Efficiency and Stability in Large Matching Markets

Yeon-Koo Che* Olivier Tercieux†

December 24, 2013

Preliminary; do not cite.

Abstract

We study efficient and stable mechanisms in many-to-one matching markets when the number of agents is large and individuals' preferences are drawn randomly from a class of distributions allowing for both common value and idiosyncratic components. In this context, as the market grows large, all Pareto efficient mechanisms (including top trading cycles, serial dictatorship, and their randomized variants) generate total payoffs that converge to the utilitarian upper bound. This result implies that Pareto-efficient mechanisms are asymptotically payoff equivalent in the population distribution sense — that is, "up to the renaming of agents." If objects' priorities are also randomly drawn but agents' common values for objects are heterogenous, then well-known mechanisms such as deferred acceptance and top trading cycle mechanisms fail either efficiency or stability even in the asymptotic sense. We propose a new mechanism that is asymptotically efficient, asymptotically stable and asymptotically incentive compatible.

Keywords: Large matching market, Pareto efficiency, Stability, Fairness, Payoff equivalence, Random graph theory.

*Department of Economics, Columbia University, USA. Email: yeonkooche@gmail.com.
†Department of Economics, Paris School of Economics, France. Email: tercieux@pse.ens.fr.

1 Introduction

Assigning indivisible resources such as housing, public school seats, employment contracts, branch postings and human organs are of central interest in modern market design.[1] Efficiency and stability are typically two goals in designing such matching markets. Pareto efficiency is an important goal since its failure means that a reassignment would make some participants strictly better off without harming the others. Meanwhile, stability of an assignment promotes long-term sustainability of a matching market by eliminating incentives for participants to block the assignment (Roth and Sotomayor, 1990); and even when strategic blocking is not an issue (e.g., because the supply of the resources is under the control of a non-strategic entity such as a public agency), stability possesses a desirable fairness property.[2]

Due to the recent progress in matching theory research, there is by now a well-established mechanism for attaining each of these two goals and for balancing the tradeoffs between the two goals when they are in conflict. For instance, mechanisms such as a serial dictatorship and top trading cycles (henceforth, TTC) are known to generate efficient assignments (Shapley and Scarf, 1974) and Gale and Shapley's deferred acceptance algorithms (in short, DA) are known to produce stable matchings (Gale and Shapley, 1962). Further, the DA achieves stability with a minimal efficiency loss,[3] and there is a sense in which TTC using the priorities of the suppliers produces an efficient matching with a minimal incidence of instabilities (Abdulkadiroglu, Che, and Tercieux, 2013).[4]

These knowledges are clearly useful. Yet, they leave open several fundamental and

[1]Cite Nobel lecture...

[2]Stability implies the so-called "no justified envy" (see Balinski and Sönmez (1999) and Abdulkadiroglu and Sonmez (2003)), namely that whenever a participant envies another, then the supplier of the object that the envied agent receives prefers that agent over the one who envies.

[3]It is well known that – with strict preferences – DA yields a stable matching that Pareto dominates all other stable matchings for the participants (on the proposing side) (Gale and Shapley, 1962). Further, there is no individually rational assignment that makes all participants on the proposing side strictly better off relative to the DA.

[4]More precisely, in the one-to-one matching, any mechanism that is efficient, strategy-proof and weakly dominates in stability TTC in the sense that pairs that do not block under TTC do not wish to block must coincide with the TTC. This result does not extend to the many-to-one matching, however. See Abdulkadiroglu, Che, and Tercieux (2013).

practical questions. First, Pareto efficiency is a very weak standard for efficiency, compatible with many different outcomes, including some apparently unreasonable and/or unfair. There are multitude of efficient mechanisms that lead to vastly different outcomes that treat individual participants very differently. For instance, a serial dictatorship can span the entire set of Pareto efficient outcomes, depending on the serial order chosen, and likewise TTC can lead to different outcomes depending on how that participants' endowments or suppliers' priorities are set. We do not yet know how they differ in terms of the aggregate profile of participants' payoffs (e.g., payoff distribution of participants or utilitarian welfare), and the literature has yet to produce a clear prescription on which efficient mechanism should be chosen out of so many.

Second, while the tradeoff between efficiency and stability is well understood, it remains unclear how best to resolve the tradeoff when both goals are important. As noted above, the standard approach is to attain one goal with the minimal sacrifice of the other. Whether this is the best way to resolve the tradeoff is far from clear. For instance, one can imagine a mechanism that is neither stable nor efficient but may be superior to DA and TTC because it involves very little loss on each account.

The purpose of the current paper is to answer these questions and in the process provide useful insights on practical market design. These questions remain outstanding since our analytical framework is so far driven primarily by the "qualitative" notions of the two goals. To make progress, we therefore need to relax them "quantitatively." To do so requires some structures on the model. First, we consider markets that are "large" in 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 Medical Match, each year about 20,000 applicants participate to fill a position in one hospital program out of 3,000 to 4,000 programs. In NYC School Choice, about 90,000 students apply each year to 500 school programs. Second, we assume that participants' preferences are generated at random according to some reasonable distributions. Specifically, we consider a model in which each agent's utility from an object depends on a common component (i.e., that does not vary across agents) and an idiosyncratic component that is drawn at random independently (and thus varies across the agents).

Studying the limit properties of a large market with random preferences generated in this way provides a framework for answering our questions. In particular, this framework enables us to perform meaningful "quantitative" relaxations of the two desiderata: we can look for mechanisms that are **asymptotically efficient** in the sense that, as the economy becomes large, with high probability (i.e., approaching one), the fraction of agents who would gain more than some arbitrarily small amount from a Pareto improving assignment goes to zero, and mechanisms that are **asymptotically stable** in the sense that as the economy becomes large, with high probability, the fraction of agents and objects who would each gain more than some arbitrarily small payoff from forming a blocking pair goes to zero. Our findings are as follows.

First, all Pareto efficient mechanisms yield aggregate payoffs, or utilitarian welfare, that converge to the same limit—more precisely the utilitarian optimum—as the economy grows large (in the sense described above). This result implies that as the economy grows large the alternative efficient mechanisms become virtually indistinguishable in terms of the aggregate payoff distribution of the participants; with the probability approaching one, they become virtually identical. In other words, up to the "renaming" of agents, agents' payoffs are asymptotically equivalent across different efficient mechanisms. The practical implication of this result is that if one cares only about efficiency, and the assumptions of the model is valid, one need not distinguish the alternative efficient mechanisms at least in terms of the aggregate payoff profiles. For instance, SD with a random serial order or TTC with random endowments or random (non intrinsic) priorities would achieve efficiency with desirable ex ante fairness property.

Second, considering an environment in which the agents' priorities at the objects are drawn at random (e.g., possibly due to the use of lotteries), we find that the efficiency loss from the DA and the stability loss from TTC do not disappear when the objects differ significantly in qualities, namely in terms of the common components of the agents' preferences. Possible inefficiencies of DA and possible instabilities of TTC are well known from the existing literature; what we are adding here is that they remain "quantitatively" significant in the large market. The reasons can be explained in intuitive terms. Suppose the objects come in two tiers, high quality and low quality, and despite the idiosyncratic preference shocks every high-quality object dominates every low-quality objects for each agent. In this case, the (agent-proposing) DA has all agents compete first for every high-quality object before they start proposing to a low-quality object. Hence, a stable assignment—

even the agent-optimal stable matching—is largely dictated by the priorities/preferences of the objects (more precisely their suppliers), with the agents' preferences having very little influence on the outcome. In other words, the competition among agents causes the stability requirement to entail a significant efficiency loss for the agents. Meanwhile, in the same environment, under TTC, a significant fraction of agents assigned low-quality objects can form blocks with a significant number of high-quality objects whose priorities/preferences are ignored when they are traded among agents. This finding is not only an interesting theoretical finding, but it has an important implication for practical market design, since it suggests that the standard approach of achieving one goal with the minimal sacrifice of the other may not be the best.

Indeed, our third finding is that there is a novel mechanism that is both asymptotically efficient and asymptotically stable. This mechanism runs a (modified) DA in multiple stages. Specifically, all agents are ordered in some way, and in each step an agent applies *one at a time* according to the serial order to the best object that has not yet rejected him[5] and the proposed object accepts or rejects the applicant, much as in the standard DA. If at any point an agent applies to an object that holds an application, one agent is rejected, and the rejected agent in turn applies to the best object among those that have not rejected him. This process goes on until an agent is rejected by more than a certain "threshold" number of times. Then the stage is terminated at that point, and all the tentative assignments up to that point become final. The next stage then begins with the last agent (who triggered termination of the last stage) applying to the best remaining object. The stages proceed in this way until no rejection occurs.

This "staged" version of DA resembles the standard DA except for one crucial difference: the mechanism periodically terminates a stage and finalizes the tentative assignment up to that point. The event triggering the termination of a stage is that an agent is rejected a number of times during the stage exceeding a certain threshold. Intuitively, the mechanism turns on a "circuit breaker" whenever the competition "overheats" to a point that puts an agent at the risk of losing an object he ranks highly to an agent who ranks it relatively lowly (more precisely below the threshold rank). This feature ensures that an object assigned at each stage does go to an agent who ranks it relatively highly among those objects available

[5] DA where offers are made according to a serial order was first introduced by McVitie and Wilson (1971).

at that stage.

Given the independent drawing of idiosyncratic shocks, the "right" threshold turns out to be $3\log^2(n)$. Given the threshold, the DA with a circuit breaker produces an assignment that is both asymptotically stable and asymptotically efficient. The analytical case for this mechanism rests on the limit analysis, but the mechanism appears to perform well even away from the limit. Our simulation based on the case with $n = 100, 200, 500$, 1,000, 2,000 shows that our mechanism generates a significantly higher surplus to the agents with hardly any loss on the objects side.

One potential concern about this mechanism is its incentive property. While the mechanism is not strategy proof, the incentive problem does not appear to be severe. A manipulation incentive arises only when an agent is in a position to trigger the circuit breaker since then the agent may wish to apply to some object safer instead of a more popular one with high probability of rejecting him. The probability of this is one out of the number of agents assigned in the current stage, which is in the order of n, so with a sufficient number of participants, the incentive issue is rather small. Formally, we show that the mechanism induces truthful reporting as an ϵ-Bayes Nash equilibrium. Further, any symmetric Bayes Nash equilibrium is asymptotically efficient and asymptotically stable.

Our DA mechanism with a circuit breaker bears some resemblance to the features that are found in popular real-world matching algorithms. The "staged termination" feature is similar to the school choice program used to assign students to colleges in China (Chen and Kesten (2013)). More importantly, the feature that prohibits an agent from "outbidding" another over an object that the former ranks lowly but the latter ranks highly is present in the truncation of participants' choice lists, which is practiced in virtually every implementation of the DA in real settings. Our large market result could provide a potential rationale for the practice that is common in actual implementation of DA but has been so far difficult to rationalize (see Haeringer and Klijn (2009), Calsamiglia, Haeringer, and Klijn (2010) and Pathak and Sömez (2013)).

Relation to the Literature

The present paper is connected with several strands of literature. First, it is related to the literature that studies large matching markets, particularly those with large number of object types and random preferences; see Immorlica and Mahdian (2005), Kojima and Pathak (2008), Lee (2012), Knuth (1996), Pittel (1989) and Ashlagi, Kanoria, and Leshno (2013). The first three papers are concerned largely with the incentive issues arising in DA. The last three papers are concerned with the ranks of the partners achieved by the agents on the two sides of the market under DA, so they are closely related to the current paper whose focus is on the payoffs achieved by the agents. In particular, our asymptotic inefficiency result of DA follows directly from Ashlagi, Kanoria, and Leshno (2013). Unlike these papers, our paper considers not just DA but also other mechanisms and also has broader perspectives dealing with efficiency and stability.

Another 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 (2008), 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 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 both the analysis and the insights quite different. The two strands of large matching market models capture issues that are relevant in different real-world settings and thus complement each other.[6]

Methodologically, the current paper utilizes the framework developed in the random graph and random mapping theory; see Bollobas (2001) and Dawande, Keskinocak, Swaminathan, and Tayur (2001) for instance.

2 Set-up

We consider a model in which a finite set of agents are matched with a finite set of objects, at most one object for each agent. Since our exercise will involve studying the limit of

[6] 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 feel. School choice in some district such as Boston Public Schools could be a suitable application, since only a handful of schools each enroll hundreds of students. 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 district such as New York Public Schools would fit the description.

a sequence of infinite economy, 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 ease.

The object types can be interpreted as schools or housing types. Each object type o has $q_o \geq 1$ **copies** or **quotas**. Since 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 $\bar{q} \geq 1$ s.t. $q_o \leq \bar{q}$ for all $o \in O^n$ and all n. The assumption that $Q^n = n$ is only for convenience as long as it grows at order n our results will go through. Similarly, the assumption that the number of copies of each object is uniformly bounded is not necessary as long as it grows at a rate which is low enough.[7]

Throughout, we shall consider a general class of random preferences that allows for a positive correlation among agents on the objects. Specifically, each agent $i \in I^n$ receives 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 $[\underline{\xi}, \bar{\xi}] \subset \mathbb{R}_+$ according to the uniform distribution.[8]

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.

We assume that the agents' common value for object $o \in O$, u_o, takes an arbitrary value in $[0, 1]$ in an n-economy, and its population distribution is given by DDF,

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

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

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

[7] As we will see (footnote 13), we can allow $\bar{q}$ to evolve with n but then $\bar{q}$ must be $O(n/log(n))$.

[8] The uniform distribution is without loss. See Lee (2012).

interpreted as the fraction of the object types whose common value exceeds u.

We assume that these DDFs have well-defined limits. That is, X^n and Y^n converge uniformly to X and Y, respectively, where X, Y are nonincreasing and left-continuous, and $X(\cdot) - Y(\cdot) \geq 0$. We further assume that $X(0) = 1$ and $Y(0) > 0$.The first assumption means that each object type may possibly have multiple copies, so the model allows for many-to-one matching but also includes as a special case an one-to-one matching with $X(\cdot) = Y(\cdot)$. The last assumption means that the number of object types increase linearly in n. We allow X and Y to be fairly general, allowing for atoms, but only at finitely many points.

Two special cases of this model are of interest. The first is a **finite-tier model.** In this model, the objects are partitioned into finite tiers, $\{O_1^n, ..., O_K^n\}$, 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 denotes also 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, as will be the case in situations in which schools are distinguished in different categories or by regions, and houses may come in clearly distinguishable tiers. For the most part, the finite model serves as an analytical vehicle that will be used to analyze the general model. From this perspective, the finite model is useful to focus on since it brings out, in the most transparent way, the insight of the random-graph theory framework that we use.

Another special case of our model the **full-support model** in which the limit distribution Y is strictly increasing in its support. This model is very similar to Lee (2012), who also considers random preferences that consist of common and idiosyncratic terms. One difference is that his framework assumes that the common component of the payoff is also drawn uniform randomly from a positive interval. Our model assumes common values to be arbitrary, but with 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 (2012)'s, except that current model also allows for atoms in the distribution of Y.

Unless specified, we are referring to a general model that has 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 $\mu : I \to O \cup \{\emptyset\}$ such that

$|\mu^{-1}(o)| \leq q_o$, 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 are suppressed for notational convenience.

In practice, a particular matching chosen will depend on the realized preferences of the agents as well as other features of the economy that the matching institution may condition on. For instance, if the objects O are institutions or individuals, their preferences on their matching partners will typically impact on what matching will arise. Alternatively, one may wish the matching to respect the existing rights that the individuals may have over the objects, for instance, the objects may be housing, and some units may have existing tenants who may have priority over these units. Likewise, the objects may be schools, and the agents are students, and assignment to a school may favor the students whose siblings already attend the school or those living close to the school. Some of these factors are random depending on the features (not captured by their idiosyncratic component) that are random. 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 idiosyncratic component of payoffs, and θ is the realization of all other variables that influence the particular matching that is selected, and let Ω denote the set of all possible states.

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 matching 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.

A Pareto efficient matching may arise from a number of different matching mechanisms in a variety of circumstances. Let $\mathcal{M}^*$ be the set of all Pareto efficient mechanisms.

