound vanishes uniformly in probability as $n \rightarrow \infty$. More precisely, assuming the distribution F^* is continuous, the statement is as follows. Fix any $\epsilon, \delta > 0$. Then, with a probability of at least $1 - \delta$, the proportion of agents enjoying any payoff u or higher under any Pareto efficient mechanism is within ϵ of the proportion of agents enjoying a payoff of u or higher under the utilitarian upper bound for sufficiently large n . It is remarkable that the rate of convergence is “uniform” with respect to the entire class of Pareto efficient mechanisms.

The following corollary is immediate:

COROLLARY 1.

$$\inf_{\mu^n \in \mathcal{M}^*} \frac{\sum_{i \in I} U_i(\mu^n(i))}{|I|} \xrightarrow{p} U^*.$$

The theorem also implies that alternative Pareto efficient mechanisms become payoff equivalent uniformly as the market grows in size—that is, “up to the renaming of the agents”:

COROLLARY 2.

$$\sup_{\mu^n, \tilde{\mu}^n \in \mathcal{M}_n^*} L(F^{\mu^n}, F^{\tilde{\mu}^n}) \xrightarrow{p} 0.$$

These results suggest that, as long as agents are ex ante symmetric in their preferences, there is little ground to favor one Pareto efficient mechanism over another in terms of the total welfare of participants or aggregate payoff distribution, at least in a large economy. This has important implications for market design. As we already noted, designers often face extra constraints arising from the existing rights or priorities of some participants over some objects, or there may be a need to treat some target group of participants affirmatively. In addition, there is a concern that accommodating such constraints or needs may sacrifice utilitarian welfare or adversely impact the aggregate distribution of payoffs. Our result implies that accommodating such constraints does not entail any significant loss in these terms in a large economy, as long as Pareto efficiency is maintained.

4 Sketch of the Proof

Here, we sketch the proof of Theorem 1, which is contained in Appendix A. For our current purposes, assume $X(\cdot)$ is degenerate with a single common value u^0 and that $X(\cdot) = Y(\cdot)$. In other words, the agents have only idiosyncratic payoffs, and the matching is one-to-one. As will be seen in Appendix A, the same proof argument works for the general case (with some care).

To begin, fix an arbitrary Pareto efficient mechanism $\tilde{\mu}$. We first invoke the fact that any Pareto efficient matching can be implemented by a serial dictatorship⁸ with a suitably-chosen serial order (see [Abdulkadiroglu and Sönmez \(1998\)](#)). Let \tilde{f} be the serial order, namely, a function that maps each agent in I to his serial order in $\{1, \dots, n\}$ that implements $\tilde{\mu}$ under a serial dictatorship. Because $\tilde{\mu}$ induces a Pareto efficient matching that depends on the state, the required serial order \tilde{f} is random.

Next, for arbitrarily small $\epsilon, \delta > 0$, define the random set:

$$\bar{I} := \left\{ i \in I \mid U_i(\tilde{\mu}(i)) \leq U(u^0, 1 - \epsilon) \text{ and } \tilde{f}(i) \leq (1 - \delta)n \right\}.$$

The set \bar{I} consists of agents who are within $1 - \delta$ top percentile in terms of their serial order \tilde{f} but fail to achieve payoff ϵ -close to the highest possible payoff.⁹ Because $\epsilon, \delta > 0$ are arbitrary, for the proof, it will suffice to show that

$$\frac{|\bar{I}|}{n} \xrightarrow{p} 0. \tag{1}$$

To prove this, we exploit a result in random graph theory. It is thus worth introducing the relevant random graph model. A **bipartite graph** G consists of vertices $V_1 \cup V_2$ and edges $E \subset V_1 \times V_2$ across V_1 and V_2 (with no possible edges within vertices in each side). An **independent set** is $\bar{V}_1 \times \bar{V}_2$, where $\bar{V}_1 \subseteq V_1$ and $\bar{V}_2 \subseteq V_2$ for which no element in $\bar{V}_1 \times \bar{V}_2$ is an edge of G . A random bipartite graph $B = (V_1 \cup V_2, p)$, $p \in (0, 1)$, is a bipartite graph with vertices $V_1 \cup V_2$ in which each pair $(v_1, v_2) \in V_1 \times V_2$ is linked by an edge with probability p (independently of edges created for all other pairs). The following result provides the crucial step for our result.¹⁰

LEMMA 1 ([Dawande, Keskinocak, Swaminathan, and Tayur \(2001\)](#)). *Consider a random bipartite graph $B = (V_1 \cup V_2, p)$ where $0 < p < 1$ is a constant and for each $i \in \{1, 2\}$ and $|V_1| = n$ and $|V_2| = m = O(n)$. There is $\kappa > 0$,*

$$\Pr \left[\exists \text{ an independent set } \bar{V}_1 \times \bar{V}_2 \text{ with } \min\{|\bar{V}_1|, |\bar{V}_2|\} \geq \kappa \ln(n) \right] \rightarrow 0 \text{ as } n \rightarrow \infty.$$

⁸A serial dictatorship mechanism specifies an order over individuals and then lets the first individual—according to the specified ordering—receive his favorite object; the next individual receives his favorite item of the remaining objects, etc.

⁹Strictly speaking, we should focus on individuals receiving payoffs lower than $U(u^0, 1) - \epsilon$. However, given that the utility functions are continuous, there is little loss in focusing our attention on agents receiving less than $U(u^0, 1 - \epsilon)$. This point will be made clear in the proof.

¹⁰ The original statement by [Dawande, Keskinocak, Swaminathan, and Tayur \(2001\)](#) assumes that $|V_1| = |V_2| = n$. It is easily verified that their arguments also apply under our more general assumptions.

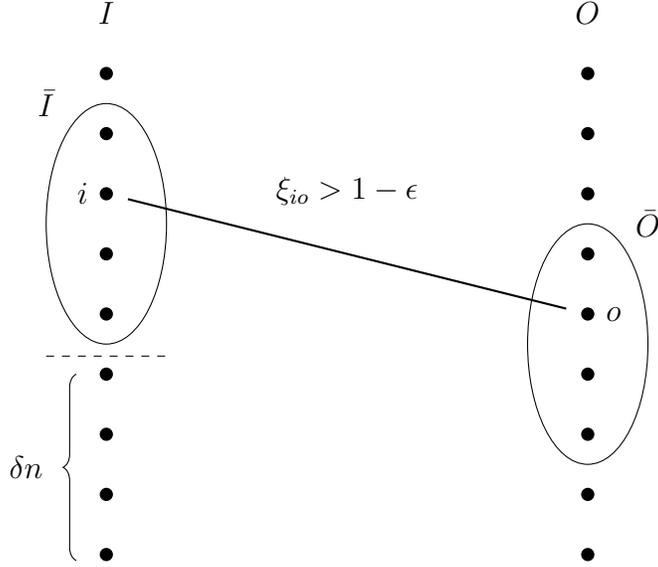


Figure 1: Illustration of a random graph and sets \bar{I} and \bar{O}

This result implies that with probability converging to 1, for every independent set, at least one side of that set vanishes in relative size as $n \rightarrow \infty$.

