Abundance of Competent Choices in Assignment Markets

Afshin Nikzad*

PRELIMINARY DRAFT

Abstract

A common utility specification model in assignment problems determines the utility of agent a having an object o by a common-value component v_o for the object and an idiosyncratic component v_o^a that depends on both the object and the agent. The assumption that the idiosyncratic components corresponding to each agent are independently and identically distributed (iid) across objects is ubiquitous in the literature on assignment mechanisms and matching markets. We show that, under this assumption, a large family of assignments—including all Pareto-efficient assignments as well as assignments with large Pareto inefficiencies—become asymptotically *payoff-equivalent*: all of these assignments attain the "best" possible payoff distribution, in the sense of first-order stochastic dominance. We introduce *abundance of competent choices*, a by-product of the iid assumption, as a driving force of such equivalence results and show that this phenomenon leads to similar equivalence results in other setups, such as Gale and Shapley's model of marriage markets [Gale and Shapley 1962].

On the other hand, we show that the abundance of competent choices does not emerge, and the subsequent payoff equivalence results do not generally hold, when the iid assumption is replaced with *negative correlation* (which means that if an agent has a "high" idiosyncratic utility component for an object, it becomes more likely that she has "lower" idiosyncratic utility components for other objects). Consequently, models with iid idiosyncratic components may not be sufficient for identifying market inefficiencies. In particular, when there is no abundance of competent choices in a marketplace, models with negatively correlated idiosyncratic components are likely to approximate that marketplace better than models with iid idiosyncratic components.

*Department of Economics, USC. Email: afshin.nikzad@usc.edu

Contents

1 Introduction

A common utility specification model in assignment problems writes the utility of an agent a having an object o as the sum of a common-value component v_o for the object and an idiosyncratic component that depends on both the agent and the object, namely v_o^a. The assumption that the idiosyncratic components corresponding to each agent are independently and identically distributed across objects (henceforth, the *iid assumption*) is a prevalent assumption used alongside this model in the literature.

We show that in various settings the iid assumption leads to the payoff-equivalency of a large family of assignments, including all Pareto-efficient assignments as well as assignments with large Pareto inefficiencies. In fact, all of these assignments attain the "best" possible payoff distribution, in the sense of first-order stochastic dominance. The driving force of this effect is a phenomenon that we dub the *abundance of competent choices*, a by-product of the iid assumption, as explained next.

It is easier to understand the abundance of competent choices in one-to-one assignment markets where the common-value components are identical. There, the abundance of competent choices boils down to the abundance of "top" choices, which means that, under the iid assumption, for any agent there would be many objects that provide her a utility "close" to the maximum achievable. We will see that this effect creates many assignments in which almost all agents attain their maximum conceivable utility,[1] asymptotically. These assignments include all Pareto-efficient assignments, as well as assignments with large Pareto inefficiencies.[2]

A similar effect can happen when the common-value components are heterogeneous. For some intuition, consider a one-to-one assignment market with an equal number of agents and objects. In such a market, common-value components always contribute a fixed amount to any assignment. Therefore, they can be ignored altogether when comparing the agents' payoff distributions in different assignments. Thus, payoff-equivalence results similar to the case of identical common-value components would hold.

On the other hand, we show that the abundance of competent choices does not emerge, and the subsequent payoff equivalence results do not hold, when the idiosyncratic components become *negatively correlated* (which means if an agent has a "high" idiosyncratic utility component for an object, it becomes more likely that she has "lower" idiosyncratic utility

[1] We say an agent's maximum conceivable utility is $\overline{u}$ when her valuation has support $[\underline{u}, \overline{u}]$.

[2] By an assignment with large Pareto inefficiencies we mean an assignment in which a large fraction of agents can be made strictly better off without making any other agent worse off.

components for the other objects). Hence, the iid assumption can lead to payoff-equivalency results for assignment mechanisms (e.g., [Che and Tercieux 2018, Lee and Yariv 2018]), which would not hold in general otherwise.

Theoretical models in economics sometimes leverage the iid idiosyncratic components assumption, possibly for its analytical simplicity (e.g., see [Kanoria et al. 2015, Arnosti and Shi 2017, Che and Tercieux 2018, Lee and Yariv 2018, Che and Tercieux 2019], among others).[3] Our findings suggest that models with iid idiosyncratic components may not be sufficient for identifying market inefficiencies, and those with (negatively) correlated idiosyncratic components can be sharper in distinguishing between the assignments and capturing the assignment inefficiencies. In particular, when there is no abundance of competent choices in a marketplace, models with negatively correlated idiosyncratic components are likely to approximate that marketplace better than models with iid idiosyncratic components.