3 Cardinal Payoff implication of Pareto Efficiency

The first question of the present paper is whether a Pareto efficient matching approaches a utilitarian efficient outcome. To answer this question, it is useful to consider a utilitarian benchmark. For this purpose, the following thought experiment is useful. Suppose first we assign the agents to objects just to maximize their common component of their payoffs. Clearly, this can be done by assigning the agents with objects in the order of their common values, starting from the object (type) with the highest common value, and once exhausting those, assigning those with lower common values. (Since such an assignment does not address the idiosyncratic component of agents' payoffs, there is no guaranteeing that the assignment will attain utilitarian or even Pareto efficiency.) Since the maximum value of the idiosyncratic payoff cannot exceed $\overline{\xi}$, however, we can at least say that the per-capital payoff for agent cannot exceed the utilitarian upper bound for the n-economy $\int_0^1 U(u,\overline{\xi})dX^n(u)$ which converges to the **limit utilitarian upper bound**:

$$U^* := \int_0^1 U(u,\overline{\xi})dX(u).$$

The aggregate payoff distribution of an economy, whether it is a finite n-economy or its limit, can be represented by a decumulative distribution function, i.e., a nonincreasing left-continuous function $\bar{F}$ mapping from $[0, U(1,\overline{\xi})]$ to $[0,1]$ with $\bar{F}(U(0,\underline{\xi})) = 1$. The number $\bar{F}(z)$ is interpreted as the fraction of the agents with payoff no less than z. Any matching μ also induces an aggregate payoff distribution representable by a DDF. Hence, a mechanism μ induces a random DDF, denoted $\bar{F}^{\mu}$. For a later purpose, it is convenient to introduce a notion of distance between two DDFs, called Lévy metric: For any DDFs, $\bar{F}$ and $\bar{G}$, let

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

One can simply think of this as the distance of two DDFs measured on the points of continuity.[9]

Consider a sequence of mechanisms $\{\mu^n\}$ and let $\{\bar{F}^{\mu^n}\}$ denote the sequence of aggregate payoff distributions resulting from them. For any (non-random) decumulative distribution

[9] Convergence of DDFs in Lévy metric boils down to a pointwise or weak convergence, by the Portmanteau's lemma.

function (DDF) $\bar{F}$, we say that a sequence of mechanisms $\{\mu^n\}$ **converges to $\bar{F}$ in payoff distribution** if for any $\epsilon > 0, \upsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$,

$$\Pr\left[L(\bar{F}^{\mu^n}, \bar{F}) \geq \epsilon\right] < \upsilon.$$

In words, the notion means that the fraction of the agents enjoying any given payoff under μ^n can be made arbitrarily similar to the fraction of agents enjoying an arbitrarily similar payoff under the benchmark distribution $\bar{F}$ with an arbitrarily high probability if n is made sufficiently large. The definition thus speaks to the asymptotic probabilistic convergence of population distribution of payoffs under a sequence of mechanisms. Importantly, it does not speak about the convergence of a payoff of a particular agent or agent types. It thus more apt to interpret the convergence in payoffs as a *convergence of payoffs up to renaming of agents.*

It is of particular interest to consider mechanisms converging in payoff to the limit utilitarian upper bound. Let $\bar{F}^*$ denote the population payoff distribution corresponding to the limit utilitarian upper bound. That is, $\bar{F}^*(z) = X(U^{-1}(z; \overline{\xi}))$ for each z. We say that a sequence of mechanisms $\{\mu^n\}$ **converges to the limit utilitarian upper bound** if it converges in payoff to $\bar{F}^*$.

We shall also consider a class of mechanisms converging to a benchmark payoff distribution in a certain uniform sense. **A sequence of families of mechanisms**, $\{\hat{\mathcal{M}}^n\}$, where $\hat{\mathcal{M}}^n \subset \mathcal{M}^n$, **equi-converges in payoff to** $\bar{F}$ if, for any $\epsilon > 0, \upsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$,

$$\Pr\left[\sup_{\mu^n \in \hat{\mathcal{M}}^n} L(\bar{F}^{\mu^n}, \bar{F}) \geq \epsilon\right] < \upsilon.$$

Uniform convergence to the limit utilitarian upper bound is analogously defined. An implication of uniform convergence for a sequence of families of mechanisms is that the payoff distribution within that family becomes very similar to each other.

We say that **a class of mechanisms**, $\hat{\mathcal{M}} \subset \mathcal{M}$, **is asymptotically payoff equivalent** if, for any $\epsilon > 0, \upsilon > 0$, there exists $N \in \mathbb{N}$ such that for all $n > N$,

$$\Pr\left[\sup_{\mu^n, \tilde{\mu}^n \in \hat{\mathcal{M}}^n} L(\bar{F}^{\mu^n}, \bar{F}^{\tilde{\mu}^n}) \geq \epsilon\right] < \upsilon.$$

Again it is important to note that the payoff equivalence in the sense of the aggregate payoff distribution. Roughly speaking, if a class of mechanisms are asymptotically payoff equivalent, the payoffs of the agents are similar "up to renaming" across the different mechanisms within that class uniformly in the limit as the economy grows large.

Our key argument exploits a result in random graph theory. Hence, we begin with the relevant model of random graph and the result. A **bipartite graph** G consists vertices, $V_1 \cup V_2$, and edges $E \subset V_1 \times V_2$ across V_1 and V_1 (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. An independent set is **balanced** if $|\bar{V}_1| = |\bar{V}_2|$ and we refer to a balanced independent set with the maximum size (i.e., the cardinality of $\bar{V}_i$) as a **maximal balanced independent set**. 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 will prove crucial for our subsequent results.

LEMMA 1. *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 $|V_2| = m = O(n)$. For any $\gamma \in (0, 1)$,*

$$\Pr\left[\exists \text{ an independent set } \hat{V}_1 \times \hat{V}_2 \text{ with } |\hat{V}_1| = |\hat{V}_2| \geq \gamma n\right] \leq \kappa \left(\frac{1}{n!}\right)^2$$

for some strictly positive constant κ.

PROOF. See Appendix A. □

REMARK 1. *Given a graph with vertices $V_1 \cup V_2$, $\bar{V}_1 \times \bar{V}_2$ is a biclique if and only if $\bar{V}_1 \times \bar{V}_2$ is an independent set in its complement graph. Thus, in an environment where $|V_1| = |V_2| = n$, Theorem 2.6. in Dawande, Keskinocak, Swaminathan, and Tayur (2001), yields that the probability that there exists a balanced independent set of size $|\hat{V}_1| = |\hat{V}_2| \geq \gamma \log(n)$ tends to 0 as n goes to infinity. Our argument follows their proof. Beyond the fact that our environment allows for the possibility that $|V_1| \neq |V_2|$, we focus on balanced independent sets of size $|\hat{V}_1| = |\hat{V}_2| \geq \gamma n$, while they focus on balanced independent sets of size $|\hat{V}_1| = |\hat{V}_2| \geq \gamma \log(n)$. Because our requirement on the size of the balanced independent set is weaker, the rate of convergence (of order $\left(\frac{1}{n!}\right)^2$ in our case) turns out to be much faster than in their case (of order $\left(\frac{1}{\log(n)!}\right)^2$).*

We now study the limit implication of any Pareto efficient mechanism. To this end, we first partition the set of objects in each n-economy based on their common values into finite tiers. That is, let $\sup(supp(Y)) =: u^1 > u^2 > ... > u^K = 0$. In the finite-tier economy, the tiers here can be defined to correspond to the finite common values. In the general model, any such tiers will induce a DDF which will approximate the true distribution Y^n from below (as K increases). Define $O^n_{\le k} := \{o \in O^n | u_o \ge u^k\}$ be the set of objects in tier k or better, and let $Y^n_{\le k} := Y^n(u^k)$ and $X^n_{\le k} := X^n(u^k)$, denote the associated mass of objects and the associated mass of copies of objects. Define similarly $Y_{\le k} := Y(u^k)$ and $X_{\le k} := X(u^k)$ for the limit economy. From now on, for notational ease, we shall suppress n except for X^n and Y^n to avoid confusion with their limit counterparts.

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 $f_\mu(\omega) : I \to I$, 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 $SD^{f_\mu(\omega)}_\omega(i) = \mu_\omega(i)$ for each $i \in I$. Since the matching μ arising from the mechanism depends on the random state ω, so is the serial order f implementing μ. In the sequel, we shall study a Pareto efficient matching mechanism μ via the associated SD^{f_μ}. 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

$$I_{\le k}(\mu) := \{i \in I | f(i) \le Q^n X^n_{\le k}\}.$$

be the set of agents who have a serial order within the total supply of objects in tiers k or better (in the equivalent serial dictatorship implementation). For any ϵ, the set

$$\bar{I}_{\le k}(\mu) = \big\{i \in I_{\le k}(\mu) | U_i(SD^f(i)) \le U(u^k, \bar{\xi} - \epsilon)\big\},$$

consists of the agents who realize payoff no greater than $U(u^k, \bar{\xi} - \epsilon)$ while having a serial order within $Q^n X^n_{\le k}$. The following lemma will be crucial for the main result.

Lemma 2. *For any $\epsilon > 0$ and $\gamma > 0$,*

$$\Pr\left[\exists \mu \in \mathcal{M}^* \ s.t. \ \frac{|\bar{I}_{\le k}(\mu)|}{|I_{\le k}(\mu)|} \ge \gamma\right] \le \kappa \left(\frac{1}{n!}\right)^2$$

for some constant $\kappa > 0$.

PROOF. See Appendix A. □

THEOREM 1. *The Pareto efficient mechanisms equi-converge in payoff to the limit utilitarian upper bound.*

PROOF. See Appendix A. □

A simple corollary is:

COROLLARY 1.

$$\inf_{\mu \in \mathcal{M}^*} \mathbb{E}\left[\frac{\sum_{i \in I} U_i(\mu(i))}{|I|}\right] \to U^* \text{ as } n \to \infty.$$

As mentioned earlier, the following payoff equivalence holds across Pareto efficient mechanisms.

COROLLARY 2. *The class of Pareto efficient mechanisms is asymptotically payoff equivalent.*

4 Balancing Efficiency and Stability

We now consider matching mechanisms that rely on ordinal preference messages on both sides of the market. Most common mechanisms used in practice in fact use ordinal preferences messages, with the deferred acceptance mechanism and the top trading cycles mechanism being two notable examples (these two mechanisms are defined in section 4.1 below).

For simplicity, we suppose that objects have priorities (in the sequel, we use the term priorities which could be understood as preferences if, for instance, objects are interpreted as being institutions like firms) that are random. Specifically, each object $o \in O$ receives utility from getting matched with individual $i \in I$:

$$V_i(o) = V(\eta_{i,o}),$$

where *idiosyncratic shock* $\eta_{i,o}$ is a random variable drawn independently and identically from $[\underline{\eta}, \bar{\eta}] \subset \mathbb{R}_+$ according to the uniform distribution.[10] We further assume that the

[10]Again, the uniform distribution is without loss.

function $V(\cdot)$ takes values in $\mathbb{R}_+$, is strictly increasing and continuous in the idiosyncratic shock. The utility of remaining unmatched is assumed to be 0 so that all objects find all individuals acceptable. In addition to this, we will restrict out attention to the finite-tier model where the objects are partitioned into finite tiers, $\{O_1, ..., O_K\}$. For each k, we denote the proportion of objects in O_k, i.e., $\frac{|O_k|}{|O|}$, by x_k. We also restrict our attention to the one-to-one case.

4.1 Two mechanisms

□ **Top trading cycle (TTC) mechanisms:**

Following Abdulkadiroglu and Sonmez (2003), we define the **top trading cycles** mechanism, denoted TTC, as follows:[11] Fix $t \geq 1$. Then, Round t proceeds as follows. Each individual $i \in I$ points to his most preferred object (if any). Each object $o \in O$ points to the individual to which it assigns the highest priority. Since the number of individuals and objects are finite, the directed graph so obtained has at least one cycle, Every individual who belongs to a cycle is assigned to the object he is pointing at. Any individual who has been assigned an object as well as any object which has been assigned an individual is removed. The algorithm terminates when all individuals have been assigned; otherwise, it proceeds to Round $t + 1$.

This algorithm terminates in a finite number of rounds. Indeed, at the end of each round, at least one individual is removed and there are finitely many individuals. The TTC mechanism is defined as a function which for each realization of individuals' preferences as well as objects' priorities selects a matching obtained by the above algorithm.

As is well-known, TTC is a Pareto-efficient matching mechanism and is not stable. It is also strategy-proof, i.e., it is a dominant strategy to report preferences truthfully. The top trading cycles mechanism is one of the few mechanisms regularly used in practice. For instance, in the school choice context, it was used until recently in New Orleans and recently, San Francisco announced plans to implement a top trading cycles mechanism. A generalized version of TTC is also used to assign kidneys to sick patients (see Sonmez and Unver (2011)).

[11] Shapley and Scarf (1974) first introduced a simplified version of the top trading cycle mechanism in the housing market setting. The original idea is attributed to David Gale by Shapley and Scarf (1974).

□ The deferred acceptance (DA) mechanism

The original deferred acceptance (DA) algorithm was defined by Gale and Shapley (1962). An alternative equivalent algorithm has been proposed by McVitie and Wilson (1971). For convenience, we only refer to this alternative formulation. The DA algorithm is defined recursively as follows.

Step 0: Linearly order individuals in I.

Step 1**:** Let individual 1 make an offer to his most favorite object in O. This object tentatively holds individual 1, and go to Step 2.

Step $i \geq 2$**:** Let agent i make an offer to his most favorite object o in O among the objects to which he has not yet made an offer. If o is not tentatively holding any individual, then o tentatively holds i. whenever $i = n$, end the algorithm; otherwise iterate to Step $i + 1$. If however o is holding an individual tentatively—call him i^*—object o chooses between i and i^* accepting tentatively the one who is higher in its preference list, and rejecting the other. The rejected agent is named i and we go back to the beginning of Step i.

This process terminates in finite time and yields to a matching μ. The DA mechanism is defined as a function which for each realization of individuals' preferences as well as objects' priorities selects a matching obtained by the above algorithm.

While the DA mechanism is stable, it is not Pareto-efficient. However, it is known to be student-optimal stable: given individuals' preferences as well as objects' priorities, it produces a stable matching which is the matching that is weakly preferred to any other stable matching by all individuals. It is also strategy-proof, i.e., it is a dominant strategy to report preferences truthfully (Dubins and Freedman (1981); Roth (1982)). Here again the DA mechanism is one mechanism regularly used in practice. It has been implemented in both New York City (Abdulkadiroglu, Pathak, and Roth (2005)) and Boston (Abdulkadiroglu, Pathak, Roth, and Sonmez (2005)).

□ Asymptotic Notions of Efficiency and Stability

Now, we define two notions that will be central to our analysis. A matching mechanism μ is **asymptotically efficient** if for any mechanism μ' that weakly Pareto-dominates μ

for the agents I, for any $\epsilon, \gamma, \upsilon$, there exists $N \in \mathbb{N}$ such that for all $n > N$, we have

$$\Pr\left\{\frac{|I_\epsilon(\mu'|\mu)|}{n} \geq \gamma\right\} < \upsilon,$$

where

$$I_\epsilon(\mu'|\mu) := \{i \in I | U_i(\mu(i)) < U_i(\mu'(i)) - \epsilon\}.$$

In words, a matching is asymptotically efficient if, as the economy gets large, with high probability any Pareto improving rematching, if there is any, could make only an arbitrarily small fraction of agents more than ϵ better off.

A matching mechanism μ is **asymptotically stable** if, for any $\epsilon, \gamma, \upsilon$, there exists $N \in \mathbb{N}$ such that for all $n > N$, we have

$$\Pr\left\{\frac{|J_\epsilon(\mu)|}{n(n-1)} \geq \gamma\right\} < \upsilon,$$

where

$$J_\epsilon(\mu) := \{(i, o) \in I \times O | U_i(o) > U_i(\mu(i|\omega)) + \epsilon \text{ and } V_o(i) > V_o(\mu(o|\omega)) + \epsilon\}.$$

To apply this notion, we first consider an **ϵ-block,** i.e., a pair of unmatched agent and object who each would gain ϵ or more from matching each other rather than matching its original partner. Asymptotic stability then requires that for any $\epsilon > 0$, with high probability the fraction of these ϵ-blocks out of all $n(n-1)$ "possible" blocking pairs is vanishing as the economy grows large. Hence, in a large market, an asymptotically stable matching will not admit a large number of agents and objects with discrete motive to block them. It is still possible that a large number of agents may be willing to form blocks with some objects, but in that case the number of such objects will be small relative to the willing agents, accommodating only a small number. In such a case, the By contrast, if a matching is not asymptotically stable, with some some probability bounded below from zero, a non-vanishing fraction of agents and objects can form blocks simultaneously, so it can be a destabilizing force. In this sense, the notion is natural and useful for large markets.