To prove our result, it therefore suffices to show that \bar{I} forms a vanishing side of an independent set in an appropriately-defined random graph. Consider a random bipartite graph consisting of I on one side and O on the other side where an edge is created between $i \in I$ and $o \in O$ if and only if $\xi_{i,o} > 1 - \epsilon$. Let

$$\bar{O} := \left\{ o \in O \mid \tilde{f}(\tilde{\mu}(o)) \geq (1 - \delta)n \right\}$$

be the (random) set of objects that are assigned to the agents who are at the bottom δ percentile in terms of the serial order \tilde{f} . See Figure 1 for a graphical representation of the construction, where the set I is ordered according to (a realization of) the serial order.

The key observation is that the (random) subgraph $\bar{I} \times \bar{O}$ is an independent set.

To see this, suppose the contrary—there is an edge between an agent $i \in \bar{I}$ and an object $o \in \bar{O}$ in some state ω (as illustrated in Figure 1). By construction of \bar{I} , agent $i \in \bar{I}$ must realize less than $1 - \epsilon$ of idiosyncratic payoff from $\tilde{\mu}_\omega(i)$. However, the fact that there is an edge between i and o means that i would gain more than $1 - \epsilon$ in idiosyncratic payoff from o . Thus, agent i must prefer o to his match $\tilde{\mu}_\omega(i)$. However, the fact that $o \in \bar{O}$ means that o is not yet claimed and is thus available when agent i (who is within top $1 - \delta$ of serial order \tilde{f}_ω) picks $\tilde{\mu}_\omega(i)$. This is a contradiction, proving that $\bar{I} \times \bar{O}$ is an independent set.

Next we observe that $|\bar{O}| \geq \delta n$, meaning that \bar{O} never vanishes in probability. Lemma 1 then implies that set \bar{I} must vanish in probability. Importantly, this result applies uniformly to all mechanisms in \mathcal{M}^* : If we define the sets $\bar{I}(\tilde{\mu})$ and $\bar{O}(\tilde{\mu})$ for each $\tilde{\mu} \in \mathcal{M}^*$ as above, for each $\tilde{\mu} \in \mathcal{M}^*$, $\bar{I}(\tilde{\mu}) \times \bar{O}(\tilde{\mu})$ forms an independent set of the *same* random graph! This explains the uniform convergence.

REMARK 1. If the mechanism $\tilde{\mu}$ were a serial dictatorship with a “deterministic” serial order f , a simple direct argument proves the result. First, let us note that we can think of each agent as drawing his preferences “along the algorithm,” i.e., he draws his preferences for the stage when it is his turn to make a choice. Obviously, the distribution of i ’s preferences is not affected by the choices of agents ahead of that agent in the serial order. Fix any arbitrary $\epsilon, \delta > 0$ and let E_i be the event that at agent i ’s turn to make a choice there remains at least one object o such that $U_i(o) \geq U(u^0, 1 - \epsilon)$. Then, all agents except those in the bottom δ -percentile serial orders enjoy idiosyncratic payoffs ϵ -close to the upper bound with probability:

$$\begin{aligned} & \Pr\{U_i(\tilde{\mu}(i)) \geq U(u^0, 1 - \epsilon) \text{ for all } i \text{ with } f(i) < (1 - \delta)n\} \\ & \geq \Pr\{\cap_{i \in I: f(i) < (1 - \delta)n} E_i\} \\ & = 1 - \Pr\{\cup_{i \in I: f(i) < (1 - \delta)n} E_i^c\} \\ & \geq 1 - (1 - \delta)n(1 - \epsilon)^{\delta n} \rightarrow 1 \text{ as } n \rightarrow \infty. \end{aligned}$$

However, this argument does not work for an arbitrary Pareto efficient mechanism. For a general Pareto efficient mechanism, the serial order implementing the mechanism is, in general, not independent of the agents’ preferences (which is required in the last inequality of the above string). Our general proof using random graph theory avoids this difficulty.

5 Implications and Robustness

We explore several implications of our findings. First, the utilitarian efficiency of Pareto efficient allocations stated in Corollary 1 is remarkable and surprising given the fact that monetary transfers are not allowed. One interpretation is that a large market makes utilities virtually transferable by creating “rich” opportunities for agents to trade on idiosyncratic payoffs. In other words, objects that are uniformly valued by the participants can be transferred from one set of agents to another set without entailing much loss in terms of the idiosyncratic payoffs. In this sense, a large market can act as a “substitute” for monetary transfers. Such a result, while plausible, is neither obvious nor universally true. This strong payoff equivalence does not extend when the welfare of both sides are relevant. DA is Pareto efficient across

both sides of the market—i.e., taking the objects as welfare-relevant entities—but, as shown in [Che and Tercieux \(2015\)](#), it does not attain the highest total payoffs across the market, and a different matching (which is Pareto inefficient) yields a higher utilitarian welfare.

Second, the current equivalence result is reminiscent of a similar equivalence result obtained by [Abdulkadiroglu and Sönmez \(1998\)](#) between two well-known mechanisms, random serial dictatorship and TTC with random ownership, and of the large market equivalence result obtained by [Che and Kojima \(2010\)](#) between random serial dictatorship and a probabilistic serial mechanism as well as their extensions ([Pathak and Sethuraman \(2011\)](#), [Carroll \(2014\)](#), [Lee and Sethuraman \(2011\)](#) and [Bade \(2014\)](#)). While these results consider arbitrary preferences on the agents, they assume ex ante symmetric random priorities with respect to the objects. By contrast, our equivalence result does not impose any structure on the priorities on the object side, allowing them to be arbitrary, but it does impose a certain structure on the agents’ preferences (to consist of common values and iid idiosyncratic shocks). Our result also holds only in the limit as the number of agents and objects becomes large (with the number of object types bounded or growing at a slower rate), whereas the equivalence result by [Abdulkadiroglu and Sönmez \(1998\)](#) and others holds for any finite economy. Ultimately, the current result complements the existing focus on (ex post) Pareto efficient mechanisms that treat agents symmetrically in terms of tie-breaking or ex ante assignment of property rights. One implication of our result is that the fairness achieved by the symmetric treatment of agents does not entail any significant welfare loss, neither in terms of utilitarian welfare nor of the payoff distribution among agents.

Third, a similar result is known to hold under a stable matching in some large market settings. [Knuth \(1997\)](#) and [Pittel \(1989\)](#), among others, have shown that if the agents’ ordinal preferences are drawn iid uniformly and the market is balanced, the aggregate welfare of the agents under a stable matching approaches the utilitarian upper bound.¹¹ [Lee and Yariv \(2014\)](#) show a similar result in a large balanced market in which the agents on both sides have random preferences that consist of common and idiosyncratic shocks, and the common shock has full support. It turns out that these conditions are important. [Che and Tercieux \(2015\)](#) show that this result does not hold if the common value component of objects’ preferences (or agents’ priorities with the objects) do not have full support. Even with preference drawn iid uniformly on both sides, [Ashlagi, Kanoria, and Leshno \(2013\)](#) find that if the market is unbalanced, the agents on the long side compete excessively for agents (or objects) on the short side in a stable matching, which entails a significant welfare loss for the former agents.

¹¹Specifically, they show that the (preference) rank of the objects agents enjoy converges to $\log(n)$ on average, which means that the idiosyncratic payoffs will be in the order of $1 - \log(n)/n$ on average. Because the common values are degenerate in their environment, the result follows.