1.1 Summary of results

We establish the above insights through a series of results on assignment models with both ordinal and cardinal preferences, which are briefed next. First, we consider a random market with n agents and n objects, where each agent ranks all objects independently and uniformly at random. We show that, in such a market, there exists an assignment (bijection) of objects to agents with a constant average rank. That is, there exists a constant R, independent of the market size (n), such that the object assigned to each agent has an average position of at most R on her list (Proposition 4.1).

On the other hand, a well-known result by [Knuth 1996] shows that in the same market, the average rank of the assignment generated by the random serial dictatorship mechanism, in which agents are ordered randomly and choose objects one by one in that order, is $\log n$. Comparing these two results implies that the average rank of the assignment takes a heavy toll under the random serial dictatorship mechanism compared to the first-best.

We then contrast this gap between the two Pareto-efficient assignments to the payoff-equivalency result of [Che and Tercieux 2018], whose work implies that in the same market, Pareto-efficient assignments are asymptotically payoff-equivalent (up to relabeling of agents) if the utility of an agent from an object is drawn iid across all agent-object pairs.[4] One might presume that the difference between these findings stems from the fact that one of them is

[3] That said, there is empirical work on matching markets that allows for arbitrary correlations between the idiosyncratic utility components; e.g., see [Abdulkadiroglu et al. 2017].

[4] We will discuss that their result is in fact more general.

based on an ordinal notion of preferences and the other on a cardinal notion. We rule this out by showing that the gap between the two Pareto-efficient assignments can also persist in cardinal models.[5]

What accounts for the difference is, rather, the correlation between the idiosyncratic components of an agent's utilities over the objects. Consider again the utility specification model which writes the utility of an agent a having an object o as $u_a(o) = v_o + v_o^a$. When the idiosyncratic components are iid, [Che and Tercieux 2018] show that Pareto-efficient assignments are asymptotically payoff-equivalent under their assumptions. On the other hand, when the idiosyncratic components are negatively correlated, we show that Pareto-efficient assignments can have very different payoff distributions. (Proposition 5.1)

To see why Pareto-efficient assignments can have very different average ranks as discussed earlier, note that strict ordinal preferences encapsulate a notion of correlated idiosyncrasy: when an object is ranked first, any other object must be ranked second or worse. Such a notion of (negative) correlation is absent from cardinal utility models with iid idiosyncratic components, and its absence causes the *abundance of competent choices* phenomenon, formally introduced in Section 5. There, we show that this phenomenon leads to asymptotic payoff-equivalency of a large family of assignments that includes all Pareto-efficient assignments as well as assignments with large Pareto inefficiencies. In fact, we will show that all of these assignments attain the "best" possible payoff distribution, in the sense of first-order stochastic dominance.

Finally, in Section 6, we discuss the abundance of competent choices in other markets by focusing on two well-studied setups: Gale and Shapley's model of marriage markets [Gale and Shapley 1962] and the *hiring problem* [Bruss 2000]. In particular, we highlight the results of [Lee and Yariv 2018], which imply that in any stable assignment in a marriage market, any agent asymptotically attains the maximum conceivable utility when the idiosyncratic components are iid. Analogous to our previous findings, we show that this is no longer the case when the idiosyncratic components are negatively correlated (Proposition 6.1).

Besides establishing the notion of abundance of competent choices, our analysis provides an additional insight: the average rank under the RSD mechanism takes a heavy toll compared to the first-best. This is a consequence of one of our results that shows the existence of Pareto-efficient assignments with a constant average rank in random markets. The proof needs to keep track of the correlations between the objects' ranks (which is precisely why payoff-equivalence results break here). To handle these correlations, we use two applications

[5] This can happen, for instance, when agents have dichotomous preferences (Section 5).

of the FKG inequality [Fortuin et al. 1971], a fundamental correlation inequality in statistical mechanics and probabilistic combinatorics.

The rest of the paper is organized as follows. Section 2 reviews related literature. Section 3 presents the main setup. Section 4 proves the existence of Pareto-efficient assignments with a constant average rank in random markets. Section 5, as discussed above, compares this finding to recent payoff equivalence results, deciphers the apparent contrast, and introduces the *abundance of competent choices* as a driving force of such equivalence results under the iid assumption. Section 6, as noted above, discusses the abundance of competent choices in other setups.