□ **Preliminary Results**

We first consider the case in which the participants' preferences for the objects are uncorrelated. That is, the support of the common component of the agents' utilities are degenerate. In our finite-tier model, it is equivalent to saying that the number of tiers $k = 1$. In this case both DA and TTC involve little tradeoff:

Proposition 1. *If the support of $Y(\cdot)$ is degenerate, then DA is asymptotically efficient, and TTC is asymptotically stable.*

The former follows from Pittel (1992), which shows that under DA with high probability all participants are assigned objects which they rank within $3\log^2(n)$. Since this number grows much slowly relative to n, we get that the agents attain payoffs arbitrarily close to the upper bound $\bar{\xi}$ as the economy grows large. Hence, DA is asymptotically efficient.[12]

Asymptotic stability of TTC follows from Theorem 1. Since TTC is a Pareto efficient mechanism, Theorem 1 implies that for any $\epsilon > 0, \gamma > 0, \upsilon > 0$, for all sufficiently high n, the fraction of the set $I_{\leq 1}$ of agents realizing payoffs less than $u_1^0 + \overline{u} - \epsilon$ being greater than γ can occur with probability at most of υ. Since $I_{\leq 1} \supset J_\epsilon(TTC)$, the asymptotic stability of TTC follows in this case.

As we show below, Proposition 1 no longer holds when the agents' preferences are correlated, in particular, when some objects are perceived by "all" agents to be better than the other objects. This situation is quite common in many contexts, such as school assignment, since schools have distinct qualities that students and parents evaluate in a similar fashion.

To consider such an environment in a simple way, we shall suppose the objects are divided into two tiers O_1 and O_2 such that $|I| = |O_1| + |O_2| = n$. As assumed earlier, $\lim_{n\to\infty} \frac{|O_i|}{n} = x_i > 0$. In addition, we assume that each object in O_1 is considered by every agent to be better than each object in O_2: $U(u_1^0, 0) > U(u_2^0, \overline{\xi}))$, where u_1 and u_2 are common values of the objects from tier 1 and tier 2, respectively. The preferences/priorities by the objects are given by idiosyncratic random shocks, as assumed above. In this environment, we shall show that the standard tradeoff between DA and TTC extends to the large markets even in the asymptotic sense — namely, DA is not asymptotically efficient and TTC is not asymptotically stable. This observation runs counter to the common wisdom based on finite market that correlation of preferences on one side usually renders stable

[12]Our notion of efficiency focuses on one side of the market: the individuals' side. It is worth noting here that even if we were to focus only on the other side: the objects' side, efficiency would still follow from from Pittel (1992) even though we are using DA where individuals are the proposers.

allocations efficient or efficient allocations stable.

4.2 Asymptotic Instability of TTC

Our first result is that, with correlated preferences, TTC fails to be asymptotically stable.

Theorem 2. *In our model with two tiers, TTC is not asymptotically stable. More precisely, there exist $\epsilon, \upsilon, \gamma$ all strictly positive, such that, for all $n > N$,*

$$\Pr\left\{\frac{|J_\epsilon(TTC)|}{n(n-1)} \geq \gamma\right\} > \upsilon,$$

for some N.

We shall provide a proof for the remainder of this subsection. To begin, we call a cycle of length 2 — namely, an agent points to an object, which points back to the agent —- a **short cycle**, and any cycle of length greater than 2 will a **long cycle**. For the proof, we shall observe that the objects in O_1 assigned via long cycles do not enjoy high payoffs (since the trade along such cycles ignores objects' preferences), and that many of these objects can form ϵ-blocks with the agents that are assigned to objects in O_2. Since the latter agents are significant (i.e., in the order of n), the important part of the argument will be to show that the objects that are assigned via long cycles and are willing to form ϵ-blocks with them are also significant in numbers.

To begin, let $\hat{O}$ denote the random set of objects in O_1 that are assigned in TTC via long cycles, and let $I_2 := \{i \in I | TTC(i) \in O_2\}$ be the random set of agents who are assigned under TTC to objects in O_2. Appendix B establishes the following result.

Lemma 3. *There exist $\gamma > 0, \delta > 0, N > 0$ s.t.*

$$\Pr\left\{\frac{|\hat{O}|}{n} > \delta\right\} > \gamma,$$

for all $n > N$.

Proof. See Appendix B. □

While this result is quite intuitive, its proof is not trivial. Using an appropriate extension of "random mapping theory," we can compute the expected number of objects in O_1 that are assigned via long cycles in the first round of TTC. But, this turns out to be insufficient for our purpose since the number of objects that are assigned in the first round of TTC (in the the order of $\sqrt{n}$) comprises a vanishing fraction of n as the market gets large. Extending the random mapping analysis to the subsequent rounds of TTC is difficult, however, since the preferences of the agents and objects remaining after the first round are no longer i.i.d. Hence it is necessary to investigate the precise random structure of the preferences that evolve over time. Appendix B does this. In particular, we establish that the number of objects (and thus agents) assigned in each round of TTC follows a simple Markov structure, implying that the number of agents cleared in each round is not subject to the conditioning issue. The composition of the cycles, in particular short versus long cycles, is subject to the conditioning issue, however. Nevertheless, we managed to show that the number of objects assigned in each round of TTC can be bounded above. And this bound, combined with the Markov property of the number of objects assigned in each round, produces Lemma 3.

We next establish that any randomly selected (unmatched) pair from $\hat{O}$ and I_2 forms an ϵ-block with positive probability for sufficiently small $\epsilon > 0$.

LEMMA 4. *There exist $\varepsilon > 0, \zeta > 0$ such that, for all $n > N$, for any $\epsilon \in [0, \varepsilon)$,*

$$\Pr\left[\eta_{jo} \geq \eta_{TTC(o)o} + \epsilon \,\middle|\, o \in \hat{O}, j \in I_2\right] > \zeta.$$

PROOF. See Appendix A. □

Let $\hat{I}_2(o) := \{i \in I_2 | \eta_{io} > \eta_{TTC(o)o} + \epsilon\}$ be the set of agents that are assigned to objects in O_2 and but could serve partners of ϵ-blocks for objects in O_1. The implication of Lemma 4 is that there are non-vanishing number of such agents for any objects in O_1 that are assigned via long cycles.

COROLLARY 3. *For any $\epsilon > 0$ sufficiently small, there exist $\zeta > 0, N > 0$ such that, for all $n > N$,*

$$\mathbb{E}\left[\frac{|\hat{I}_2(o)|}{n} \,\middle|\, o \in \hat{O}\right] \geq x_2\zeta$$

PROOF. See Appendix A. □

The theorem follows from Lemma 3 and Corollary 3. The former implies that as the economy grows, the number of objects assigned via long cycles remain significant. The latter implies that each of such object finds many agents assigned by TTC to O_2 desirable for forming ϵ-blocks. More precisely, for any sufficiently small $\epsilon \in (0, U(u_1^0, 0) - U(u_2^0, \bar{\xi}))$, we get that, for any large n,

$$
\begin{aligned}
\mathbb{E}\left[\frac{|J_\epsilon(TTC)|}{n(n-1)}\right] &\geq \mathbb{E}\left[\sum_{o\in\hat{O}}\frac{|\hat{I}_2(o)|}{n(n-1)}\right] \\
&\geq \mathbb{E}\left[\sum_{o\in\hat{O}}\frac{|\hat{I}_2(o)|}{n(n-1)}\Big||\hat{O}|\geq\delta n\right]\Pr\{|\hat{O}|\geq\delta n\} \\
&\geq \mathbb{E}\left(\mathbb{E}\left[\sum_{o\in\hat{O}}\frac{|\hat{I}_2(o)|}{n(n-1)}\Big||\hat{O}|\geq\delta n,\hat{O}\right]\right)\gamma \\
&= \mathbb{E}\left(\sum_{o\in\hat{O}}\mathbb{E}\left[\frac{|\hat{I}_2(o)|}{n(n-1)}\Big||\hat{O}|\geq\delta n,\hat{O},o\in\hat{O}\right]\right)\gamma \\
&\geq \delta n\mathbb{E}\left[\frac{|\hat{I}_2(o)|}{n(n-1)}\Big|o\in\hat{O}\right]\gamma \\
&\geq \delta\mathbb{E}\left[\frac{|\hat{I}_2(o)|}{n}\Big|o\in\hat{O}\right]\gamma\geq\delta\zeta x_2\gamma>0.
\end{aligned}
$$

The stated result of Theorem 2 follows from this.

4.3 Asymptotic Inefficiency of DA

Our second result is the asymptotic inefficiency of DA.

Theorem 3. *In our two tier model, DA is not asymptotically efficient. More precisely, there exists a matching μ that Pareto dominates DA and strictly positive numbers $\epsilon, \upsilon, \gamma$ such that*

$$\Pr\left\{\frac{|I_\epsilon(\mu|DA)|}{|I|}\geq\gamma\right\}\geq\upsilon,$$

for all $n > N$ for some N.

Proof. See Appendix A. □

The intuition behind this result is as follows. Correlation in agents' preferences means nontrivial competition for some objects. When agents compete for an object, the object is assigned based on their priorities at the object (or the preferences by the initial supplier of the object). Hence, nontrivial competition in terms of a large number of agent preferring the same set of objects, means that the ranks and thus the idiosyncratic payoffs that the agents enjoy are likely to be high. To put it differently, the stability requirement (and the preferences on the side of the objects) do not constitute a significant constraint when agents' preferences are diverse enough not to compete each other but they do constitute significant enough constraint to undermine asymptotic efficiency in the large market, when their preferences are correlated.

4.4 DA with Circuit Breaker

As we just saw, two of the most prominent mechanisms fail to find matchings which, with high probability, are asymptotically stable and asymptotically efficient. In the sequel, we define a new mechanism which finds such matchings. To be more precise, we define a class of mechanisms indexed by some κ. We will show how an appropriate value of κ can be chosen in order to achieve our goal.

Given a value κ, the DA with Circuit Breaker algorithm (DACB) is defined recursively on three objects $\hat{I}$ and $\hat{O}$ the set of remaining agents and objects, respectively, and a counter for each agent that records the number of times the agent was rejected. We first initialize $\hat{I} = I$ and $\hat{O} = O$, and set the counter for each agent to be zero.

Step 0: Linearly order individuals in $\hat{I}$.

Step 1**:** Let individual 1 make an offer to his most favorite object in $\hat{O}$. This object tentatively holds individual 1, and go to Step 2.

Step $i \geq 2$**:** Let agent i make an offer to his most favorite object o in $\hat{O}$ among the objects to which he has not yet made an offer. If o is not tentatively holding any individual, then o tentatively holds i. whenever $i = n$, end the algorithm; otherwise iterate to Step $i + 1$. If however o is holding an individual tentatively—call him i^*—object o chooses between i and i^* accepting tentatively the one who is higher in its preference list, and rejecting the other. The counter for the rejected agent increases by one. There are two cases:

1. If counter of the rejected agent is greater than or equal to κ, then each agent who

is assigned tentatively to an object in Steps $1, ..., i$ is assigned to that object. Reset $\hat{O}$ to be the set of unassigned objects and $\hat{I}$ to be the set of unassigned individuals. Reset the counter of the agent rejected at step i to be zero. If $\hat{I}$ is non-empty, go back to step 0, otherwise, terminate the algorithm.

2. If the counter of the agent rejected at Step i is below κ. The rejected agent is named i and we go back to the beginning of Step i.

This process terminates in finite time and yields to a matching μ. This algorithm modifies the McVitie and Wilson (1971) version of DA where the tentative assignments are periodically finalized. We say that a **stage** begins whenever $\hat{O}$ is reset, and the stages are numbered $1, 2, ...$ serially. For our purpose we set $\kappa := 3\log^2(n)$. The next theorem shows that this choice of κ is appropriate: the mechanism so defined is efficient and stable in the asymptotic sense.

Theorem 4. *DACB is asymptotically efficient and asymptotically stable.*

The Theorem directly follows from the proposition below.

Proposition 2. *Fix any $k \geq 1$. As $n \to \infty$, with probability approaching one, stage k of the DACB ends at step $|O_k| + 1$ and all objects in O_k are assigned. In addition, for any $\epsilon > 0$ and γ*

$$\Pr\left[\frac{|\{i \in I_k | U_i(DACB(i)) \geq U(u_k, \bar{\xi}) - \epsilon\}|}{|I_k|} \geq \gamma\right] \to 1$$

as $n \to \infty$; where $I_k := \{i \in I | DACB(i) \in O_k\}$. Similarly,

$$\Pr\left[\frac{|\{o \in O_k | V_o(DACB(o)) \geq V(\bar{\eta}) - \epsilon\}|}{|O_k|} \geq \gamma\right] \to 1$$

as $n \to \infty$.

Before we move to the proof of this proposition, we begin with a preliminary result. In the sequel, to avoid additional notations, we sometimes assume that $\underline{\xi} = 0$ and $\bar{\xi} = 1$.

Lemma 5. *Fix any $\varepsilon > 0$. Let $\hat{I}$ and $\hat{O}$ be two sets such that both $|\hat{I}|$ and $|\hat{O}|$ are in between αn and n for some $\alpha > 0$. For each $i \in \hat{I}$, let X_i be the number of objects in $\hat{O}$ for which $\xi_{io} \geq \bar{\xi} - \varepsilon$. For any $\varepsilon' < \varepsilon$*

$$\Pr\{\exists i \text{ with } X_i \leq \varepsilon' |\hat{O}|\} \to 0$$

as $n \to \infty$.

PROOF. X_i follows a binomial distribution $B(|\hat{O}|, \varepsilon)$ (recall that ξ_{io} follows a uniform distribution with support $[0,1]$). Hence,

$$\begin{aligned}\Pr\{\exists i \text{ with } X_i \leq \varepsilon'|\hat{O}|\} &\leq \sum_{i\in\hat{I}} \Pr\{X_i \leq \varepsilon'|\hat{O}|\} \\ &= |\hat{I}| \Pr\{X_i \leq \varepsilon'|\hat{O}|\} \\ &\leq |\hat{I}|\frac{1}{2}\exp\left(-2\frac{(|\hat{O}|\varepsilon - \varepsilon'|\hat{O}|)^2}{|\hat{O}|}\right) \\ &= \frac{|\hat{I}|}{2\exp\left(2(\varepsilon-\varepsilon')^2|\hat{O}|\right)} \to 0\end{aligned}$$

where the first inequality is by the union bound while the second equality is by Hoeffding's inequality. □

COROLLARY 4. *Fix any $\varepsilon' > 0$ small enough and $k \geq 1$. With probability going to 1 as $n \to \infty$, all individuals in I only rank objects in O_k within their $\epsilon'|O_k|$ favorite objects in $O_{\geq k}$.*

PROOF. By Lemma 5 where $\epsilon := u_k - u_{k+1}$, $\hat{I} := I$ and $\hat{O} := O_k$, we get that, for any $\epsilon' < u_k - u_{k+1}$, with probability going to 1 as $n \to \infty$, all individuals in I have at least $\epsilon'|O_k|$ objects in $o \in O_k$ for which $\xi_{io} > \bar{\xi} - \epsilon$. By the choice of ϵ, all such objects are better than any other objects in $O_{>k+1}$. □

PROOF OF PROPOSITION 2. We focus on $k = 1$, as will become clear, the other cases can be treated exactly in the same way.

First, let us consider the submarket composed of the $|O_1|$ first agents (according to the ranking given in the definition of DACB) and of all objects in O_1 objects. If we were to run DA here because preferences are drawn iid, by Pittel [Theorem 6.1., (b) 1992], with probability approaching 1 as n grows, at the end of (standard) DA, *all* agents have made less than $3\log(n)^2$ offers.