The general takeaway from these papers is that the outcome of DA is likely to be bounded away from the utilitarian upper bound when there is competition among agents for desirable objects. This suggests that even with full-support distributions and balanced markets, this should be observed to some extent as long as agents’ preferences are sufficiently positively correlated. This is indeed what we observe in our simulated data. In addition, we observe a similar phenomenon in our analysis of field data from New York City, which suggests that in real markets, DA entails a significant efficiency loss compared to Pareto efficient mechanisms. See Section 6.

By contrast, our result is robust to market imbalance and to a general distribution of agent priorities of objects. Specifically, our model makes no assumption on the latter, and agents’ preferences allow for market imbalances. Suppose that the objects in our model exist in two tiers with the common value u_1 of tier-1 objects exceeding the common value u_2 of tier-2 objects by a significant margin. This will create an unbalanced market as far as tier-1 objects are concerned, and a similar competition by agents to obtain tier-1 objects will arise under DA. Indeed, [Che and Tercieux \(2015\)](#) show that a stable matching mechanism is not even approximately efficient in this situation. By contrast, the imbalance and associated competition do not entail any significant welfare loss for the agents in a Pareto efficient mechanism; specifically, competition for scarce resources does not entail significant losses for those who are fortunate enough to be assigned them if the assignment is Pareto efficient.

Let us now come back to the structure we impose on agent preferences. We assume that utilities are symmetric, i.e., the function U is not player dependent. While this serves as a useful normalization, one may wonder whether our results depend on the particular scaling of individual utilities. Scaling of individual utilities would matter at a superficial level, for instance, if we scale up the utilities of some group of agents and keep the others the same, efficient mechanisms that treat them differently would result in different aggregate utilities.¹² However, there are important senses in which our insights are robust to the rescaling of individual utilities.

First, regardless of how individual utilities are (re)scaled, in all efficient mechanisms, the fraction of agents who receive an arbitrarily high *idiosyncratic payoff* converges to 1 as the market grows large. Indeed, a large market still creates “rich” opportunities for agents to trade on idiosyncratic payoffs, and thus, as the market gets large, the set of payoffs associated with Pareto-efficient mechanisms still has a specific structure.

Second, on more normative grounds, given that we are using the utilitarian criterion

¹²Imagine two serial dictatorship mechanisms that yield systematically different serial orders based on group membership.

(which assigns equal weight to individual welfare), the symmetry assumption essentially reflects the idea that we are not willing to differentiate among agents.

Finally and most importantly, many institutions assess the performance of a mechanism based on the distributions of “relative ranks” that agents achieve, namely, the ordinal rank of the object obtained in each agent’s ranking divided by the number of objects allocated.¹³ Such a measure is clearly invariant to the scaling of the individual utilities. Our large market equivalence result implies equivalence in this measure: as the market grows large, the distribution of relative ranks becomes identical across all Pareto efficient mechanisms.¹⁴

6 Applicability of the Findings

What do our results imply for realistic markets? We study this question with two sets of data: (1) simulated data and (2) choice data supplied by the New York City Department of Education. Specifically, we shall compare across alternative efficient mechanisms using DA as the benchmark. These exercises enable us to examine our theory in an environment that involves: (1) a range of different market sizes, including relatively small market sizes, (2) a broad range of preference and priority distributions, and (3) realistic quotas for object types (e.g., schools).

6.1 Simulation

We first simulate our full-support model. Specifically, we assume that each agent’s preference is given by $U(u_o, \xi_{io}) = u_o + \xi_{io}$, where u_o and ξ_{io} are both generated according to a uniform distribution from $[0, 1]$. Some of the mechanisms we consider also depend on the agents’ priorities or object preferences for agents, which we assume are given by $V(v_i, \eta_{io}) = v_i + \eta_{io}$, where v_i is the objects’ common preference for agent i and η_{io} is object o ’s idiosyncratic preference for agent i . We also assume that both v_i and η_{io} are drawn according to a uniform distribution from $[0, 1]$. Note that this preference/priority structure is a special case of our model because priorities are arbitrary in our model.

We consider four different mechanisms: priority-based top trading cycles (PBTTC), two versions of serial dictatorship (SD min and SD max), and Gale-Shapley’s DA algorithm.

¹³Featherstone (2015) is the first paper studying “rank efficient mechanisms,” i.e., mechanisms with rank distributions that are not stochastically dominated.

¹⁴To be more precise, let us consider a cumulative distribution of normalized ranks where for each u_o , a fraction of individuals $X(u_o)$ have a relative rank smaller than $1 - (X(u_o) + \int_{u_o}^1 \Pr\{U(u, \xi_{io}) < U(u_o, 1)\}dX(u))$. The cumulative distribution of normalized ranks of any Pareto efficient mechanism converges to this.

PBTTC, proposed by [Abdulkadiroglu and Sonmez \(2003\)](#), executes Pareto-improving trades among applicants in multiple rounds.¹⁵ For the two versions of SD, we randomly draw a serial order of agents 100 times and select the best- and worst-performing assignment in terms of utilitarian welfare. The purpose is to see the range of variations in utilitarian welfare associated with different Pareto efficient mechanisms. Finally, DA, proposed by [Gale and Shapley \(1962\)](#), attains stable matching in a two-sided (i.e., agent-agent) matching environment.¹⁶ In the case of agent-object matching, the mechanism does not attain Pareto efficiency, but we focus on the mechanism as a benchmark. DA is interesting in light of the recent result by [Lee and Yariv \(2014\)](#) that it becomes approximately efficient in a large market given the full support environment (on both sides) assumed here. Given this result, a natural question arises as to how the convergence rates differ between DA and efficient mechanisms for realistic market sizes.

Figure 2 shows the utilitarian welfare performance of the alternative mechanisms averaged over 50 iterations of the preference and priority draws $\{(u_o, \xi_{io}, v_i, \eta_{io})\}_{i,o}$, measured in terms of the aggregate values of ξ_{io} 's accruing to the agents.¹⁷ (Recall that the common shock u_o affects all agents identically and thus can be subtracted without loss.)

All mechanisms, including DA, perform well and improve as the market grows large. These results are in line with our main findings (in particular, Corollary 1) and with [Lee and Yariv \(2014\)](#). What is not implied by these results and can be learned from the simulations is the speed of convergence and the uniformity across mechanisms. For market size $n = 2,000$, all efficient mechanisms—even the SD min—realize more than 98% of the highest possible surplus. The uniformity of convergences across efficient mechanisms is indeed remarkable, for they become essentially indistinguishable for any $n \geq 1000$. This shows that our uniform

¹⁵In each round, each applicant points to his most preferred object among those available, each object points to the applicant with the highest priority for the object among those available, and the applicants associated with a cycle (which must exist due to the finite number of participants) are assigned the objects they point to and exit the market, along with the seats they are assigned. The same process is repeated with the remaining participants, until all participants are exhausted.

¹⁶DA runs in multiple rounds. In each round, each agent proposes to the most preferred object that has not yet rejected that agent, and the object chooses the most preferred agent among those that have proposed up to that point. The algorithm iterates to the next round as long as some agents are rejected and have acceptable objects remaining. When the process stops, the tentative assignment at the last round becomes final.