2 Related literature

The assumption of iid idiosyncratic components is a common assumption in the literature of two-sided markets. For example, it is used to model match values in studying welfare under stable assignments [Lee and Yariv 2018, Che and Tercieux 2019], size of the core in stable assignments [Kanoria et al. 2015], agents' valuations for objects in studying Pareto-efficient assignments [Che and Tercieux 2018], agents' valuations for houses in the design of lotteries to allocate housing [Arnosti and Shi 2017], and match values in dating platforms [Kanoria and Saban 2017].

The iid idiosyncratic components assumption, as discussed in this paper, is consequential in that in can lead to payoff-equivalency results (between assignments or assignment mechanisms) which may not be present otherwise. For example, under the iid idiosyncratic components assumption, it is shown that the Pareto-efficient assignments of objects to agents are asymptotically payoff-equivalent [Che and Tercieux 2018]. Similar payoff-equivalence results have been established for stable assignments in two-sided markets. Lee [2017] shows that, under the iid idiosyncratic assumption,[6] in large matching markets most agents are close to being indifferent over different stable assignments, and therefore, they do not have significant incentives to manipulate mechanisms that produce a stable assignment. In a related work, Lee and Yariv [2018] show that the average efficiency of all stable assignments is essentially the same, asymptotically equal to the maximum efficiency achievable by any (possibly non-stable) assignment, when the idiosyncratic components are iid. Compte and Jehiel [2008] propose a modified notion of stability in a problem involving reassignments, and

[6] He assumes that the idiosyncratic components are iid from the uniform distribution over the unit interval, and mentions that "Whatever distribution we assume, there always exists a change of variables that delivers the uniform distribution, and we can transform utility functions monotonically."

design a mechanism for finding such stable reassignments optimally. Inline with the findings of Lee and Yariv [2018], they find asymptotically efficient optimal assignments when the idiosyncratic components are iid.

The "core-convergence" result of Kanoria et al. [2015] as well is relevant to the findings of Lee and Yariv [2018]. Kanoria et al. [2015] show that, under the iid idiosyncratic components assumption[7] the *size* of core is "small", where size is defined to be "the difference between the maximum and minimum total utility of workers (equivalently, firms) in the core, scaled by the (maximum) social welfare."

Che and Tercieux [2019] show that efficiency and stability can be achieved at the same time, under the iid idiosyncratic components assumption. In particular, they show that, in the absence of common-value components for objects, both efficiency and stability can be achieved in an asymptotic sense via standard mechanisms such as deferred acceptance and top trading cycles. In the presence of common-value-components, however, they design another mechanism that achieves stability and Pareto efficiency asymptotically, as the standard mechanisms become either inefficient or unstable.

Liu and Pycia [2016] show that all asymptotically efficient, symmetric, and asymptotically strategy-proof mechanisms lead to the same allocations in large markets, essentially when objects have large capacities. Their approach is different than most of the above in that, rather than making probabilistic assumptions about the agents' payoffs, they focus on mechanisms in which no small subset of agents can change the outcome significantly.

From a technical perspective, we build on techniques from Random Graph theory and probabilistic analysis. One of the results—within the results that we provide to establish the abundance of competent choices phenomenon—is of independent interest. In Proposition 4.1 we show that in *random* assignment markets there exists an assignment with a constant average rank,[8] i.e., an average rank that does not depend on the market size. This finding, together with a result of Knuth [1996], show that the average rank under the RSD mechanism takes a heavy toll compared to the first-best. Our proof uses insights from a theorem of Walkup [1980] on the existence of perfect matchings in random graphs, together with two applications of the FKG inequality [Fortuin et al. 1971], a fundamental correlation inequality in statistical mechanics and probabilistic combinatorics. Finally, we note the potential connections and the differences between Proposition 4.1 and the Parisi conjecture,

[7] More precisely, they study a two-sided market with a fixed number of agent types, where the match utilities for agents are determined by independent idiosyncratic components whose distribution only depends on the agent types.

[8] We will define rank formally later, as the position of an object on an agent's preference list.

proved by Linusson and Wästlund [2004]. The details are discussed in Section 4.1. Making a formal connection to the Parisi conjecture is a theoretical direction of independent interest.

3 Setup

In this section, we set up the assignment model and define notions that will allow us to compare efficiency of different assignments cardinally and ordinally.