Now, let us come back to the original market. By Corollary 4 (and $3\log(n)^2/n \to 0$), the event that all agents' $3\log(n)^2$ favorite objects are in O_1 has probability approaching 1 as $n \to 0$. Let us condition on the event. Observe that the conditional distribution of individuals' preferences over objects in O_1 is the same as the unconditional one (of course, this is also true for the distribution of objects' priorities over individuals). Since, given our conditioning event, the $|O_1|$ first steps of DACB go exactly in the same way as DA in the submarket composed of the $|O_1|$ first agents (according to the ranking used in DACB) and of all objects in O_1 objects, applying the result by Pittel mentioned above, with probability going to 1 as $\to \infty$, we achieve the end of step $|O_1|$ of DACB before stage 1 ends (i.e., before an agent makes an application to his $3\log(n)^2$ favorite object). The outcome so far is the one achieved in DA in the submarket composed of the $|O_1|$ first agents and of all objects in O_1 objects. Hence, with probability going to 1, all individuals in the $|O_1|$ first (again according the ranking given in the definition of DACB) have a payoff close to $U(u_1, \bar{\xi})$ and the proportion of objects in O_1 which get a payoff close to $\bar{\eta}$ converges in probability to 1. Now, observe that to end step $|O_1| + 1$, the only way is that an individual applies to an object in O_2 or an individual makes an application to his $3\log(n)^2$ favorite object. Given our conditioning event, for all individuals, objects in O_2 are ranked below the $3\log(n)^2$ favorite object. Thus, the only way step $|O_1| + 1$ can end is that an individual makes an application to his $3\log(n)^2$ favorite object which implies that stage 1 ends. Hence, at the end of step $|O_1| + 1$, objects in O_1 receive even more offers than at the end of step $|O_1|$. Thus, we obtain that the proportion of objects in O_1 which get a payoff close to $\bar{\eta}$ still converges in probability to 1. In addition, with probability going to 1, $|O_1|$ individuals and objects are matched at the end of stage 1. By construction, all these individuals enjoy a payoff arbitrarily close to $U(u_1, \bar{\xi})$. Of course, this proved under our conditioning event but since this event has probability going to 1 as $n \to \infty$, this result holds even without conditioning. Thu,s we have proved Proposition 2 for the case $k = 1$.

Now, observe that at the moment a stage k is ended, the objects present in stage $k+1$ have received no offers. Thus, by the principle of deferred decisions, we can assume that the individuals' preferences over those objects are yet to be drawn. Similarly, we can assume that priorities of those objects are also yet to be drawn. Put in another way, conditional on stage k being over, we can assume without loss that the distribution of preferences and priorities is the same as the unconditional one. Thus, we can proceed inductively to complete the proof. □

Theorem 4 shows that DACB is superior to DA or TTC in large markets when the designer cares about both (asymptotic) efficiency and (asymptotic) stability. One potential drawback of DACB is that it is not strategy-proof. In particular, the agent who triggers a stage to end may potentially gain from misreporting his preferences to include in his top κ favorite objects "safe" items which are outside his top κ favorite objects but are unlikely to be popular among other agents. But the chance of becoming in the position to trigger termination of a stage is one out of those assigned in the same stage, so it is very small particularly in a large economy. Hence, the incentive problem with the DACB is not very serious. We state two results that formalize the sense in which the DACB perform well from the incentive perspective.

THEOREM 5. *Fix any $\epsilon > 0$. Under DACB, there exists $N > 0$ such that for all $n > N$, it is an ϵ-Bayes-Nash equilibrium for agents to report truthfully their preferences.*

PROOF. Let's consider the event that all (and only) objects in O_k are assigned in stage k and that each agent who is assigned object in O_k enjoys high enough idiosyncratic payoff so that he prefers his assigned object to all objects in $O_{\geq k+1}$. Given Proposition 2, the probability of this event goes to one as $n \to \infty$. Fix any agent i and let P_i be a truthful report and P_i' be a misreport. Let us further assume that whether i reports P_i or P_i', he is assigned in stage k and is not one of the individuals who triggers the end of stage k or any previous stage. For each of these reports, the agent has at least probability $1 - (\frac{1}{|O_1|} + \cdots + \frac{1}{|O_k|}) \to 1$ (as $n \to \infty$) of being in that position. Fix any realization of preferences on both sides of the market that gives rise to those events. We show that under such realization i cannot have any (ex-post) gains from misreporting P_i'. Given that our conditioning events have a (joint) probability which goes to 1, and given that the utility gains from misreporting are bounded (uniformly across preferences' realizations), for any $\epsilon > 0$, the expected utility gains from misreporting are bounded by ϵ for n sufficiently large. If agent i reports truthfully, say P_i, then, the assignment for the agent is the same as DA applied to a subeconomy consisting of O_k and the agents assigned to O_k plus the agent who triggers stage k to end and all agents except for i submit the same preferences that are truncated to contain only $3\log(n)^2$ favorite choices among O_k, and agent i reports P_i (recall that i does not trigger the stage to end). Now, note that since i does trigger any stage before k, he cannot be assigned to an object in $O_{<k}$ by deviating. In addition, by assumption, we know that if i reports P_i', he will still be assigned in stage k. If i gets an object in O_k, his assignment is the same as DA on the same subeconomy with i's preferences replaced by P_i'. Strategyproofness of DA on the subeconomy implies that the agent does not gain from misreporting his preferences. If i is assigned to an object in $O_{\geq k+1}$ (again by our conditioning) this must be worse than the object he obtains in stage k (which is in O_k) when he reports truthfully. Thus, i has no (ex-post) incentives to misreport. □

THEOREM 6. *When the mechanism is used, any symmetric Bayesian Nash equilibrium is asymptotically efficient and stable.*

PROOF. Fix any agent i, and consider the event/type that the agent prefers any top β-th object in each tier k to all objects in any tier $j > k$. We shall argue that for n sufficiently high, it is a strict best response for i to report truthfully among top β objects in each tier and above any objects in lower tier "in that event," given that all other agents follow the same strategies. By the Hoeffding's inequality, the probability of the joint event that "all" agents have the strict best response of adopting such a strategy goes to one as $n \to 1$, establishing the desired result.

Assume the other individuals employ the hypothesized behavior. Given the symmetry of objects within each tier, agent i will order the objects within each tier truthfully among them, conditional on listing them. Hence, it suffices to show that she will never put any object in tier $j \geq k$ ahead of some top β object of tier k. Suppose to the contrary that there exists an object o' from tier j that agent i ranks ahead of some top-β object in tier k, and without loss that no other from tier below k is listed ahead of o'. We shall consider the upper bound for the gains from the such a deviation, and show that the upper bound must negative for n sufficiently large.

Consider an event $\mathcal{E}$ that all objects in O_k are assigned in stage k, for each $k = 1, ..., K$. This event arises with probability no less than $1 - K\exp(-c\log^2(n))$ (we are taking the union bound on Pittel applied to each stage/tier). Given event $\mathcal{E}$, suppose agent i gets her first turn to make an offer in stage k, and after being rejected by several of top β objects in tier k, she considers applying to either the next-most preferred top β object say o in tier k or an object o' from tier $j > k$. In the former case, with probability at least of $1/\log(n)$, she will be accepted and thus assigned o. With the remaining probability, she will be assigned her top β object in tier $k+1$ with probability $1 - \frac{1}{x_{k+1}n}$. The payoff from this is at least

$$\frac{1}{\log(n)}u(u_k, \xi - \epsilon_n) + \left(1 - \frac{1}{\log(n)}\right)\left(1 - \frac{1}{x_{k+1}n}\right)u(u_{k+1}, \xi - \epsilon_n),$$

where $\xi - \epsilon_n$ is the realization of β-th highest order statistic. Suppose instead the agent deviates to list o' instead of o. In the same situation (i.e., having been rejected by all objects ahead of o or o'), she can get at most

$$u(u_j, \xi).$$

Hence, for some $\Delta, \Delta', \delta, \delta' > 0$, the gain from deviation can be bounded above by

$$
\begin{aligned}
&(1 - Ke^{-c\log^2(n)})\left[u(u_j, \bar{\xi}) - \frac{1}{\log(n)} u(u_k, \bar{\xi} - \epsilon_n) - \left(1 - \frac{1}{\log(n)}\right)\left(1 - \frac{1}{x_{k+1} n}\right) u(u_{k+1}, \bar{\xi} - \epsilon_n)\right] \\
&\quad + Ke^{-c\log^2(n)}\Delta \\
=&\left(1 - \frac{1}{\log(n)}\right)\left(u(u_j, \xi) - u(u_{k+1}, \bar{\xi} - \epsilon_n)\right) + \frac{1}{\log(n)}\left(u(u_j, \bar{\xi}) - u(u_k, \bar{\xi} - \epsilon_n)\right) + Ke^{-c\log^2(n)}\Delta' \\
&\qquad + \text{ higher order terms} \\
=&\left(u(u_j, \xi) - u(u_{k+1}, \bar{\xi} - \epsilon_n)\right) - \frac{1}{\log(n)}\left(u(u_k, \bar{\xi} - \epsilon_n) - u(u_{k+1}, \bar{\xi} - \epsilon_n)\right) + Ke^{-c\log^2(n)}\Delta' \\
&\qquad + \text{ higher order terms} \\
\leq&\frac{\log^2(n)}{n}\delta - \frac{1}{\log(n)}\delta' + Ke^{-c\log^2(n)}\Delta' + \text{ higher order terms} \\
=&-\frac{1}{\log(n)}\delta' + o(\log(n)) \\
<&0.
\end{aligned}
$$

The inequality follows since

$$
u(u_j, \xi) - u(u_{k+1}, \bar{\xi} - \epsilon_n) \leq u(u_{k+1}, \xi) - u(u_{k+1}, \bar{\xi} - \epsilon_n) \leq \sup_{\xi} u_2(u_{k+1}, \xi)\epsilon_n \leq \sup_{\xi, \xi'} u_2(u_{k+1}, \xi) h(\xi') \frac{3\log^2(n)}{n},
$$

where u_2 is the partial derivative of u with respect to the second argument, and h is the derivative of the inverse of Γ_I; the first inequality holds since $j \geq k+1$, and the last inequality follows since $\beta = 3\log^2(n)$, so β-th lowest order statistic ϵ_n cannot exceed $\sup_{\xi'} h(\xi') \frac{3\log^2(n)}{n}$. □

Remark 2. One can construct another mechanism which achieves asymptotic efficiency and asymptotic stability based on the famous Erdös-Renyi Theorem. The theorem states that a random bipartite graph across I and O where an edge is formed between $i \in I$ and $o \in O$ if and only only if $\xi > \bar{\xi} - \epsilon$ and $\eta > \bar{\eta} - \epsilon$ admits a perfect bipartite matching with probability approaching one as $n = |I| = |O|$ tends to ∞. There are famous algorithms – like the *augmenting path algorithm* – which find maximal matchings and hence that would find such perfect matchings whenever they exist.

One can thus imagine a mechanism in which agents and objects (more precisely their suppliers) report their idiosyncratic shocks, and a maximum matching is returned. Such a mechanism is asymptotically efficient and asymptotically stable. But mechanisms would not have good incentives properties. An agent will be reluctant to report the objects in lower tiers even though they will bring high idiosyncratic preferences. Indeed, if he expects that with significant probability, he will not to get any object in the highest tier, he will have incentives to claim that he enjoys high idiosyncractic payoffs with a large number of high tier objects and that all his idiosyncractic payoffs for the other tiers are low. It is very likely that there is a perfect matching even under this misreport and this will ensure him to get matched with high tier objects.

REMARK 3. **Knowledge of Objects' Preferences** The above analyses assume that the agents do not know the preferences of the objects. If they know their priorities at the objects, the incentives problem become more difficult. In particular, the agents will find it optimal to sacrifice some surplus in order to increase their chance of getting a higher tier object.

REMARK 4. **Many to one matching** It is unclear how to extend the above mechanism to many to one matching. The first issue is how to set β. Suppose first that the number of copies of each object type is bounded. Then, one would expect that β has the same order of magnitude. What if the number of copies increase in the same order of magnitude, or even in the higher order of magnitude? To consider the latter, imagine the continuum economy with a finite object types but with a continuum of seats for each object type (and a continuum of agents). To fix an idea, suppose there are 4 colleges, two of which are in tier 1 and two of which are in tier 2. Any tier 1 college is better than any tier 2 college for all students. There is a unit mass of students and each college has a quarter seats. If $\beta = 2$, then each student will adopt a cutoff strategy whereby she will rank the true top college in tier 1 as top, but ranks the second college in tier 1 as second if and only if its value is above a certain threshold.

A Main Proofs of the Paper

Proof of Lemma 1. Let Z_a be the number of balanced independent sets $\hat{V}_1 \times \hat{V}_2$ with $|\hat{V}_1| = |\hat{V}_2| = a$. We show that the probability of having a balanced independent set $\hat{V}_1 \times \hat{V}_2$ with $|\hat{V}_1| = |\hat{V}_2| = a \geq \gamma n$ is smaller than $\kappa \left(\frac{1}{n!}\right)^2$ where κ is some strictly positive constant. Observe that if $\hat{V}_1 \times \hat{V}_2$ is an independent set, then so is any subset of $\hat{V}_1 \times \hat{V}_2$. Thus, whenever there is a balanced independent set of size a, there must be an independent set of smaller size. Otherwise stated, $\Pr\{\exists a \geq \gamma n \text{ s.t. } Z_a \geq 1\} \leq \Pr\{Z_{\gamma n} \geq 1\}$. Thus, we can assume without loss of generality that $a = \gamma n$ and just show that $\Pr\{Z_a \geq 1\} \leq \kappa \left(\frac{1}{n!}\right)^2$. We have

$$\Pr\{Z_a \geq 1\} \leq \mathbb{E}(Z_a) = \binom{n}{a}\binom{m}{a}((1-p)^a)^a \leq \left(\frac{n^a}{a!}\right)\left(\frac{m^a}{a!}\right)((1-p)^a)^a .$$

The computation of $\mathbb{E}(Z_a)$ follows from the following argument. A set of vertices $A \cup B$ where $A \subseteq V_1$,and $B \subseteq V_2$ forms an independent set if there is no edge between every pair of vertices $v_1 \in A$ and $v_2 \in B$. Suppose that $|A| = |B| = a$. Since the probability of an edge is p, the probability that a given $A \cup B$ forms an independent set is $((1-p)^a)^a$. There are $\binom{n}{a}$ different ways of choosing a subset $A \subseteq V_1$ of size a and $\binom{m}{a}$ ways of choosing a subset $B \subseteq V_2$ of size a. Hence, the number of pairs of subsets (A, B) we are considering is $\binom{n}{a}\binom{m}{a}$.

Now, since $m = O(n)$, note that there is N such that for all $n > N$,

$$a = \gamma n \geq \frac{\log(n) + \log(m)}{\log(1/(1-p))} = \log_{(1-p)}(n^{-1}) + \log_{(1-p)}(m^{-1}).$$

Hence, for $n > N$,

$$(1-p)^a \leq (1-p)^{\frac{\log(n)+\log(m)}{\log(1/(1-p))}} = (1-p)^{\log_{1-p}(n^{-1})}(1-p)^{\log_{1-p}(m^{-1})} = n^{-1}m^{-1}.$$

Thus, for $n > N$, we get that $\binom{n}{a}\binom{m}{a}((1-p)^a)^a \leq \left(\frac{1}{a!}\right)^2$. Therefore, there must exist a strictly positive constant κ such that for all values of n, we have $\Pr\{Z_a \geq 1\} \leq \kappa \left(\frac{1}{a!}\right)^2 \leq \kappa \left(\frac{1}{n!}\right)^2$. □

Proof of Lemma 2. Fix any $\epsilon > 0$ and $\gamma > 0$. We first build a random graph on $I \cup O$ where an edge (i, o) is added if and only if $\xi_{i,o} > \bar{\xi} - \epsilon$.

Now choose any $\delta \in (0,1)$. For each $\mu \in \mathcal{M}^*$, define random sets $I^{\delta}_{\leq k}(\mu) := \{i \in I \,|\, f(i) \leq Q^n X^n_{\leq k}(1-\delta)\}$, $\overline{I}^{\delta}_{\leq k}(\mu) := \{i \in I^{\delta}_{\leq k} \,|\, U_i(SD^f(i)) \leq U(u^k, \overline{\xi} - \epsilon)\}$, and

$$\bar{O}^{\delta}_{\leq k}(\mu) := \{o \in O_{\leq k} \,|\, \exists i \in \mu^{-1}(o) \text{ s.t. } f(i) > Q^n X^n_{\leq k}(1-\delta)\},$$

which consists of objects in $O_{\leq k}$ assigned to the agents with serial order worse than $Q^n X^n_{\leq k}(1-\delta)$.