¹⁷For the two versions of SD, for each preference draw, $\{(u_o, \xi_{io})\}_{i,o}$, we run 100 serial dictatorship mechanisms based on 100 randomly drawn serial orders of the agents, and we select from among them the best- and worst-performing assignments in terms of utilitarian welfare. (Hence, the priority draws $\{(v_i, \eta_{io})\}_{i,o}$ play no role for SDs.) We then average the utilitarian welfare performances of the selected SDs over 50 random preference drawings.

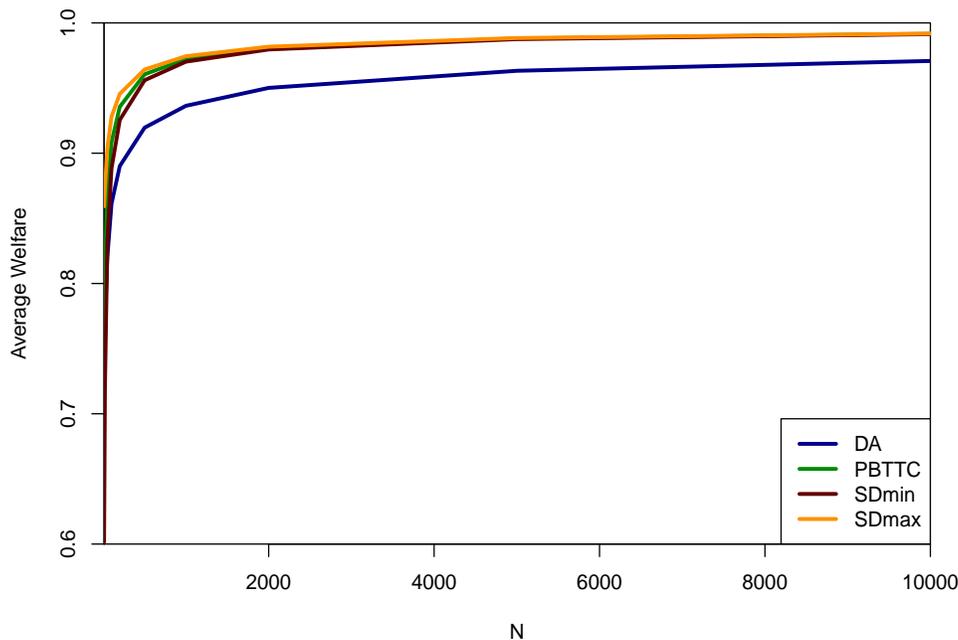


Figure 2: Welfare Comparison across Mechanisms

asymptotic equivalence result “kicks in” for even moderately sized markets. Note that the current equivalence result holds across different priorities and therefore is distinct from—not implied by—the equivalence result by [Abdulkadiroglu and Sönmez \(1998\)](#) and its extensions. Indeed, for very small n , the difference between SD min and SD max is appreciable, suggesting that the classical equivalence result is not applicable here. But even for a modestly large n , the difference between the two mechanisms vanishes. Finally, although DA performs well in utilitarian efficiency, there is a clear difference relative to the efficient mechanisms. For $n = 2,000$, there is a 3% point difference relative to SD min, and the tangible difference of at least 2% point remains even for $n = 10,000$. This suggests that the convergence rates differ nontrivially between efficient mechanisms and DA.

We next study the robustness of our comparison. Figure 3 draws the preferences and priorities in the same way as the baseline case but assumes that there are twice as many agents as the objects. The performance of all efficient mechanisms in terms of utilitarian welfare and their equivalence remain largely unchanged. The performance of the DA has gotten significantly worse. This result is in line with [Ashlagi, Kanoria, and Leshno \(2013\)](#)

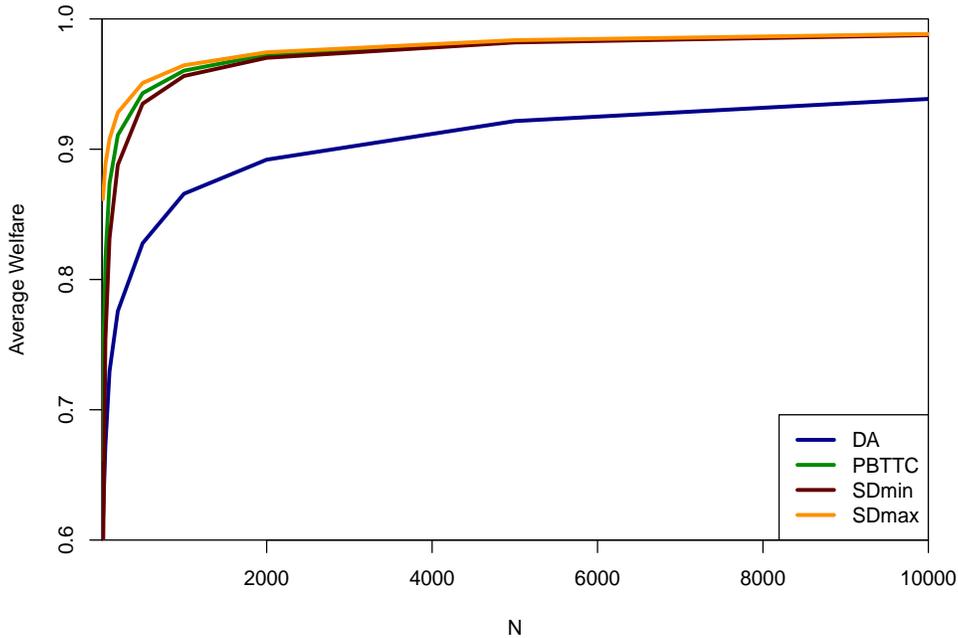


Figure 3: Welfare Comparison across Mechanisms

and [Che and Tercieux \(2015\)](#), which suggests that a significant welfare loss would result from DA, as agents compete excessively for scarce resources.¹⁸

Such competition arises even in a balanced market if agents' preferences are highly correlated or equivalently if agents' common preferences shocks become much more important than idiosyncratic shocks. Figure 4 depicts such a scenario: the model is the same as above except that agents' common preference shock u_o is now uniform on $[0, 3]$. Again, the performances of the efficient mechanisms are quite robust to this change in the model, but the DA performance is substantially worse than it is under the baseline model. In fact, the welfare loss is comparable to that under the market imbalance. In other words, competition by agents to realize high common value objects requires a significant sacrifice of their idiosyncratic payoffs under a stable mechanism but not under efficient mechanisms.

¹⁸Although these papers show that the agents' payoffs under DA no longer converge to the utilitarian upper bound given market imbalance, this result is obtained in a simpler environment and does not directly apply to the simulated setting.