An *assignment market* contains a set of agents A and a set of objects O. We use n, m to denote $|A|, |O|$, respectively. Each object has a capacity c. Throughout the paper we assume that $n = cm$. An assignment is a function $\mu : A \to O$ that assigns each agent to an object without assigning more than c agents to any object.

Each agent a receives utility $u_a(o)$ from being assigned to an object $o \in O$. We suppose that $u_a(o) = v_o + v_o^a$, where v_o is called the *common-value component* for object o and v_o^a is called the *idiosyncratic utility component* of agent a for object o.[9] An assignment μ is called *Pareto-efficient* if there exists no other assignment in which all agents receive weakly higher utilities than in μ, with at least one agent receiving strictly higher utility. The *utility-optimal assignment* is the assignment that maximizes the sum of agents' utilities. The *average utility* of an assignment is defined by the sum of agents' utilities normalized by the number of agents.

The agents' ordinal preferences

The utilities of an agent from the objects *induce* a (possibly non-strict) preference ordering over the objects for that agent. This ordering is also called the *preference list* of the agent. For each agent a, the first position of her preference list contains all of the objects that provide her with the highest utility from the set $\{u_a(o)\}_{o \in O}$, the second position of the list contains all of the objects that provide the agent with the second highest utility, and so on. The position of an object o on the preference list of an agent a is called the *rank* of the object for that agent, and is denoted by $r_a(o)$. The average rank of an assignment μ is defined as

$$\overline{r}(\mu) = \frac{1}{|A|} \cdot \sum_{a \in A} r_a(\mu(a)).$$

The set of preference lists of all agents is called a *preference profile*. When A, O are

[9] Our results hold more generally where $u_a(o) = U(v_o, v_o^a)$ and U is monotone and Lipschitz continuous. The focus here is on the case of $U(v_o, v_o^a) = v_o + v_o^a$ for expositional simplicity.

known from the context, we denote the set of all preference profiles by Π and typically denote a member of it by π. For a preference profile π, let the *rank-optimal* assignment, $r^\star(\pi)$, be the assignment with the minimum average rank. Observe that, given any π, $r^\star(\pi)$ is a Pareto-efficient assignment.

Random assignment markets

A *random market* is an assignment market in which the induced preference list of each agent is distributed uniformly over the set of all strict orderings of objects, independently of all the other agents. An example of a random market is an assignment market wherein the common-value components are the same for all objects and the idiosyncratic components are drawn iid from a continuous CDF.

Assignment mechanisms

The *random serial dictatorship* (RSD) mechanism is an assignment mechanism in which agents are ordered randomly and then, one by one in that order, choose their most preferred item from the set of remaining items. We call the assignment generated by RSD the *RSD assignment*.

Asymptotic notations

For two functions $f, g : \mathbb{R}_+ \to \mathbb{R}_+$ we write $f = o(g)$ when $\lim_{n\to\infty} f(n)/g(n) = 0$, $f = \omega(g)$ when $g = o(f)$, $f = \Omega(g)$ when $\lim_{n\to\infty} f(n)/g(n) > 0$, $f = O(g)$ when $g = \Omega(f)$, and $f = \Theta(g)$ when $f = O(g)$ and $g = O(f)$.

We say a sequence of events $E_1, E_2, \ldots$ happen *with high probability* (whp) if $\mathbb{P}[E_n]$ approaches 1 as n approaches infinity.

4 Average rank of Pareto-efficient assignments

This section performs a preliminary motivating analysis that compares Pareto-efficient assignments with respect to their average ranks. The next result shows that, in a random market, the expected average rank of the rank-optimal assignment is bounded from above by a constant independent of the market size.

Proposition 4.1. *There exists a constant R independent of c, n such that the expected average rank of the rank-optimal assignment in a random market is at most R.*

We sketch the proof in Section 4.1, where we also illustrate that the constant R is quite small (less than 2). These findings hold for any capacity parameter c.

When the object capacities are 1, [Knuth 1996] shows that the expected average of the RSD assignment is almost equal to $\ln n$, and therefore approaches infinity as n approaches infinity, in contrast to the average rank in the rank-optimal assignment. The gap between the rank-optimal assignment and the RSD assignment persists even when the object capacities are greater than 1; in particular, when $c = o(\log n)$.