Then, the set $\overline{I}^{\delta}_{\leq k}(\mu) \cup \bar{O}^{\delta}_{\leq k}(\mu)$ must be an independent set. If not, there would exist an edge $(i,o) \in \overline{I}^{\delta}_{\leq k} \times \bar{O}^{\delta}_{\leq k}$. Then,

$$U_i(o) > U(u^k, \overline{\xi} - \epsilon) \geq U_i(SD^f(i))$$

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

$$f(i) \leq Q^n X^n_{\leq k}(1-\delta) < f(i'), \text{ for some } i' \in \mu^{-1}(o)$$

where the first inequality comes from the fact that $i \in I^{\delta}_{\leq k}(\mu)$ while the second from the fact that $o \in \bar{O}^{\delta}_{\leq k}(\mu)$. Thus, this means that when i becomes the dictator under SD^f, object o is still available, and the agent does not choose it. But $U_i(o) > U_i(SD^f(i))$ means that i chooses an object worse than o, which yields a contradiction.

In particular, for each $\mu \in \mathcal{M}^*$, $\overline{I}^{\delta}_{\leq k}(\mu) \cup \bar{O}^{\delta}_{\leq k}(\mu)$ contains a balanced independent set with size $\min\left\{|\overline{I}^{\delta}_{\leq k}(\mu)|, |\bar{O}^{\delta}_{\leq k}(\mu)|\right\}$. Since $|I| = n$ and $|O|$ is in the order of n, applying Lemma 1, we get that, for any $\tilde{\gamma} > 0$:

$$\Pr\left[\exists \mu \in \mathcal{M}^* \text{ s.t. } \min\left\{|\overline{I}^{\delta}_{\leq k}(\mu)|, |\bar{O}^{\delta}_{\leq k}(\mu)|\right\} \geq \tilde{\gamma} n\right] \leq \kappa_1 \left(\frac{1}{n!}\right)^2 \tag{1}$$

for some constant $\kappa_1 > 0$.

Since $\left|\bar{O}^{\delta}_{\leq k}(\mu)\right| \bar{q} \geq \sum_{o \in \bar{O}^{\delta}_{\leq k}(\mu)} q_o \geq \left\lfloor \delta Q^n X^n_{\leq k} \right\rfloor = \left\lfloor \delta n X^n_{\leq k} \right\rfloor$ for each $\mu \in \mathcal{M}^*$, and since $X^n_{\leq k} \to X_{\leq k} \geq Y_{\leq k} > 0$ as $n \to \infty$, one can find $\beta > 0$ and $N_1 \in \mathbb{N}$ such that for all $n > N_1$,

$\left|\bar{O}^{\delta}_{\leq k}(\mu)\right| \geq \beta n$ for each $\mu \in \mathcal{M}^*$.[13] Hence, for any $\gamma' > 0$ and for any $n > N_1$:

$$\begin{aligned}
\Pr\left[\exists \mu \in \mathcal{M}^* \text{ s.t. } |\overline{I}^{\delta}_{\leq k}(\mu)| \geq \gamma' n\right] &\leq \Pr\left[\exists \mu \in \mathcal{M}^* \text{ s.t. } |\overline{I}^{\delta}_{\leq k}(\mu)| \geq \min\{\gamma', \beta\} n\right] \\
&= \Pr\left[\exists \mu \in \mathcal{M}^* \text{ s.t. } \min\left\{|\overline{I}^{\delta}_{\leq k}(\mu)|, |\bar{O}^{\delta}_{\leq k}(\mu)|\right\} \geq \min\{\gamma', \beta\} n\right] \\
&\leq \kappa_2 \left(\frac{1}{n!}\right)^2,
\end{aligned}$$

for some κ_2, where the equality comes from the choice of β and N_1 while the last inequality holds by (1).

Since $|I_{\leq k}(\mu)|/n \to X_{\leq k}$, and $X_{\leq k} \geq Y_{\leq k} > 0$ as $n \to \infty$, we get that for any $c > 0$,

$$\Pr\left[\exists \mu \in \mathcal{M}^* \text{ s.t. } \frac{\left|\bar{I}^{\delta}_{\leq k}(\mu)\right|}{|I_{\leq k}(\mu)|} \geq c\right] \leq \kappa \left(\frac{1}{n!}\right)^2$$

for some constant $\kappa > 0$.

Finally, by construction, $|\bar{I}^{\delta}_{\leq k}(\mu)| \geq |\bar{I}_{\leq k}(\mu)| - \left\lfloor \delta Q^n X^n_{\leq k} \right\rfloor \geq |\bar{I}_{\leq k}(\mu)| - \delta Q^n X^n_{\leq k}$. Since $Q^n X^n_{\leq k} = |I_{\leq k}(\mu)|$, we get that

$$\frac{\left|\bar{I}^{\delta}_{\leq k}(\mu)\right|}{|I_{\leq k}(\mu)|} \geq \frac{\left|\bar{I}_{\leq k}(\mu)\right|}{|I_{\leq k}(\mu)|} - \delta$$

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

$$\Pr\left[\exists \mu \in \mathcal{M}^* \text{ s.t. } \frac{\left|\bar{I}_{\leq k}(\mu)\right|}{|I_{\leq k}(\mu)|} \geq c + \delta\right] \leq \Pr\left[\exists \mu \in \mathcal{M}^* \text{ s.t. } \frac{\left|\bar{I}^{\delta}_{\leq k}(\mu)\right|}{|I_{\leq k}(\mu)|} \geq c\right] \leq \kappa \left(\frac{1}{(n)!}\right)^2.$$

Set δ and c such that $\delta + c = \gamma$. Then, we have

$$\Pr\left[\frac{|\bar{I}_{\leq k}(\mu)|}{|I_{\leq k}(\mu)|} \geq \gamma\right] \leq \kappa \left(\frac{1}{n!}\right)^2$$

as was to be shown. □

Proof of Theorem 1. To prove the statement, we will show that the payoff distributions induced by Pareto efficient mechanisms converge to $\bar{F}^*$ in the sense defined earlier.

[13] Here we use the assumption that $\bar{q}$ does not increase in n. If it were to depend on n and we further assume that it is $O(n/log(n))$, one can check that there is $\beta > 0$ and $N_1 \in \mathbb{N}$ s.t. for all $n > N_1$, $\left|\bar{O}^{\delta}_{\leq k}(\mu)\right| \geq \beta log(n)$ for each $\mu \in \mathcal{M}^*$. Using the alternative version of Lemma 1 by Dawande, Keskinocak, Swaminathan, and Tayur (2001) mentioned in Remark 1, one can show that Theorem 1 below goes through.

Fix any $\epsilon' > 0$ and $\upsilon' > 0$. We shall show that there exists $N \in \mathbb{N}$ such that for all $n > N$,

$$\Pr\left[\sup_{\mu\in\hat{\mathcal{M}}^*}\sup_{z\in Z}\inf_{\hat{z}\in[z-\epsilon',z+\epsilon']}\left\{|\bar{F}^{\mu}(\hat{z})-\bar{F}^*(z)|\right\}\geq\epsilon'\right]<\upsilon', \tag{2}$$

where $\bar{F}^*$ and $\bar{F}^{\mu}$ are respectively the DDF of the payoff induced by the limit utilitarian upper bound and the DDF induced by mechanism μ in $\mathcal{M}^*$. We later show that this is sufficient for the proof.

To this end, we partition the common value space $[0,1]$ into intervals $\cup_{k=1}^{K}(u^{k+1},u^k]$, where $\sup(supp(Y)) =: u^1 > u^2 > ... > u^K = 0$ are such that $\sup_k \sup_{u,u'\in(u^{k+1},u^k]}|X(u) - X^n(u')| < \frac{\epsilon'}{2}$, for any $n > \hat{N}$ for some $\hat{N} \in \mathbb{N}$. Such a partition exists since $X^n \to X$ uniformly and since one can select all points of discontinuity of X to be a subset of the threshold values for the partition. Define as before $O_{\leq k} = \{o \in O | u_o \geq u^k\}$ and $Y^n_{\leq k} := Y^n(u^k)$, $X^n_{\leq k} := X^n(u^k)$.

The partition induces a corresponding partition of the payoff space $Z := [0, U(u^1,\overline{\xi})]$ into intervals $Z_k := (U(u^{k+1},\overline{\xi}), U(u^k,\overline{\xi})]$, $k = 1, ..., K-1$. Next, let $\epsilon > 0$ be such that $U(u^k,\overline{\xi}-\epsilon) > \max\{U(u^k,\overline{\xi})-\epsilon', U(u^{k+1},\overline{\xi})\}$ for all $k = 1, ..., K-1$. Then, for each Pareto efficient mechanism $\mu \in \mathcal{M}^*$ and for each $z \in Z_k$, $k \in \{1, ..., K\}$, let $z' := \min\{z, U(u^k,\overline{\xi}-\epsilon)\}$. Clearly, given our choice of ϵ, we have that $z' \in Z_k$. In addition, given this choice of ϵ, $z' \in [z-\epsilon', z+\epsilon']$. Indeed, this is trivially true if $z \leq U(u^k,\overline{\xi}-\epsilon)$ (in which case $z' = z$) and if $z > U(u^k,\overline{\xi}-\epsilon)(= z')$, we have that $z - \epsilon' < U(u^k,\overline{\xi}) - \epsilon' < U(u^k,\overline{\xi}-\epsilon) = z'(< z)$ where the first inequality comes from the fact that $z \in Z_k$ while the second is by the choice of ϵ. Define

$$J^{\mu}(z) = \{i \in I | U_i(\mu(i)) \geq z\}.$$

be the set of agents enjoying payoff of at least z under matching μ. Let u_z be such that $U(u_z,\overline{\xi}) = z$. Clearly, each agent in $J^{\mu}(z)$ must be obtaining an object with common value no less than u_z. This means that $|J^{\mu}(z)| \leq Q^n X^n(u)$ for all $\mu \in \mathcal{M}^*$.

By definition, for each z, $\bar{F}^*(z) = X(u_z)$. Now, because $Q^n = n = |I|$, for each z,

$\frac{|J^\mu(z')|}{|I|} - \bar{F}^*(z) \leq X^n(u_{z'}) - X(u_z)$. Then, for all $n > \hat{N}$,

$$
\begin{aligned}
&\Pr\left[\sup_{\mu\in\mathcal{M}^*}\sup_{z\in Z_k}\left(\frac{|J^\mu(z')|}{|I|} - \bar{F}^*(z)\right) \geq \epsilon'\right] \\
\leq &\Pr\left[\sup_{\mu\in\mathcal{M}^*}\sup_{z\in Z_k}\left(X^n(u_{z'}) - X(u_z)\right) \geq \epsilon'\right] \\
\leq &\Pr\left[\sup_{\mu\in\mathcal{M}^*}\sup_{u,u'\in(u^{k+1},u^k]}\left(X^n(u') - X(u)\right) \geq \epsilon'\right] = 0,
\end{aligned} \tag{3}
$$

where the last inequality comes from the fact that by definition of z and z', both u_z and u'_z are in $(u^{k+1}, u^k]$ while the equality to 0 is from the definition of $\hat{N}$.

For each k, let $I_{\leq k}(\mu) := \left\{i \in I | f(i) \leq Q^n X^n_{\leq k}\right\}$ and $\bar{I}_{\leq k}(\mu) := \left\{i \in I_{\leq k} | U_i(SD^f(i)) \leq U(u^k, \bar{\xi} - \epsilon)\right\}$, where SD^f is the SD rule implementing μ. Note that, for any z, any agent in $I_{\leq k}(\mu) \setminus \bar{I}_{\leq k}(\mu)$ must be enjoying payoff at least of $U(u^k, \bar{\xi} - \epsilon) \geq z'$, so such an agent must belong to $J^\mu(z')$. In other words, for any z, $J^\mu(z') \supset I_{\leq k}(\mu) \setminus \bar{I}_{\leq k}(\mu)$. Hence, there exists $N_k \in \mathbb{N}$, with $N_k \geq \hat{N}$, such that for all $n > N_k$,

$$
\begin{aligned}
&\Pr\left[\left(\sup_{\mu\in\mathcal{M}^*}\sup_{z\in Z_k} -\frac{|J^\mu(z')|}{|I|}\right) + \bar{F}^*(z) \geq \epsilon'\right] \\
\leq &\Pr\left[\left\{\sup_{\mu\in\mathcal{M}^*} -\left(\frac{|I_{\leq k}(\mu)| - |\bar{I}_{\leq k}(\mu)|}{|I_{\leq k}(\mu)|}\right)\left(\frac{|I_{\leq k}(\mu)|}{|I|}\right)\right\} + X(u) \geq \epsilon'\right] \\
= &\Pr\left[\left\{\sup_{\mu\in\mathcal{M}^*} -\left(1 - \frac{|\bar{I}_{\leq k}(\mu)|}{|I_{\leq k}(\mu)|}\right) X^n(u^k)\right\} + X(u) \geq \epsilon'\right] \\
\leq &\Pr\left[\left(\sup_{\mu\in\mathcal{M}^*} \frac{|\bar{I}_{\leq k}(\mu)|}{|I_{\leq k}(\mu)|}\right) + \left|X(u) - X^n(u^k)\right| \geq \epsilon'\right] \\
< &\upsilon'/K,
\end{aligned} \tag{4}
$$

where the last inequality follows from Lemma 2 (with $\gamma = \epsilon'/2$) and an appropriate choice of N_k.

Combining (3) and (4), we get that for each $k = 1, \dots, K$, and $n > N_k$,

$$
\Pr\left[\sup_{\mu\in\mathcal{M}^*}\sup_{z\in Z_k} \left|\frac{|J^\mu(z')|}{|I|} - \bar{F}^*(z)\right| \geq \epsilon'\right] < \upsilon'/K,
$$

Since $\bar{F}^\mu(z') = \frac{|J^\mu(z')|}{|I|}$, we obtain that for all $n > \max_k N_k$,

$$
\begin{aligned}
&\Pr\left[\sup_{\mu\in\hat{\mathcal{M}}^*} L(\bar{F}^\mu, \bar{F}^*) \geq \epsilon'\right] \\
\leq &\Pr\left[\sup_{\mu\in\hat{\mathcal{M}}^*} \left\{\sup_{z\in Z} \inf_{\hat{z}\in[z-\epsilon', z+\epsilon']} |\bar{F}^\mu(\hat{z}) - \bar{F}^*(z)|\right\} \geq \epsilon'\right] \\
\leq &\Pr\left[\sup_{\mu\in\hat{\mathcal{M}}^*} \left\{\sup_{z\in Z} |\bar{F}^\mu(z') - \bar{F}^*(z)|\right\} \geq \epsilon'\right] \\
\leq &\sum_{k=1}^{K} \Pr\left[\sup_{\mu\in\hat{\mathcal{M}}^*} \left\{\sup_{z\in Z_k} |\bar{F}^\mu(z') - \bar{F}^*(z)|\right\} \geq \epsilon'\right] \\
< &K\upsilon'/K = \upsilon',
\end{aligned}
$$

where the first inequality holds since if $L(\bar{F}^\mu, \bar{F}^*) \geq \epsilon'$, then there exists z such that $\inf_{\hat{z}\in[z-\epsilon', z+\epsilon']} |\bar{F}^\mu(\hat{z}) - \bar{F}^*(z)| \geq \epsilon'$; the second inequality holds since $z' \in [z-\epsilon', z+\epsilon']$; and the third follows from Boole's inequality. This completes the proof. □

Proof of Lemma 4. Note first that since there are large common value differences, if $o \in \hat{O} \subset O_1$ and $j \in I_2$, it must be that o does not point to j in the cycle to which o belongs under TTC (otherwise, if j is part of the cycle in which o is cleared, since $o \in O_1$, this means that j must be pointing to an object in O_1 when she is cleared, which is a contradiction with $j \in I_2$). Note also that j is still in the market when o is cleared.

Define $E_1 := \{\eta_{jo} \geq \eta_{TTC(o)o}\} \wedge \{o \in \hat{O}\} \wedge \{j \in I_2\}$ and $E_2 := \{\eta_{jo} \leq \eta_{TTC(o)o}\} \wedge \{o \in \hat{O}\} \wedge \{j \in I_2\}$. We first show that $\Pr E_1 = \Pr E_2$.