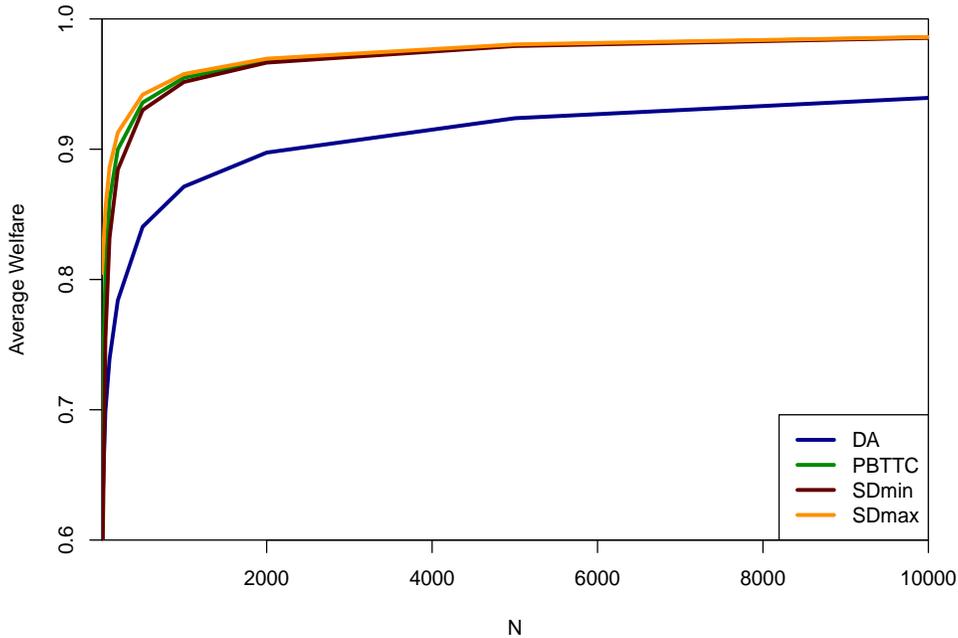


Figure 4: Welfare Comparison across Mechanisms

6.2 Calibration Based on NYC School Choice Data

In New York City, approximately 90,000 students (mostly in the 8th grade) are assigned to over 700 public high school programs through an annual centralized matching process. We focus on the main round (round 2) of assignment. In that round, each student submits a rank ordered list (ROL) of up to 12 programs, and each program ranks applicants who listed it on their ROLs, according to its priority criteria, which depend on the types of the program.¹⁹ The priorities are coarse for many programs, and a single (uniform) lottery is used to break ties for all programs. Given the ROLs and priorities, a student-proposing DA algorithm is used to generate an assignment.

We use the 2009-2010 choice data to calibrate the assignments that would arise under alternative Pareto efficient mechanisms: priority-based top trading cycles (PBTTTC), random serial dictatorship (RSD), and priority-based serial dictatorship (PBSD). RSD is a serial

¹⁹The programs are categorized into several types in terms of admissions method: screened, limited un-screened, un-screened, ed-op, zoned and audition. See [Che and Tercieux \(2015\)](#) for a detailed description of the data and the institutional details.

dictatorship in which the applicants’ serial orders are determined at random. PBSO is a serial dictatorship in which the applicants’ serial orders are based on the average priority rank of the agents.²⁰ We study the distribution of the ranks enjoyed by the participants under these mechanisms. These mechanisms differ in the way that they treat the participants, so there is no a priori expectation for the relation of the distributions.

Before proceeding, a couple of remarks are in order. First, following the existing literature, we assume that the observed ROLs of the applicants represent their truthful preference ranking of top programs. This assumption is not entirely innocuous because the strategyproofness of DA does not apply when the applicants’ ROLs are truncated (see [Haeringer and Klijn \(2009\)](#)). Nevertheless, approximately 80% of the participants did not fill their ROLs, suggesting that truncation was not a binding constraint (see [Abdulkadiroglu, Pathak, and Roth \(2009\)](#) and [Abdulkadiroglu, Agarwal, and Pathak \(2015\)](#) for the same assumption). Second, under the current DA algorithm, programs do not specify priorities for students unless they rank them in their ROLs. To calibrate PBTTTC, we assume that programs assign lower priorities to students who do not rank them than to those who do rank them. This does not pose a serious problem for our purposes because Pareto efficiency of PBTTTC does not depend on this assumption.²¹

Table 1 and Figure 5 present the distribution of preference ranks achieved by the applicants under alternative efficient mechanisms, using DA as a control mechanism. They exhibit striking resemblance in the rank distribution across alternative Pareto efficient mechanisms. While the source of the resemblance is not immediately clear, the noticeable difference these mechanisms exhibit relative to the DA outcome suggests that the resemblance is not driven by the special nature of the underlying preferences.²² The difference also suggests that even though DA may perform well in terms of efficiency in a large market, as suggested by [Lee and Yariv \(2014\)](#), the convergence rate may be slow enough to entail an appreciable difference

²⁰Specifically, we compute a student’s normalized priority rank (after a lottery draw) at each program to which he/she applied. We then take the average across all programs that a student applied to. This yields the average priority rank for each student. We then form a serial order of the students based on this order (in the ascending order).

²¹If PBTTTC were introduced, the priority information would be collected for all programs, regardless of whether a student lists them on their ROLs. So, the outcome would not be the same. At the same time, our finding below (as does our theoretical result) suggests that the difference in the distribution of preference ranks achieved by the applicants will not be significant.

²²For example, if all applicants submit the same ROL of programs, the rank distribution would be identical across all assignments, and there would be no difference between DA assignment and efficient assignments. Likewise, if there are no conflicts of interests, again, all agents will be assigned to their top choice under both an efficient mechanism and DA. The difference between the DA and efficient mechanisms suggests that neither scenario holds here.

Table 1: Rankings achieved under 4 different algorithms

| | DA | PBTTC | RSD | PBSD |
|-------------|------------------|------------------|------------------|------------------|
| #1 | 35200.87 (53.67) | 38090.25 (36.58) | 37657.08 (51.03) | 37784.47 (53.24) |
| #2 | 14006.8 (53.01) | 13256.99 (46.08) | 13307.03 (61.02) | 13265.9 (53.12) |
| #3 | 8168.72 (41.93) | 7157.68 (41.24) | 7103.74 (51.6) | 7190.1 (52.82) |
| #4 | 4882.67 (35.32) | 4025.68 (31.32) | 3983.21 (41.23) | 4100.22 (39.46) |
| #5 | 2976.64 (29.75) | 2382.62 (25.83) | 2374.91 (34.81) | 2484.59 (35.67) |
| #6 | 1716.71 (20.81) | 1347.35 (21.12) | 1343.15 (25.16) | 1433.89 (28.34) |
| #7 | 996.4 (19.27) | 746.87 (17.07) | 789.61 (20.66) | 851.97 (22.41) |
| #8 | 592.47 (16.46) | 443.39 (12.92) | 471.48 (15.55) | 511.6 (16.47) |
| #9 | 336.74 (11.8) | 265.24 (11) | 287.17 (14.51) | 304.6 (14.02) |
| #10 | 190.38 (9) | 150.22 (8.26) | 174.88 (11.12) | 186.05 (11.11) |
| #11 | 122.17 (6.34) | 100.79 (6.02) | 112.97 (9.16) | 126.66 (9.23) |
| #12 | 66.22 (5.41) | 54.22 (4.9) | 69.58 (7.64) | 74.36 (7.01) |
| #Unassigned | 8458.21 (29.31) | 9693.7 (31.31) | 10040.19 (47.17) | 9400.59 (37.87) |

Note: We ran 100 iterations of each algorithm with independent draws of lotteries and focused on the average performance of each algorithm, including DA. Standard errors are in parentheses.

compared with efficient mechanisms.

Recall also that the programs have intrinsic priorities in the data, and the alternative Pareto efficient mechanisms differ in the way that the programs' priority information is used to generate the assignment. Hence, the resemblance across alternative Pareto efficient mechanisms cannot be explained by the equivalence result of [Abdulkadiroglu and Sönmez \(1998\)](#) or its extensions by recent authors. These authors focus on an environment in which programs have no intrinsic priorities and find equivalence of efficient mechanisms that treat agents randomly in an *ex ante* symmetric manner. Importantly, equivalence holds only in the *ex ante* sense (in terms of the lotteries the agents receive), and it does not imply that a similar rank distribution would result from different Pareto efficient mechanisms using different priority systems.

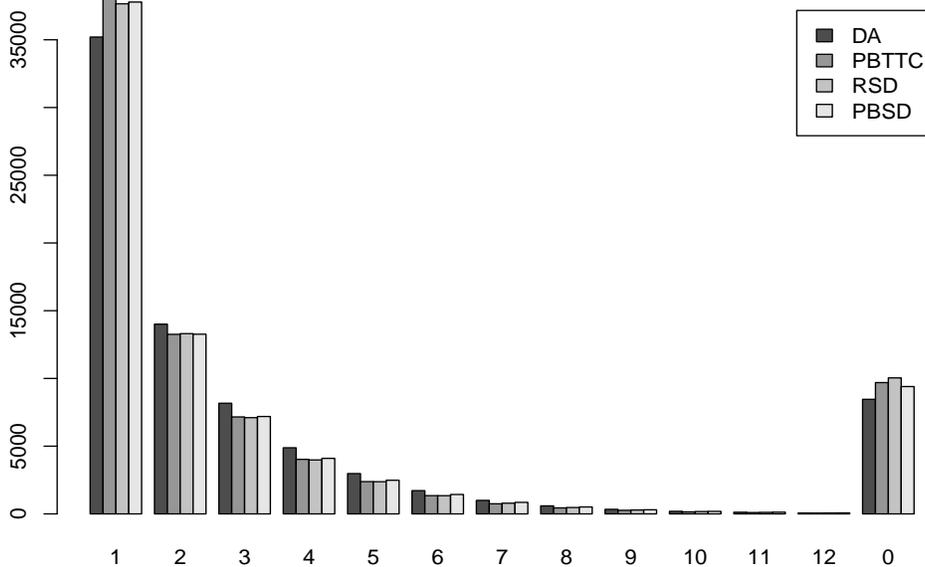


Figure 5: Rank Distribution under Alternative Mechanisms (averaged across 100 iterations).

A Appendix: Proof of Theorem 1

A.1 Preliminaries

For an n -economy and for each $u \in [0, 1]$, let $O_{\geq u}^n := \{o \in O^n | u_o \geq u\}$ and $O_{\leq u}^n := \{o \in O^n | u_o \leq u\}$ denote the set of object types that yield the common value no less than u and the set of object types that yield the common value no greater than u , respectively. The numbers of objects with types in $O_{\geq u}^n$ and $O_{\leq u}^n$ are respectively denoted by $Q_{\geq u}^n$ and $Q_{\leq u}^n$. For notational simplicity, we shall suppress the dependence of these sets on n , with the exception of X^n .

Now, consider any Pareto efficient mechanism $\mu \in \mathcal{M}^*$. By a well known result (e.g., [Abdulkadiroglu and Sönmez \(1998\)](#)), any Pareto efficient matching can be equivalently implemented by a serial dictatorship mechanism with a suitably chosen serial order. Let SD^{f_μ} be the serial dictatorship mechanism where for each state ω a serial order $f_\mu(\omega) : I \rightarrow I$, a bijective mapping, is chosen so as to implement $\mu_\omega(\cdot)$. That is, for each state $\omega \in \Omega$, the serial order f_μ is chosen so that $SD_\omega^{f_\mu(\omega)}(i) = \mu_\omega(i)$ for each $i \in I$. Since the matching μ

arising from the mechanism depends on the random state ω , the serial order f implementing μ is a random variable. In the sequel, we shall study a Pareto efficient matching mechanism μ via the associated SD^f . To avoid clutter, we shall now suppress the dependence of f on μ .

Given an n -economy, for any Pareto efficient mechanism μ and the associated serial order f , let

$$I_{\geq u}(\mu) := \{i \in I \mid f(i) \leq Q_{\geq u}\}.$$

be the set of agents who have a serial order within the total supply of objects whose common values are at least u (in the equivalent serial dictatorship implementation). For any ϵ , the set

$$\bar{I}_{\geq u}^\epsilon(\mu) = \{i \in I_{\geq u}(\mu) \mid U_i(SD^f(i)) \leq U(u, 1 - \epsilon)\},$$

consists of the agents who realize payoff no greater than $U(u, 1 - \epsilon)$ while having a serial order within $Q_{\geq u}$. The following lemma will be crucial for the main result.

LEMMA 2. For any $\epsilon, \gamma > 0$,

$$\Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \text{ such that } \frac{|\bar{I}_{\geq u}^\epsilon(\mu)|}{|I|} \geq \gamma \right] \rightarrow 0$$

as $n \rightarrow \infty$.

PROOF. Fix any $\epsilon > 0$ and $\gamma > 0$. We first build a random bipartite graph on $I \cup O$ where an edge (i, o) is added if and only if $\xi_{i,o} > 1 - \epsilon$.

Now choose any $\delta \in (0, 1)$. For each $\mu \in \mathcal{M}^*$ and u , define random sets $I_{\geq u}^\delta(\mu) := \{i \in I \mid f(i) \leq (1 - \delta)Q_{\geq u}\}$, $\bar{I}_{\geq u}^{\epsilon, \delta}(\mu) := \{i \in I_{\geq u}^\delta(\mu) \mid U_i(SD^f(i)) \leq U(u, 1 - \epsilon)\}$, and

$$\bar{O}_{\geq u}^\delta(\mu) := \{o \in O_{\geq u} \mid \exists i \in \mu^{-1}(o) \text{ s.t. } f(i) > (1 - \delta)Q_{\geq u}\},$$

which consists of object types in $O_{\geq u}$ assigned to the agents with serial order worse than $(1 - \delta)Q_{\geq u}$.

We argue that the set $\bar{I}_{\geq u}^{\epsilon, \delta}(\mu) \times \bar{O}_{\geq u}^\delta(\mu)$ must be an independent set of the random bipartite graph on $I \cup O$. To prove this, suppose otherwise. Then, there would exist an edge $(i, o) \in \bar{I}_{\geq u}^{\epsilon, \delta}(\mu) \times \bar{O}_{\geq u}^\delta(\mu)$. Then,

$$U_i(o) > U(u, 1 - \epsilon) \geq U_i(SD^f(i))$$

where the strict inequality holds since $\xi_{i,o} > 1 - \epsilon$ (i.e., (i, o) is an edge), $o \in O_{\geq u}$, and since $U(\cdot, \cdot)$ is monotonic (in particular strictly increasing in idiosyncratic component). The weak inequality holds because $i \in \bar{I}_{\geq u}^{\epsilon, \delta}$. In addition, we must have

$$f(i) \leq (1 - \delta)Q_{\geq u} < f(i'), \text{ for some } i' \in \mu^{-1}(o)$$

where the first inequality comes from the fact that $i \in \bar{I}_{\geq u}^{\epsilon, \delta}(\mu) \subset I_{\geq u}^{\delta}(\mu)$ while the second from the fact that $o \in \bar{O}_{\geq u}^{\delta}(\mu)$. Thus, this means that when i becomes the dictator under SD^f , object o is still available. But $U_i(o) > U_i(SD^f(i))$ means that i chooses an object worse than o , which yields a contradiction.