Assume that under the realizations $\boldsymbol{\xi} := (\xi_{io})_{io}$ and $\boldsymbol{\eta} := (\eta_{io})_{io}$ event E_1 is true. Define $\hat{\boldsymbol{\eta}} := (\hat{\eta}_{io})_{io}$ where $\hat{\eta}_{jo} := \eta_{TTC(o)o}$ and $\hat{\eta}_{TTC(o)o} := \eta_{jo}$ – while $\hat{\boldsymbol{\eta}}$ and $\boldsymbol{\eta}$ coincide otherwise. It is easily checked that under the realizations $\boldsymbol{\xi}$ and $\hat{\boldsymbol{\eta}}$, event E_2 is true. Indeed, that $\{\hat{\eta}_{jo} \leq \hat{\eta}_{TTC(o)o}\}$ holds true is trivial. Now, since, as we already claimed, under the realizations $\boldsymbol{\xi}$ and $\boldsymbol{\eta}$, j and $TTC(o)$ are never pointed by o, when j and $TTC(o)$ are switched in o's priorities, by definition of TTC, o still belongs to the same cycle and, hence, TTC runs exactly in the same way. This shows that $\{o \in \hat{O}\} \wedge \{j \in I_2\}$ also holds true under the realizations $\boldsymbol{\xi}$ and $\hat{\boldsymbol{\eta}}$,

Given that $\Pr(\boldsymbol{\xi}, \boldsymbol{\eta}) = \Pr(\boldsymbol{\xi}, \hat{\boldsymbol{\eta}})$, we get that $\Pr E_1 = \Pr E_2$.

Next, let $E_\epsilon := \{\eta_{jo} \geq \eta_{TTC(o)o} + \epsilon\}$. Note that

$$\cup_{\epsilon>0} E_\epsilon = \{\eta_{jo} > \eta_{TTC(o)o}\} =: E.$$

Since the distribution $\Pr[\cdot]$ of η_{io} has no atom, $\Pr\left[\cdot \middle| o \in \hat{O}, j \in I_2\right]$ has no atom as well ($\Pr(\eta_{jo} = \eta) = 0 \Rightarrow \Pr(\eta_{jo} = \eta \middle| o \in \hat{O}, j \in I_2) = 0$). Thus, we must have

$$\Pr\left[E \middle| o \in \hat{O}, j \in I_2\right] = \Pr\left[\{\eta_{jo} \geq \eta_{TTC(o)o}\} \middle| o \in \hat{O}, j \in I_2\right] = \frac{1}{2}.$$

Since E_ϵ is increasing when ϵ decreases, combining the above, we get[14]

$$\lim_{\epsilon \to 0} \Pr\left[E_\epsilon \middle| o \in \hat{O}, j \in I_2\right] = \Pr\left[\cup_{\epsilon>0} E_\epsilon \middle| o \in \hat{O}, j \in I_2\right] = \Pr\left[E \middle| o \in \hat{O}, j \in I_2\right] = \frac{1}{2}.$$

Thus, one can fix $\delta \in (0, 1/2)$ arbitrarily close to 0 and find $\varepsilon > 0$ so that for any $\epsilon \in (0, \varepsilon)$, $\Pr\left[E_\epsilon \middle| o \in \hat{O}, j \in I_2\right] \geq \frac{1}{2} - \delta > 0$. □

Proof of Corollary 3. Then, for any ϵ sufficiently small, we have $\zeta > 0$ and $N > 0$ such that

$$\begin{aligned}
\mathbb{E}\left[|\hat{I}_2(o)| \middle| o \in \hat{O}\right] &= \mathbb{E}\left[\sum_{i \in I_2} \mathbf{1}_{\{\eta_{io} > \eta_{TTC(o)o} + \epsilon\}} \middle| o \in \hat{O}\right] \\
&= \mathbb{E}_{I_2}\left(\mathbb{E}\left[\sum_{i \in I_2} \mathbf{1}_{\{\eta_{io} > \eta_{TTC(o)o} + \epsilon\}} \middle| o \in \hat{O}, I_2\right]\right) \\
&= \mathbb{E}_{I_2}\left(\sum_{i \in I_2} \mathbb{E}\left[\mathbf{1}_{\{\eta_{io} > \eta_{TTC(o)o} + \epsilon\}} \middle| o \in \hat{O}, I_2, i \in I_2\right]\right) \\
&= \mathbb{E}_{I_2}\left(x_2 n \mathbb{E}\left[\mathbf{1}_{\{\eta_{io} > \eta_{TTC(o)o} + \epsilon\}} \middle| o \in \hat{O}, I_2, i \in I_2\right]\right) \\
&= x_2 n \left(\mathbb{E}\left[\mathbf{1}_{\{\eta_{io} > \eta_{TTC(o)o} + \epsilon\}} \middle| o \in \hat{O}, i \in I_2\right]\right) \\
&= x_2 n \Pr(\eta_{io} > \eta_{TTC(o)o} + \epsilon \middle| o \in \hat{O}, i \in I_2) \\
&= x_2 n \Pr(\eta_{io} \geq \eta_{TTC(o)o} + \epsilon \middle| o \in \hat{O}, i \in I_2) \\
&\geq x_2 \zeta n,
\end{aligned}$$

for all $n > N$. □

[14] Recall the following property. Let $\{E_n\}_n$ be an increasing sequence of events. Let $E := \cup_n E_n$ be the limit of $\{E_n\}_n$. Then: $\Pr(E) = \lim_{n\to\infty} \Pr(E_n)$.

Proof of Theorem 3. Since $U(u_1^0, 0) > U(u_2^0, \bar{\xi}))$, all objects in O_1 are assigned before any agent starts applying to objects in O_2. Hence, the assignment achieved by individuals matched to objects in O_1 is the same as the one obtained when we run DA in the submarket with individuals in I and objects in O_1. This submarket has been studied by Ashlagi, Kanoria, and Leshno (2013). Their Corollary 2.3. shows that there exists a constant $c > 0$ such that, with probability converging to one as n grows, the average rank of individuals matched to objects in O_1 is at least n/c. Thus, this means that there exists $\epsilon' > 0, \upsilon' > 0, \gamma' > 0$ such that for all $n > N'$ for some $N' > 0$,

$$\Pr\left\{\frac{|\tilde{I}_\epsilon|}{|I|} \geq \gamma'\right\} \geq \upsilon',$$

where $\tilde{I}_{\epsilon'} := \{i \in I | DA(i) \in O_1, U_i(DA(i)) \leq U(u_1, \bar{\xi} - \epsilon')\}$ is the set of agents assigned to objects in O_1 but receive payoffs bounded above by $U(u_1, \bar{\xi} - \epsilon')$.

Now consider a matching mechanism that first runs DA and then runs a Shapley-Scarf TTC afterwards, namely the TTC with the DA assignments serving as the initial endowments for the agents. This mechanism μ clearly Pareto dominates DA. In particular, if $DA(i) \in O_1$, then $\mu(i) \in O_1$. For any ϵ'', let

$$\check{I}_{\epsilon''} := \{i \in I | \mu(i) \in O_1, U_i(DA(i)) \geq U(u_1, \bar{\xi} - \epsilon'')\},$$

be those agents attain at least the payoff of $U(u_1, \bar{\xi} - \epsilon'')$. By Lemma 2, we have for any ϵ'', γ'' and υ'', such that

$$\Pr\left\{\frac{|\check{I}_{\epsilon''}|}{|I|} > \gamma''\right\} < \upsilon'',$$

for all $n > N''$ for some $N'' > 0$.

Now set ϵ', ϵ'' such that $\epsilon = \epsilon' - \epsilon'' > 0$, γ', γ'' such that $\gamma := \gamma' - \gamma'' > 0$, and υ', υ'' such that $\upsilon := \upsilon' - \upsilon'' > 0$. Observe that $I_\epsilon(\mu|DA) \supset \tilde{I}_{\epsilon'} \setminus \check{I}_{\epsilon''}$, so $|I_\epsilon(\mu|DA)| \geq |\tilde{I}_{\epsilon'}| - |\check{I}_{\epsilon''}|$.

It then follows that for all $n > N := \max\{N', N''\}$,

$$\begin{aligned}
\Pr\left\{\frac{|I_\epsilon(\mu|DA)|}{|I|} \geq \gamma\right\} &\geq \Pr\left\{\frac{|\tilde{I}_{\epsilon'}|}{|I|} - \frac{|\check{I}_{\epsilon''}|}{|I|} \geq \gamma\right\} \\
&\geq \Pr\left\{\frac{|\tilde{I}_{\epsilon'}|}{|I|} \geq \gamma' \text{ and } \frac{|\check{I}_{\epsilon''}|}{|I|} \leq \gamma''\right\} \\
&\geq \Pr\left\{\frac{|\tilde{I}_{\epsilon'}|}{|I|} \geq \gamma'\right\} - \Pr\left\{\frac{|\check{I}_{\epsilon''}|}{|I|} > \gamma''\right\} \\
&\geq \upsilon' - \upsilon = \upsilon.
\end{aligned}$$

□

B Analysis of TTC

In this section, we provide an analysis of TTC in our random environment. Our final goal is to prove Lemma 3. For our purpose, it is sufficient to consider the TTC assignment arising from the market consisting of the agents in I and the objects in top tier O_1 (recall that, irrespective of the realizations of the idiosyncratic values, all agents consider objects in O_1 better than objects in O_2. Hence, we shall simply consider an **unbalanced market** consisting of a set I of agents and a set O of objects such that (1) *the preferences of each side with respect to the other side are drawn i.i.d. uniformly,* and (2) *both* $|O|$ *and* $|I| - |O|$ *increase in the order of* $|O|$*, as the market size* $|O|$ *grows to infinity.* The analysis of this market requires a preliminary result on bipartite random mapping.

B.1 Preliminaries

Here, we develop a couple of preliminary results that we shall later invoke. Through, we shall consider two finite sets I and O, with cardinalities $|I| = n, |O| = o$.

Number of Spanning Rooted Forests. A **rooted tree** is a connected directed bipartite digraph where all vertices have out-degree 1 except the root which has out-degree 0.[15]

[15]Sometimes, a tree is defined as an acyclic undirected connected graph. In such a case, a tree is rooted when we name one of its vertex a "root". Starting from such a rooted tree, if all edges now have a direction

A **rooted forest** is a bipartite graph which consists of a collection of disjoint rooted trees. A **spanning rooted forest over** $I \cup O$ is a forest with vertices $I \cup O$ under which all vertices are part of at least one edge. From now on, a spanning forest will be understood as being over $I \cup O$. We will be using the following result.

Lemma 6 (Jin and Liu (2004)). *Let $V_1 \subset I$ and $V_2 \subset O$ where $|V_1| = \ell$ and $|V_2| = k$. There are $o^{n-\ell-1}n^{o-k-1}(kn + \ell o - k\ell)$ spanning rooted forests having in total $\ell + k$ roots, ℓ in V_1 and k in V_2.*

Random Bipartite Mapping. We now consider arbitrary mappings, $g : I \to O$ and $h : O \to I$, defined over our finite sets I and O. Note that such mapping naturally induces bipartite digraphs with vertices $I \cup O$ and directed edges with the number of outgoing edges equal to one for each vertex. In this digraph, $i \in I$ points to $g(i) \in O$ while $o \in O$ points to $h(o) \in I$. A **random bipartite mapping** selects a composite map $g \circ h$ uniformly from a set $\mathcal{G} \times \mathcal{H} = O^I \times I^O$. Note that a random bipartite mapping induces a random bipartite digraph consisting of vertices $I \cup O$ and directed edges with the number of outgoing edges equal to one for each vertex. We say that a vertex in a digraph is **cyclic** if it is in a cycle. The following lemma states the number of cyclic vertices in a random bipartite digraph induced by a random bipartite mapping.

Lemma 7 (Jaworski (1985), Corollary 3). *The number q of the cyclic vertices in a random bipartite digraph induced by a random bipartite mapping $g : I \to O$ and $h : O \to I$ has an expected value of*

$$\mathbb{E}[q] := 2 \sum_{i=1}^{o} \frac{(o)_i (n)_i}{o^i n^i},$$

where $(x)_j := x(x-1)\cdots(x-j-1)$.

For the next result, consider agents I' and objects O' such that $|I'| = |O'| = m > 0$. We say a mapping $f = g \circ h$ is a **bipartite bijection**, if $g : I' \to O'$ and $h : O' \to I'$ are both bijections. Note that a bipartite bijection consists of disjoint cycles. A **random bipartite bijection** is a (uniform) random selection of a bipartite bijection from the set of all bipartite bijections. The following result will prove useful for a later analysis.

leading toward the root, then the out-degree of any vertex (except the root) is 1. So the two definitions are actually equivalent.

Lemma 8. *Fix sets I' and O' with $|I'| = |O'| = m > 0$, and a subset $K \subset I' \cup O'$, with its $a \geq 0$ vertices in I' and its $b \geq 0$ vertices in O'. The probability that each cycle in a random bipartite bijection contains at least one vertex from K is*

$$\frac{a+b}{m} - \frac{ab}{m^2}.$$

Proof. We shall invoke Lovasz (1979) Exercise 3.6, which establishes that the probability that each cycle of a random permutation[16] of a finite set $|X|$ contains at least one element of a set $K \subset X$ is $|K|/|X|$.

To this end, observe first that a bipartite bijection $f = g \circ h$ induces a permutation of set I'. Thus, a random bipartite bijection induces a random permutation of set I'. The probability that each cycle of the randomly selected bipartite bijection contains at least one vertex in K is identical to the probability that each cycle of the induced permutation of I' contains at least one of $a + X$ vertices, where X is the (random) number of vertices in $I' \setminus K$ that point to $K \cap O'$. For any $\max\{b-a, 0\} \leq x \leq \min\{m-a, b\}$,

$$\Pr\{X = x\} = \frac{\binom{a}{b-x}\binom{m-a}{x}}{\binom{m}{b}}.$$

The above formula can be understood as follows. $\binom{a}{b-x}\binom{m-a}{x}$ is the number of ways one can choose $b-x$ nodes in $K \cap I'$ and x nodes in $I' \setminus K$. Thus, the total number of bipartite bijection having exactly x vertices in $I' \setminus K$ that point to $K \cap O'$ is $\binom{a}{b-x}\binom{m-a}{x}w$ where w is the total number of bipartite bijections where our given $b-x$ nodes in $K \cap I'$ point to nodes in $K \cap O'$ and the given x nodes in $I' \setminus K$ point to nodes in $K \cap O'$. Note that w is equal to the number of bipartite bijections where our given b nodes point to nodes in $K \cap O'$. Hence, the total number of bipartite bijections having b nodes in I' pointing to $K \cap O'$ is $\binom{m}{b}w$. Thus, we get the above formula.

[16] Formally, a permutation of X is a bijection $f : X \to X$. A random permutation chooses randomly a permutation f in the set of all possible permutations.

Applying the earlier result, the desired probability is

$$\begin{aligned}
&\sum_{x=\max\{b-a,0\}}^{\min\{m-a,b\}} \Pr\{X=x\}\frac{a+x}{m} \\
=&\frac{a}{m}+\sum_{x=\max\{b-a,0\}}^{\min\{m-a,b\}} \Pr\{X=x\}\frac{x}{m} \\
=&\frac{a}{m}+\sum_{x=\max\{b-a,0\}}^{\min\{m-a,b\}} \frac{\binom{a}{b-x}\binom{m-a}{x}}{\binom{m}{b}}\left(\frac{x}{m}\right) \\
=&\frac{a}{m}+\left(\frac{m-a}{m\binom{m}{b}}\right)\sum_{x=\max\{b-a,1\}}^{\min\{m-a,b\}} \binom{a}{b-x}\binom{m-a-1}{x-1} \\
=&\frac{a}{m}+\left(\frac{m-a}{m\binom{m}{b}}\right)\binom{m-1}{b-1} \\
=&\frac{a}{m}+\frac{b(m-a)}{m^2} \\
=&\frac{a+b}{m}-\frac{ab}{m^2},
\end{aligned}$$

where the fourth equality follows from the Vandermonde's identity. □

B.2 Markov Chain Property of TTC

Again consider a TTC in an unbalanced market with agents I and objects O. As is well known, TTC assigns agents to objects via cycles formed recursively in multiple rounds. We shall call a cycle of length 2—an agent points to an object, which in turn points to the original agent—a **short-cycle**. Any cycles of length greater than 2 shall be called **long-cycles**. Our aim is to prove that the number of the agents assigned via long-cycles in TTC grows in the same order n as the size of the market n grows. The difficulty with proving this result stems from the fact that the preferences of the agents and objects remaining after the first round of TTC need not be uniform, with their distributions affected nontrivially by the realized event of the first round TTC, and the nature of the conditioning is difficult to analyze in the large market.[17] Our approach is to prove that, even though the exact

[17]To understand why there is a conditioning issue here, let us consider an example where we have three individuals $\{1,2,3\}$ and three objects $\{o_1,o_2,o_3\}$ and that both preferences and priorities are drawn iid

composition of cycles are subject to the conditioning issue, the number of agents assigned in each round follows a Markov chain, and is thus free from the conditioning issue. We then combine this observation with the bound we shall establish on the number of agents assigned via short-cycles, to produce a desired result.