Since $\bar{I}_{\geq u}^{\epsilon, \delta}(\mu) \times \bar{O}_{\geq u}^{\delta}(\mu)$ is an independent set for each $\mu \in \mathcal{M}^*$ and $u \in [0, 1]$ and since $|I| = n$, applying Lemma 1, we have that, for any $\gamma' > 0$:

$$\Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \in [0, 1] \text{ s.t. } \min \left\{ |\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)|, |\bar{O}_{\geq u}^{\delta}(\mu)| \right\} \geq \gamma' n \right] \rightarrow 0 \quad (2)$$

as n goes to infinity.

Fix any $\gamma' > 0$. Recall $|\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)| \leq |I_{\geq u}^{\delta}(\mu)| \leq (1 - \delta)Q_{\geq u}$ and $|\bar{O}_{\geq u}^{\delta}(\mu)| \geq \delta Q_{\geq u} / \bar{q}$. Hence, if $|\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)| \geq \gamma' n$, then we must have $|\bar{O}_{\geq u}^{\delta}(\mu)| \geq \frac{\delta \gamma'}{(1 - \delta) \bar{q}} n$, where recall \bar{q} is the upper bound for the number of copies for each object type.

Hence, as $n \rightarrow \infty$,

$$\begin{aligned} & \Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \text{ s.t. } |\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)| \geq \gamma' n \right] \\ &= \Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \text{ s.t. } |\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)| \geq \gamma' n \text{ and } |\bar{O}_{\geq u}^{\delta}(\mu)| \geq \frac{\delta \gamma'}{(1 - \delta) \bar{q}} n \right] \\ &\leq \Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \text{ s.t. } \min \left\{ |\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)|, |\bar{O}_{\geq u}^{\delta}(\mu)| \right\} \geq \min \left\{ 1, \frac{\delta}{(1 - \delta) \bar{q}} \right\} \gamma' n \right] \\ &\rightarrow 0, \end{aligned}$$

where the convergence follows from (2).²³

Finally, by construction, $|\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)| \geq |\bar{I}_{\geq u}^{\epsilon}(\mu)| - \delta Q_{\geq u} - 1$. Since $Q_{\geq u} \leq n$, we get that

$$\frac{|\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)|}{|I|} \geq \frac{|\bar{I}_{\geq u}^{\epsilon}(\mu)|}{|I|} - \delta - \frac{1}{|I|}$$

for each $\mu \in \mathcal{M}^*$. Hence, it follows that

$$\begin{aligned} & \Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \text{ s.t. } \frac{|\bar{I}_{\geq u}^{\epsilon}(\mu)|}{|I|} \geq \gamma' + \delta \right] \\ &\leq \Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \text{ s.t. } \frac{|\bar{I}_{\geq u}^{\epsilon, \delta}(\mu)|}{|I|} \geq \gamma' + \frac{1}{|I|} \right] \rightarrow 0. \end{aligned}$$

²³ Here we use the assumption that \bar{q} does not increase in n . If \bar{q} increases in n at the rate of $O(n/\log(n))$, then one can check that the lower bound in the above equation is $\omega(\log(n))$. Using Lemma 1, one can show that Lemma 2—and thus Theorem 1—holds.

Set δ and γ' such that $\delta + \gamma' = \gamma$. Then,

$$\Pr \left[\exists \mu \in \mathcal{M}^* \text{ and } u \text{ s.t. } \frac{|\bar{I}_{\geq u}^\epsilon(\mu)|}{|I|} \geq \gamma \right] \rightarrow 0.$$

□

We are now ready to prove Theorem 1.

A.2 Proof of Theorem 1

To prove the statement, we will show that the payoff distributions induced by Pareto efficient mechanisms converge to F^* in the sense defined earlier.

Fix any $\epsilon > 0$. We shall show that, as $n \rightarrow \infty$,

$$\Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z \max \{F^\mu(z - \epsilon) - F^*(z), F^*(z) - F^\mu(z + \epsilon)\} \geq \epsilon \right] \rightarrow 0, \quad (3)$$

where F^* and F^μ are respectively the CDF of the payoff induced by the limit utilitarian upper bound and the CDF of the payoffs induced by mechanism μ in \mathcal{M}^* .

Let

$$J^\mu(z) := \{i \in I \mid U_i(\mu(i)) \leq z\}.$$

denote the set of agents enjoying payoff of at most z under matching mechanism μ . Obviously, $F^\mu(z) = |J^\mu(z)|/n$. Let $u(z)$ be such that $U(u(z), 1) = z$ for each $z \in \hat{Z} := [U(0, 1), U(1, 1)]$. (This is well defined since $U(\cdot, 1)$ is continuous and strictly increasing.) Note that the function $u : \hat{Z} \rightarrow [0, 1]$ so defined is continuous and increasing. Let us fix ϵ' such that for any common value $u \leq u(z) + \epsilon'$ we have $U(u, 1) \leq z + \epsilon$ for each $z \in \hat{Z} := [U(0, 1), U(1, 1)]$. Note that this is well-defined since $U(u(z), 1) = z$ and $U(\cdot, 1)$ is continuous and strictly increasing. Further observe that ϵ' is strictly positive. Clearly, for any $z \in \hat{Z}$, any agent matched with an object having common value no greater than $u(z) + \epsilon'$ must be in $J^\mu(z + \epsilon)$. This means that $|J^\mu(z + \epsilon)| \geq Q_{\leq u(z) + \epsilon'} = X^n(u(z) + \epsilon')n$ for all $\mu \in \mathcal{M}^*$. By definition, for each z , $F^*(z) = X(u(z))$.

Then,

$$\begin{aligned}
& \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z \left(-\frac{|J^\mu(z + \epsilon)|}{|I|} + F^*(z) \right) \geq \epsilon \right] \\
&= \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_{z \in \hat{Z}} \left(-\frac{|J^\mu(z + \epsilon)|}{|I|} + F^*(z) \right) \geq \epsilon \right] \\
&\leq \Pr \left[\sup_{z \in \hat{Z}} (-X^n(u(z) + \epsilon') + X(u(z))) \geq \epsilon \right] \\
&= \Pr \left[\sup_u (-X^n(u + \epsilon') + X(u)) \geq \epsilon \right] \\
&\leq \Pr \left[\sup_u (-X^n(u + \min\{\epsilon, \epsilon'\}) + X(u)) \geq \min\{\epsilon, \epsilon'\} \right] \\
&\rightarrow 0, \tag{4}
\end{aligned}$$

as $n \rightarrow \infty$. The first equality comes from the fact that for any $z < U(0, 1)$, $F^*(z) = 0$; the first inequality is by definition of ϵ' and the convergence follows since $L(X^n, X) \rightarrow 0$ as $n \rightarrow \infty$.

For the next part, recall $I_{\geq u}(\mu) := \{i \in I | f(i) \leq Q_{\geq u}\}$ and $\bar{I}_{\geq u}^\epsilon(\mu) := \{i \in I_{\geq u}(\mu) | U_i(SD^f(i)) \leq U(u, 1 - \epsilon)\}$, where SD^f is the SD rule implementing μ . Let $I_{\leq u}(\mu) := \{i \in I | f(i) \leq Q_{\leq u}\}$. We also extend the function u such that $u(z) = 0$ for any $z \in [U(0, 0), U(0, 1)]$.

Fix ϵ'' such that for any common value $u \geq u(z) - \epsilon''$ we have $z - \epsilon \leq U(u, 1 - \epsilon'')$ for all $z \in [U(0, 0), U(1, 1)]$. Note that this is well-defined since $U(u(z), 1) \geq z$ and U is continuous and strictly increasing in both components. In addition, $\epsilon'' > 0$. Observe that $U(\mu(i)) \leq z - \epsilon$ implies $i \in I_{\leq u(z) - \epsilon''} \cup \bar{I}_{\geq u(z) - \epsilon''}^{\epsilon''}(\mu)$. Hence,

$$\frac{|J^\mu(z - \epsilon)|}{|I|} \leq \frac{|I_{\leq u(z) - \epsilon''}|}{|I|} + \frac{|\bar{I}_{\geq u(z) - \epsilon''}^{\epsilon''}(\mu)|}{|I|}.$$

We obtain

$$\begin{aligned}
& \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z \left(\frac{|J^\mu(z - \epsilon)|}{|I|} - F^*(z) \right) \geq \epsilon \right] \\
& \leq \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z \left(\frac{|I_{\leq u(z) - \epsilon''}|}{|I|} + \frac{|\bar{I}_{\geq u(z) - \epsilon''}^{\epsilon''}(\mu)|}{|I|} - F^*(z) \right) \geq \epsilon \right] \\
& \leq \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z (X^n(u(z) - \epsilon'') - X(u(z))) \geq \frac{\epsilon}{2} \right] + \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z \frac{|\bar{I}_{\geq u(z) - \epsilon''}^{\epsilon''}(\mu)|}{|I|} \geq \frac{\epsilon}{2} \right] \\
& = \Pr \left[\sup_u (X^n(u - \epsilon'') - X(u)) \geq \frac{\epsilon}{2} \right] + \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z \frac{|\bar{I}_{\geq u(z) - \epsilon''}^{\epsilon''}(\mu)|}{|I|} \geq \frac{\epsilon}{2} \right] \\
& \leq \Pr \left[\sup_u \left(X^n(u - \min\{\epsilon'', \frac{\epsilon}{2}\}) - X(u) \right) \geq \min\{\epsilon'', \frac{\epsilon}{2}\} \right] + \Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_u \frac{|\bar{I}_{\geq u}^{\epsilon''}(\mu)|}{|I|} \geq \frac{\epsilon}{2} \right] \\
& \rightarrow 0, \tag{5}
\end{aligned}$$

where the convergence holds by $L(X^n, X) \rightarrow 0$ and Lemma 2.

Combining (4) and (5) and the fact that $F^\mu(z) = |J^\mu(z)|/n$, we conclude that

$$\Pr \left[\sup_{\mu \in \mathcal{M}^*} \sup_z \max\{F^\mu(z - \epsilon) - F^*(z), -F^\mu(z + \epsilon) + F^*(z)\} \geq \epsilon \right] \rightarrow 0,$$

as $n \rightarrow \infty$.

References

- ABDULKADIROGLU, A., N. AGARWAL, AND P. A. PATHAK (2015): “The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match,” MIT, unpublished mimeo.
- ABDULKADIROGLU, A., Y.-K. CHE, AND Y. YASUDA (2015): “Expanding ‘Choice’ in School Choice,” *American Economic Journal: Microeconomics*, 7, 1–42.
- ABDULKADIROGLU, 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(5), 1954–1978.
- ABDULKADIROGLU, A., AND T. SÖNMEZ (1998): “Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems,” *Econometrica*, 66, 689–701.

- ABDULKADIROGLU, A., AND T. SONMEZ (2003): “School Choice: A Mechanism Design Approach,” *American Economic Review*, 93, 729–747.
- ASHLAGI, I., Y. KANORIA, AND J. D. LESHNO (2013): “Unbalanced Random Matching Markets,” MIT, Unpublished mimeo.
- AZEVEDO, E. M., AND J. W. HATFIELD (2012): “Complementarity and Multidimensional Heterogeneity in Matching Markets,” mimeo.
- AZEVEDO, E. M., AND J. D. LESHNO (2011): “A supply and demand framework for two-sided matching markets,” *Unpublished mimeo, Harvard Business School*.
- BADE, S. (2014): “Random Serial Dictatorship: the One and Only,” Royal Holloway College, Unpublished mimeo.
- CARROLL, G. (2014): “A General Equivalence Theorem for Allocation of Indivisible Objects,” *Journal of Mathematical Economics*, 51, 163–177.
- CHE, Y.-K., J. KIM, AND F. KOJIMA (2013): “Stable Matching in Large Economies,” mimeo.
- CHE, Y.-K., AND F. KOJIMA (2010): “Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms,” *Econometrica*, 78(5), 1625–1672.
- CHE, Y.-K., AND O. TERCIEUX (2015): “Efficiency and Stability in Large Matching Markets,” Columbia University and PSE, Unpublished mimeo.
- DAWANDE, M., P. KESKINOCAK, J. SWAMINATHAN, AND S. TAYUR (2001): “On Bipartite and Multipartite Clique Problems,” *Journal of Algorithms*, 41, 388403.
- FEATHERSTONE, C. (2015): “Rank Efficiency: Investigating a Widespread Ordinal Welfare Criterion,” Penn University, Unpublished mimeo.
- GALE, D., AND L. S. SHAPLEY (1962): “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 69, 9–15.
- HAERINGER, G., AND F. KLIJN (2009): “Constrained School Choice,” *Journal of Economic Theory*, 144, 1921–1947.
- 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*, pp. 53–62.
- KNUTH, D. E. (1997): *Stable marriage and its relation to other combinatorial problems*. American Mathematical Society, Providence.
- KOJIMA, F., AND M. MANEA (2010): “Incentives in the Probabilistic Serial Mechanism,” *Journal of Economic Theory*, 145, 106–123.
- KOJIMA, F., AND P. A. PATHAK (2008): “Incentives and Stability in Large Two-Sided Matching Markets,” forthcoming, *American Economic Review*.
- LEE, S. (2014): “Incentive Compatibility of Large Centralized Matching Markets,” University of Pennsylvania, Unpublished mimeo.
- LEE, S., AND L. YARIV (2014): “On the Efficiency of Stable Matchings in Large Markets,” University of Pennsylvania, Unpublished mimeo.
- LEE, T., AND J. SETHURAMAN (2011): “Equivalence Results in the Allocation of Indivisible Objects: a Unified View,” Columbia University, Unpublished mimeo.
- PATHAK, P., AND J. SETHURAMAN (2011): “Lotteries in Student Assignment: An Equivalence Result,” *Theoretical Economics*, 6, 1–17.
- PITTEL, B. (1989): “The Average Number of Stable Matchings,” *SIAM Journal on Discrete Mathematics*, 2, 530–549.
- SHAPLEY, L., AND H. SCARF (1974): “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1, 22–37.