We shall begin with the Markov Chain result. This result parallels the corresponding result by Frieze and Pittel (1995) on the Shapley-Scarf version of TTC. The difference between the two versions of TTC is not trivial, so their proofs do not carry over.

Theorem 7. *Suppose any round of TTC begins with n agents and o objects remaining in the market. Then, the probability that there are $m \leq \min\{o, n\}$ agents assigned at the end of that round is*

$$p_{n,o;m} = \left(\frac{m}{(on)^{m+1}}\right)\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right)(o+n-m).$$

Thus, denoting n_i and o_i the number of individuals and objects remaining in the market at any round i, the random sequence (n_i, o_i) is a Markov chain.

B.2.1 Proof of Theorem 7

We begin by noting that TTC induces a random sequence of spanning rooted forests. Indeed, one could see the begining of the first round of TTC as a stituation where we have the trivial forest consisting of $|I| + |O|$ trees with isolated vertices. Within this step each vertex in I will randomly point to a vertex in O and each vertex in O will randomly point to a vertex in I. Note that once we delete the realized cycles, we again get a spanning rooted forest. So we can think again of the begining of the second round of TTC as a stituation where we start with a spaning rooted forest where the agents and objects remaining from

uniformly. Condition on the event E that 3 and o_3 are matched together in the first round of TTC. If preferences and priorities were still drawn iid uniformly conditional on E, then for any realized priorities of the remaining objects, the probability that 1 ranks o_1 above o_2 would be $1/2$. To get an intuition of why this is not true, let us consider the event F under which o_1 ranks 1 above 2 while o_2 does this opposite. Thus, given E and F, it is not very likely that 1 had ranked o_1 first in the first round of TTC (given F, with probability $1/2$, o_1 initially ranked 1 first in which case, given E, 1 cannot be ranking o_1 first). Given F and E, it is more likely that 1 had ranked o_2 first in the first round of TTC. This is the reason why, given E and F, the probability that 1 ranks o_1 above o_2 is smaller than the the probability that 1 ranks o_2 above o_1 (and so smaller than $1/2$).

the first round form this spanning rooted forest, where the roots consist of those agents and objects that had pointed to the entities that were cleared via cycles. Here again objects that are roots randomly point to a remaining individual and individuals that are roots randomly point to a remaining object. Once cycles are cleared we again obtain a forest and the process goes on like this. Formally, the random sequence of forests, $F_1, F_2, ...$ is defined as follows. First, we let F_1 be a trivial unique forest consisting of $|I|+|O|$ trees with isolated vertices, forming their own roots. For any $i = 2, ...,$ we first create a random directed edge from each root of F_{i-1} to a vertex on the other side, and then delete the resulting cycles (these are the agents and objects assigned in around $i-1$) and F_i is defined to be the resulting rooted forest. We let (N_i, k_i) denote the vertex set and the number of trees in the rooted forest F_i, and let $\mathcal{F}_{N_i,k_i}$ denote the set of all rooted forests with N_i as the vertex set and k_i as the number of trees.

LEMMA 9. *Given $(N_j, k_j), j = 1, ..., i$, every (rooted) forest of $\mathcal{F}_{N_i,k_i}$ is equally likely.*

PROOF. We prove this result by induction on i. Since for $i = 1$, by construction, the trivial forest is the unique forest which can occur, this is trivially true for $i = 1$. Fix $i \geq 2$, and assume our statement is true for $i-1$. Fix $N_i = I_i \cup O_i \subset N_{i+1} = I_{i+1} \cup O_{i+1}$, $\kappa = k_i$, and $\lambda = k_{i+1}$. We start by showing that each forest $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$ arises from the same number of pairs (F', ϕ) where $F' \in \mathcal{F}_{N_i,k_i}$ and ϕ maps the roots of F' in I_i to its vertices in O_i as well as the roots of F' in O_i to its vertices in I_i. Given F, we can construct all such pairs by choosing a quadruplet (a, b, c, d) of four non-negative integers with $a+b+c+d = \kappa$,

1. choosing c old roots from I_{i+1}, and similarly, d old roots from O_{i+1},

2. choosing a old roots from $I_i \backslash I_{i+1}$ and similarly, b old roots from $O_i \backslash O_{i+1}$,

3. choosing a partition into cycles of $N_i \backslash N_{i+1}$, each cycle of which contains at least one old root from (2),

4. choosing a mapping of the λ new roots to $N_i \backslash N_{i+1}$ satisfying the bipartite graph constraint.

Note that $\alpha(n, o, k; m, \lambda)$ does not depend on $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$. For future use, we compute the exact value of $\alpha(n, o, k; m, \lambda)$. Letting $n := |I_i|$, $o := |O_i|$, and $|I_i| - |I_{i+1}| = |O_i| -$

$|O_{i+1}| = m$, the number of choices is

$$\begin{aligned}
&\alpha(n,o,k;m,\lambda) \\
:=&\sum_{a+b+c+d=\kappa}\binom{n-m}{c}\binom{o-m}{d}\binom{m}{a}\binom{m}{b}\left(\frac{a+b}{m}-\frac{ab}{m^2}\right)(m!)^2m^\lambda \\
=&(m!)^2m^\lambda\times\left(\sum_{a+b+c+d=\kappa}\binom{n-m}{c}\binom{o-m}{d}\binom{m-1}{a-1}\binom{m}{b}\right. \\
&\left.+\sum_{a+b+c+d=\kappa}\binom{n-m}{c}\binom{o-m}{d}\binom{m}{a}\binom{m-1}{b-1}-\sum_{a+b+c+d=\kappa}\binom{n-m}{c}\binom{o-m}{d}\binom{m-1}{a-1}\binom{m-1}{b-1}\right) \\
=&(m!)^2m^\lambda\left(2\binom{n+o-1}{\kappa-1}-\binom{n+o-2}{\kappa-2}\right).
\end{aligned}$$

The first equality follows from Lemma 8, along with the fact that there are $(m!)^2$ possible bipartite bijections between $|I_i|-|I_{i+1}|$ agents and $|O_i|-|O_{i+1}|$ objects, and the fact that there are $m^\lambda = m^{\lambda_1}m^{\lambda_2}$ ways in which new roots (λ_1 in I_{i+1} and λ_2 in O_{i+1}) could have pointed to $2m$ cyclic vertices (m on the individuals' side and m on the objects' side), and the last equality follows from Vandermonde's identity.

The rest of the proof mimics Frieze and Pittel (1995). If $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$ then the inductive assumption and the Markov property of $\{F_j\}$ implies that

$$P(F_{i+1}=F|(N_1,k_1),...,(N_i,k_i)) = \frac{1}{|\mathcal{F}_{N_i,k_i}|}\sum_{F'\in\mathcal{F}_{N_i,k_i}}P(F_{i+1}=F|F_i=F').$$

Now, let $\phi_i = (\phi_i^I,\phi_i^O)$ where ϕ_i^I is the random mapping from the roots of F_i in I_i to O_i and ϕ_i^O is the random mapping from the roots of F_i in O_i to I_i. Let $\phi = (\phi^I,\phi^O)$ be a generic mapping of that sort. Since, conditioned on $F_i = F'$, the mappings ϕ_i^I and ϕ_i^O are uniform, we get (where k_i^I and k_i^O denote the number of roots of F_i in I_i and O_i respectively)

$$P(F_{i+1}=F|F_i=F') = \frac{1}{|O_i|^{k_i^I}}\frac{1}{|I_i|^{k_i^O}}\sum_{\phi}P(F_{i+1}=F|F_i=F',\phi_i=\phi).$$

The conditional probability in the sum above is 1 or 0, dependent upon whether the forest F arises from the pair (F',ϕ) or not. Given $F\in\mathcal{F}_{N_{i+1},k_{i+1}}$, we can construct all such pairs by choosing a quadruplet (a,b,c,d) of four non-negative integers satisfying (1)-(4) above with the additional constraint that $a+c=k_i^I$ and $b+d=k_i^O$. Clearly, the number of such pairs depends only on $|I_i|$, $|O_i|$, $\mathbf{k}_i=(k_i^I,k_i^O)$, k_{i+1}, and $|N_{i+1}|-|N_i|$. We note the number of such pairs by $\beta(|I_i|, |O_i|, \mathbf{k}_i; |I_{i+1}| - |I_i|, k_{i+1})$. Given a pair of integers $\mathbf{k}_i = (k_i^I, k_i^O)$ and a set N_i, we also note $\mathcal{F}_{N_i,\mathbf{k}_i}$ for the set of forests with N_i as the set of vertices, k_i^I roots in I and k_i^O roots in O. Notice that $\mathcal{F}_{N_i,k_i} = \cup_{\mathbf{k}_i : k_i^I + k_i^O = k_i} \mathcal{F}_{N_i,\mathbf{k}_i}$. Therefore, we obtain,

$$
\begin{aligned}
P(F_{i+1} &= F|(N_1, k_1), ..., (N_i, k_i)) = \frac{1}{|\mathcal{F}_{N_i,k_i}|} \sum_{F' \in \mathcal{F}_{N_i,k_i}} P(F_{i+1} = F|F_i = F') \\
&= \frac{1}{|\mathcal{F}_{N_i,k_i}|} \sum_{\mathbf{k}_i : k_i^I + k_i^O = k_i} \sum_{F' \in \mathcal{F}_{N_i,\mathbf{k}_i}} P(F_{i+1} = F|F_i = F') \\
&= \frac{1}{|\mathcal{F}_{N_i,k_i}|} \sum_{\mathbf{k}_i : k_i^I + k_i^O = k_i} \sum_{F' \in \mathcal{F}_{N_i,\mathbf{k}_i}} \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \sum_{\phi} P(F_{i+1} = F|F_i = F', \phi_i = \phi) \\
&= \frac{1}{|\mathcal{F}_{N_i,k_i}|} \sum_{\mathbf{k}_i : k_i^I + k_i^O = k_i} \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \sum_{F' \in \mathcal{F}_{N_i,\mathbf{k}_i}} \sum_{\phi} P(F_{i+1} = F|F_i = F', \phi_i = \phi) \\
&= \frac{1}{|\mathcal{F}_{N_i,k_i}|} \sum_{\mathbf{k}_i : k_i^I + k_i^O = k_i} \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \beta(|I_i|, |O_i|, \mathbf{k}_i; |I_{i+1}| - |I_i|, k_{i+1}).
\end{aligned}
$$

This probability is independent of $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$. Then so is $P(F_{i+1} = F|(N_1, k_1), ..., (N_i, k_i), (N_{i+1}, k_{i+1}))$, since it equals the ratio of the above probability and $P(F \in \mathcal{F}_{N_{i+1},k_{i+1}}|(N_1, k_1), ..., (N_i, k_i))$.
□

Lemma 9 implies that the random sequence (n_i, o_i, k_i^1, k_i^2) is a Markov chain, where n_i is the number of agents, o_i is the number of objects, and k_i^1 is the number of roots in I_i and k_i^2 is the number of roots in I_i, in TTC round i. The arguments are precisely the same as Corollaries 1 and 2 of Frieze and Pittel (1995) and are omitted. We now show that in fact (n_i, o_i) forms a Markov chain.

Lemma 10. *The random sequence (n_i, o_i) is a Markov chain, with transition probability given by*

$$
\begin{aligned}
p_{n,o;m} &:= \Pr\{n_i - n_{i+1} = o_i - o_{i+1} = m | n_i = n, o_i = o\} \\
&= \left(\frac{m}{(on)^{m+1}}\right)\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right)(o + n - m).
\end{aligned}
$$

Proof. We first compute the probability of transition from (n_i, o_i, k_i^1, k_i^2) to $(n_{i+1}, o_{i+1}, k_{i+1}^1, k_{i+1}^2)$:

$$P(n, o, \kappa; m, \lambda_1, \lambda_2) := \Pr\{n_i - n_{i+1} = o_i - o_{i+1} = m, k_{i+1}^1 = \lambda_1, k_{i+1}^2 = \lambda_2 | n_i = n, o_i = o, k_i^1 + k_i^2 = \kappa\}.$$

This will be computed as a fraction $\frac{\Theta}{\Upsilon}$. The denominator Υ counts the number of rooted forests in the bipartite digraph with k_i^1 roots in I_i and k_i^2 roots in O_i, multiplied by the ways in which k_i^1 roots of I_i could point to O_i and k_i^2 roots of O_i could point to I_i.[18] Hence, letting $f(n, o, k_1, k_2)$ denote the number of rooted forests in a bipartite graph (with n and o vertices on both sides) containing k_1 and k_2 roots on both sides.

$$\begin{aligned}\Upsilon =& \sum_{k_1+k_2=\kappa} o^{k_1} n^{k_2} f(n, o, k_1, k_2) \\ =& \sum_{k_1+k_2=\kappa} o^{k_1} n^{k_2} \binom{n}{k_1}\binom{o}{k_2} o^{n-k_1-1} n^{o-k_2-1}(nk_2 + ok_1 - k_1k_2) \\ =& \sum_{k_1+k_2=\kappa} \binom{n}{k_1}\binom{o}{k_2} o^{n-1} n^{o-1}(nk_2 + ok_1 - k_1k_2) \\ =& o^n n^o \left(2\binom{n+o-1}{\kappa-1} - \binom{n+o-2}{\kappa-2}\right).\end{aligned}$$

The first equality follows from the fact that there are $o^{k_1}n^{k_2}$ ways in which k_1 roots in I_i point to O_i and k_2 roots in O_i could point to I_i. The second equality follows from Lemma 6. The last uses Vandermonde's identity.

The numerator Θ counts the number of ways in which m agents are chosen from I_i and m objects are chosen from O_i to form a bipartite bijection each cycle of which contains at least one of κ old roots, and for each such choice, the number of ways in which the remaining vertices form a spanning rooted forest and the λ_1 roots in I_{i+1} point to objects in $O_i \setminus O_{i+1}$ and λ_2 roots in O_{i+1} point to agents in $O_i \setminus O_{i+1}$.

$$\begin{aligned}\Theta =& \binom{n}{m}\binom{o}{m} f(n-m, o-m, \lambda_1, \lambda_2)\alpha(n, o, \kappa; m, \lambda_1+\lambda_2) \\ =& \binom{n}{m}\binom{o}{m} f(n-m, o-m, \lambda_1, \lambda_2)(m!)^2 m^{\lambda_1+\lambda_2}\left(2\binom{n+o-1}{\kappa-1} - \binom{n+o-2}{\kappa-2}\right). \\ =& \left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right) m^{\lambda_1+\lambda_2} f(n-m, o-m, \lambda_1, \lambda_2)\left(2\binom{n+o-1}{\kappa-1} - \binom{n+o-2}{\kappa-2}\right).\end{aligned}$$

Collecting terms, let us compute

$$P(n, o, \kappa; m, \lambda_1, \lambda_2) = \frac{1}{o^n n^o}\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right) m^{\lambda_1+\lambda_2} f(n-m, o-m, \lambda_1, \lambda_2).$$

[18] Given that we have $n_i = n$ individuals, $o_i = o$ objects and $k_i^1 + k_i^2 = \kappa$ roots at the begining of step i under TTC, one may think of this as the total number of possible directed bipartite digraph one may obtain via TTC at the end of step i when we let k_i^1 roots in I point to their remaining most favorite object and k_i^2 roots in O point to their remaining most favorite individual.

A key observation is that this expression does not depend on κ, which implies that (n_i, o_i) forms a Markov chain.

Its transition probability can be derived by summing the expression over all possible (λ_1, λ_2)'s:

$$p_{n,o;m} := \sum_{0\le\lambda_1\le n-m, 0\le\lambda_2\le o-m} P(n, o, \kappa; m, \lambda_1, \lambda_2).$$

To this end, we obtain:

$$\begin{aligned}
&\sum_{0\le\lambda_1\le n-m}\sum_{0\le\lambda_2\le o-m} m^{\lambda_1} m^{\lambda_2} f(n-m, o-m, \lambda_1, \lambda_2)\\
=&\sum_{0\le\lambda_1\le n-m}\sum_{0\le\lambda_2\le o-m} m^{\lambda_1} m^{\lambda_2}\binom{n-m}{\lambda_1}\binom{o-m}{\lambda_2}\times\\
&\qquad (o-m)^{n-m-\lambda_1-1}(n-m)^{o-m-\lambda_2-1}((n-m)\lambda_2+(o-m)\lambda_1-\lambda_1\lambda_2)\\
=&m\left(\sum_{0\le\lambda_1\le n-m}\binom{n-m}{\lambda_1}m^{\lambda_1}(o-m)^{n-m-\lambda_1}\right)\left(\sum_{1\le\lambda_2\le o-m}\binom{o-m-1}{\lambda_2-1}m^{\lambda_2-1}(n-m)^{o-m-\lambda_2}\right)\\
+&m\left(\sum_{1\le\lambda_1\le n-m}\binom{n-m-1}{\lambda_1-1}m^{\lambda_1-1}(o-m)^{n-m-\lambda_1}\right)\left(\sum_{0\le\lambda_2\le o-m}\binom{o-m}{\lambda_2}m^{\lambda_2}(n-m)^{o-m-\lambda_2}\right)\\
-&m^2\left(\sum_{1\le\lambda_1\le n-m}\binom{n-m-1}{\lambda_1-1}m^{\lambda_1-1}(o-m)^{n-m-\lambda_1}\right)\left(\sum_{1\le\lambda_2\le o-m}\binom{o-m-1}{\lambda_2-1}m^{\lambda_2-1}(n-m)^{o-m-\lambda_2}\right)\\
=&mo^{n-m}n^{o-m-1}+mo^{n-m-1}n^{o-m}-m^2o^{n-m-1}n^{o-m-1}\\
=&mo^{n-m-1}n^{o-m-1}(n+o-m),
\end{aligned}$$

where the first equality follows from Lemma 6, and the third follows from the Binomial Theorem.

Multiplying the term $\frac{1}{o^n n^o}\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right)$, we get the formula stated in the Lemma.

□

This concludes the proof of Theorem 7.

B.3 The Number of Objects Assigned via Short Cycles

We begin with the following question: What is the expected number of agents that are assigned via short cycles, conditional on the event that m agents and m objects are cleared and that they contain a old roots on the agent side and b old roots on the object side? The answer is

$$
\begin{aligned}
\Phi(a,b) :=& \binom{m}{1}\binom{m}{1}\times \\
& \Bigg(\left(\frac{a}{m}\right)\left(\frac{m-b}{m}\right)\left(\frac{a+b-1}{m-1}-\frac{(a-1)b}{(m-1)^2}\right)(m-1)!(m-1)! \\
& + \left(\frac{m-a}{m}\right)\left(\frac{b}{m}\right)\left(\frac{a+b-1}{m-1}-\frac{(b-1)a}{(m-1)^2}\right)(m-1)!(m-1)! \\
& + \left(\frac{a}{m}\right)\left(\frac{b}{m}\right)\left(\frac{a+b-2}{m-1}-\frac{(a-1)(b-1)}{(m-1)^2}\right)(m-1)!(m-1)!\Bigg).
\end{aligned}
$$

The formula is explained as follows. The first line counts the number of ways one can choose an agent and an object to form a pair. For the chosen agent-object pair, the subsequent lines count the number of bipartite bijections each cycle of which contains at least one of the $a+b$ roots *and* the chosen pair forms a short-cycle. For instance, the second line equals the probability that the agent in the chosen pair is one of a roots but the object is not an old root, times the expected number of bipartite bijections between the remaining agents and the remaining objects, each cycle of which contains at least one old root among $a-1$ roots on the agents side and b roots on the object side. This last part uses the formula stated in Lemma 8. The subsequent lines are explained similarly. Note that we exclude the possibility that the chosen pair does not contain any root.

Rearrangement of terms produces

$$
\Phi(a,b) = (m!)^2\left(\frac{a(a-1)}{m(m-1)}+\frac{b(b-1)}{m(m-1)}+\frac{2ab}{m^2}-\frac{2a(a-1)b}{m^2(m-1)}-\frac{2ab(b-1)}{m^2(m-1)}+\frac{a(a-1)b(b-1)}{m^2(m-1)^2}\right).
$$

Using this formula, we next compute the total expected number of short cycles for all instances consistent with $(n,o,\kappa;m,\lambda_1,\lambda_2)$.

$$
\begin{aligned}
\frac{\Lambda(\kappa;\lambda_1,\lambda_2)}{\Upsilon} =& \frac{1}{\Upsilon}\binom{n}{m}\binom{o}{m} f(n-m,o-m,\lambda_1,\lambda_2)\sum_{a+b+c+d=\kappa}\binom{n-m}{c}\binom{o-m}{d}\binom{m}{a}\binom{m}{b}\Phi(a,b)m^{\lambda} \\
=& \frac{1}{\Upsilon}\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right)m^{\lambda} f(n-m,o-m,\lambda_1,\lambda_2)\delta(\kappa),
\end{aligned}
$$

where

$$
\delta(\kappa) := 4\binom{n+o-2}{\kappa-2} - 4\binom{n+o-3}{\kappa-3} + \binom{n+o-4}{\kappa-4}.
$$

(Note: Here and elsewhere, the combinatoric terms involving $\kappa - j$ are relevant for $j \geq \kappa$. The terms vanish when $j > \kappa$.)

Recall the term:

$$\Delta(\kappa) := 2\binom{n+o-1}{\kappa-1} - \binom{n+o-2}{\kappa-2}.$$

Then,

$$\rho(\kappa) := \frac{\delta(\kappa)}{\Delta(\kappa)} = \left(\frac{\kappa-1}{2(n+o)-(\kappa-1)}\right)\left(4 - 4\frac{\kappa-2}{n+o-2} + \frac{(\kappa-2)(\kappa-3)}{(n+o-2)(n+o-3)}\right).$$

It can be shown that $\rho(\kappa) \leq 1$ for κ, n, o.

Let σ_m be the expected number of agents/objects assigned via short cycles in any round given that m agent/object pairs are assigned in that round. Let 1_m be the indicator function that takes one if m agent-object pairs are cleared and zero otherwise. Then, we let

Then,

$$\begin{aligned}
&\mathbb{E}[\sigma_m \cdot 1_m | n, o] \\
\leq &\sup_{\kappa} \sum_{0 \leq \lambda_1 \leq n-m, 0 \leq \lambda_2 \leq o-m} \frac{\Lambda(\kappa, \lambda_1, \lambda_2)}{\Upsilon} \\
= &\sup_{\kappa} \left(\frac{m}{(on)^{m+1}}\right)\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right)(o+n-m)\rho(\kappa) \\
= &p_{n,o,m}\left(\sup_{\kappa} \rho(\kappa)\right) \\
\leq &p_{n,o,m}.
\end{aligned}$$

This means that for each $m \leq \min\{o, n\}$,

$$\mathbb{E}[\sigma_m | n, o, m] := \frac{\mathbb{E}[\sigma_m \cdot 1_m | n, o]}{\mathbb{E}[1_m | n, o]} \leq 1.$$

We thus reach the following conclusion:

Theorem 8. *The expected number of agents/objects assigned via short cycles cannot exceed one at any round of TTC.*

B.4 The Number of Objects Assigned via Long Cycles

Again consider the unbalanced market in which $|I| - |O|$ is in the same order of magnitude as $|O|$. Theorem 7 means that the number of objects assigned in any round coincide with the half of the cyclic vertices in a random bipartite graph consisting of the agents and the objects remaining at the beginning of that round. According to Lemma 7, this means that if n agents and $o\,(< n)$ objects begin at any round the expected number of objects that would be assigned by the end of that round would be

$$\mathbb{E}[m] = \sum_{i=1}^{o} \frac{(o)_i (n)_i}{o^i n^i}.$$

We can make two observations: First, the expected number is increasing in (n, o). This can be seen easily by the fact that $\frac{k-l}{k}$ is increasing in k for any $k > l$. Second, there exists $\hat{n}$ such that for any $n > \hat{n}$,

$$\mathbb{E}[m] \geq 2 \text{ if and only if } o \geq 4.$$

This observation follows from the fact that in case of $o = 3$,

$$\sum_{i=1}^{3} \frac{(3)_i (n)_i}{3^i n^i} = 1 + \frac{6n(n-1)}{9n^2} + \frac{6n(n-1)(n-2)}{27n^3} < 2$$

but in case of $o = 4$,

$$\sum_{i=1}^{4} \frac{(4)_i (n)_i}{4^i n^i} \geq 1 + \frac{12n(n-1)}{16n^2} + \frac{24n(n-1)(n-2)}{64n^3} > 2,$$

provided that n is sufficiently large.

We are now ready to present the main result. Recall that $\hat{O}$ is the (random) set of objects that are assigned via long cycles in TTC.

Theorem 9. *For $|I|$ sufficiently large,* $\mathbb{E}\left[\frac{|\hat{O}|}{|O|}\right] \geq \frac{1}{2} - \frac{5}{|O|}$.

Proof. Consider the following sequence of random variables $\{\mathbb{E}(Z_k\,|o_k)\}_{k=1}^{|O|}$ where o_k is the random variable which corresponds to the number of remaining objects at round k while Z_k is the random variable corresponding to the number of objects assigned at round k via long cycles. Thus, $o_1 = |O|$. By the above observation, since $\{o_k\}$ is a decreasing sequence, this sequence is decreasing as well: $\mathbb{E}(Z_1\,|o_1) \geq \ldots \geq \mathbb{E}(Z_{|O|}\,\big|o_{|O|}) = 0$. By Theorem 7, we are defining here the process $\{\mathbb{E}(Z_k\,|o_k)\}_{k=1}^{|O|}$ induced by this Markov chain. Note also that $\mathbb{E}(Z_k\,|o_k) \geq 1$ if and only if $o_k \geq 4$. Denote the random integer T for the first integer at which the decreasing sequence is smaller than 1, i.e., $\mathbb{E}(Z_k\,|o_k) \leq 1$ if and only if $k \geq T$. Note that $o_T \leq 4$.

For each $k = 1, ..., |O|$, let $\hat{o}_k|t$ denote the number of agents remaining at round k, conditional on $T = t$. This is a random variable. Note that $\hat{o}_k|t \geq 4$ for $k \leq t-1$.

Now we obtain:

$$
\begin{aligned}
\mathbb{E}[|\hat{O}|] &= \mathbb{E}(\sum_{k=1}^{|O|} Z_k) \\
&= \sum_t \Pr\{T=t\}\mathbb{E}\left[\sum_{k=1}^{t-1} Z_k + \sum_{k=t}^{|O|} Z_k \middle| T=t\right] \\
&= \sum_t \Pr\{T=t\}\left(\sum_{k=1}^{t-1}\mathbb{E}\left[Z_k \middle| T=t\right] + \sum_{k=t}^{n}\mathbb{E}\left[Z_k \middle| T=t\right]\right) \\
&= \sum_t \Pr\{T=t\}\left(\sum_{k=1}^{t-1}\mathbb{E}\left[\mathbb{E}\left[Z_k\middle|\hat{o}_k|t\right]\middle| T=t\right] + \sum_{k=t}^{|O|}\mathbb{E}\left[\mathbb{E}\left[Z_k\middle|\hat{o}_k|t\right]\middle| T=t\right]\right) \\
&= \sum_t \Pr\{T=t\}\left(\mathbb{E}\left[\sum_{k=1}^{t-1}\mathbb{E}\left[Z_k\middle|\hat{o}_k|t\right]\middle| T=t\right] + \mathbb{E}\left[\sum_{k=t}^{|O|}\mathbb{E}\left[Z_k\middle|\hat{o}_k|t\right]\middle| T=t\right]\right) \\
&\geq \sum_t \Pr\{T=t\}\mathbb{E}\left[\sum_{k=1}^{t-1}\mathbb{E}\left[Z_k\middle|\hat{o}_k|t\right]\middle| T=t\right] \\
&\geq \sum_t \Pr\{T=t\}\mathbb{E}\left[t-1\middle| T=t\right] \\
&= \mathbb{E}[T]-1.
\end{aligned}
$$

Since once we reach round T, at most four more short cycles can be formed, so the expected number of short cycles should be smaller than $\mathbb{E}(T)+4$. Indeed, the expected number of short cycles is smaller than the expected number of rounds for TTC to converge (the expected number of short cycle at each round is at most one) which itself is smaller than $\mathbb{E}(T)+4$. It follows that the

$$\mathbb{E}[|\hat{O}|] \geq \mathbb{E}[T]-1 \leq \mathbb{E}[|O|-|\hat{O}|]-5,$$

from which the result follows. □

Corollary 5. *There exists $\gamma > 0, \delta > 0, N > 0$*

$$\Pr\left\{\frac{|\hat{O}|}{|O|} > \delta\right\} > \gamma,$$

for all $|O| > N$.

References

Abdulkadiroglu, A., Y.-K. Che, and O. Tercieux (2013): "The Role of Priorities in Top Trading Cycles," Columbia University, Unpublished mimeo.

Abdulkadiroglu, A., Y.-K. Che, and Y. Yasuda (2008): "Expanding 'Choice' in School Choice," mimeo.

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

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

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.*

Balinski, M., and T. Sönmez (1999): "A tale of two mechanisms: student placement," *Journal of Economic Theory*, 84, 73–94.

Bollobas, B. (2001): *Random Graphs*. Cambridge University Press, Cambridge, United Kingdom.

Calsamiglia, C., G. Haeringer, and F. Klijn (2010): "Constrained School Choice: An Experimental Study," *American Economic Review*, 100, 1860–1874.

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.

Chen, Y., and O. Kesten (2013): "From Boston to Chinese Parallel to Deferred Acceptance: Theory and Experiments on a Family of School Choice Mechanisms," Unpublished mimeo.

Dawande, M., P. Keskinocak, J. Swaminathan, and S. Tayur (2001): "On Bipartite and Multipartite Clique Problems," *Journal of Algorithms*, 41, 388403.

Dubins, L. E., and D. A. Freedman (1981): "Machiavelli and the Gale-Shapley algorithm," *American Mathematical Monthly*, 88, 485–494.

Frieze, A., and B. Pittel (1995): "Probabilistic Analysis of an Algorithm in the Theory of Markets in Indivisible Goods," *The Annals of Applied Probability*, 5, 768–808.

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.

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

Jaworski, J. (1985): "A Random Bipartite Mapping," *The Annals of Discrete Mathematics*, 28, 137–158.

Jin, Y., and C. Liu (2004): "Enumeration for spanning forests of complete bipartite graphs," *Ars Combinatoria - ARSCOM*, 70, 135–138.

Knuth, D. E. (1996): *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. (2012): "Incentive compatibility of large centralized matching markets," URL http://www. hss. caltech. edu/˜ sangmok/research. html.

Lovasz, L. (1979): *Combinatorial Problems and Exercises*. North Holland, Amsterdam.

McVitie, D. G., and L. Wilson (1971): "The stable marriage problem," *ACM*, 14, 486–490.

Pathak, P. A., and T. Sömez (2013): "School Admissions Reform in Chicago and England: Comparing Mechanisms by their Vulnerability to Manipulation," *American Economic Review*, 103, 80–106.

Pittel, B. (1989): "The Average Number of Stable Matchings," *SIAM Journal on Discrete Mathematics*, 2, 530–549.

——— (1992): "On Likely Solutions of a Stable Marriage Problem," *The Annals of Applied Probability*, 2, 358–401.

Roth, A. E. (1982): "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 7, 617–628.

Roth, A. E., and M. A. O. Sotomayor (1990): *Two-sided matching: a study in game-theoretic modeling and analysis*. Cambridge University Press, Cambridge.

Shapley, L., and H. Scarf (1974): "On Cores and Indivisibility," *Journal of Mathematical Economics*, 1, 22–37.

Sonmez, T., and U. Unver (2011): "Market Design for Kidney Exchange," in *Oxford Handbook of Market Design*, ed. by M. N. Z. Neeman, and N. Vulkan. forthcoming.