Proposition 4.2. *The expected average rank of the RSD assignment is at least* $\frac{\ln m - 1}{c}$*.*

In Section 5, we show that the gap between different Pareto-efficient assignments also persists in cardinal utility models. On the other hand, Pareto-efficient assignments are shown to be asymptotically payoff-equivalent by [Che and Tercieux 2018]. This apparent contradiction has a simple explanation, which we discuss in Section 5.

4.1 Discussion of the result and proof ideas

To prove Proposition 4.1, we first show that the expected average rank of the rank-optimal assignment is smaller in a random market with n agents and n/c objects each with capacity c than in a random market with n agents and n objects each with capacity 1. (Lemma B.5 in the appendix shows that this holds in a stronger sense: stochastic dominance of the rank distributions.) Given this fact, it suffices to prove the claim for unit capacities: any constant R that bounds the expected average rank in markets with unit capacities will also be a valid upper bound in markets with larger capacities.

Before presenting a proof sketch for the case of unit capacities, we discuss how large the constant R can be. Our proof shows that $R < 7\frac{3}{4}$, whereas our simulations demonstrate that $R < 2$. (See Figure 1) The simulations are for markets with unit capacities. As discussed above, the upper bound for the case of unit capacities is also a valid upper bound for the general case.

The proof idea for the case of unit capacities is inspired by a result of [Walkup 1980] in random graph theory. We use two applications of the FKG inequality[10] to make a similar proof approach applicable. [Walkup 1980] shows that in a random bipartite graph with n nodes on each side, where each node is connected to two distinct neighbors independently and uniformly chosen from the other side, there exists a perfect matching, with high probability.

[10]We recall the definition of the FKG inequality in the appendix, Section A.1.

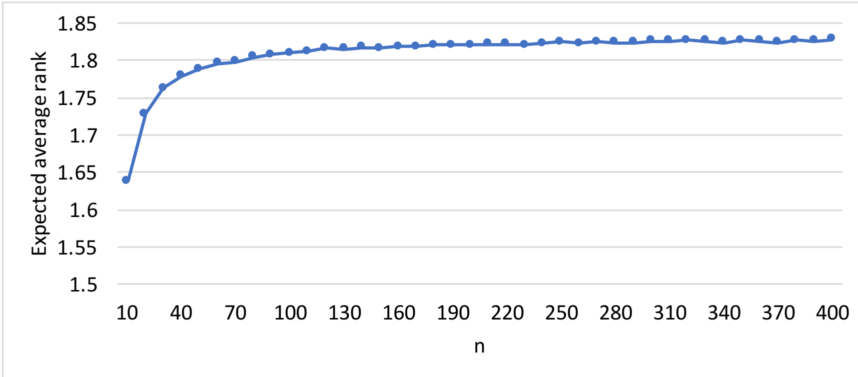

Figure 1: For each $n \in \{10, 20, 30, \ldots, 400\}$, we report the average over 1000 independent samples. The largest reported average is below 1.83. In addition, for $n = 1000$ and $n = 10000$ we take the average over 10 independent draws due to computational limitations; the simulations report average ranks 1.831 and 1.827, respectively.

We construct a bipartite graph $G(A, O)$ from a given preference profile π. Connect each node $a \in A$ to the first 3 items on the preference list of a. Also, connect each node in O to "almost" 3 other neighbors in A, chosen as follows. For any $o \in O$, among the agents who rank o after the third position, choose the 3 agents who rank o the highest. Connect o to those agents (on the other side of the bipartite graph).

We prove that $G(A, O)$ has a perfect matching with probability at least $1 - n^{-3}$ for all $n \geq 50$. The proof uses the König theorem, which implies that the size of the maximum matching (number of its edges) plus the size of the maximum independent set[11] in a bipartite graph is equal to the total number of nodes. Therefore, to prove that $G(A, O)$ has a perfect matching, it suffices to show that it has no independent set of size $n+1$. We show that this holds with high probability using a union bound on all the subsets of nodes of size $n+1$.

In the graph that we construct, unlike in the graph in [Walkup 1980], the neighbors of nodes $o \in O$ are not chosen independently. In particular, there are two types of correlations involved: (i) Correlations across the partitions: The "almost" 3 neighbors that are chosen for a node in O are not independent of the 3 neighbors that are chosen for a node in A, and (ii) Correlations within each partition: The "almost" 3 neighbors that are chosen for a node $o_1 \in O$ are not independent of the "almost" 3 neighbors that are chosen for a node $o_2 \in O$. The bulk of the proof involves handling these correlations, mainly using the FKG inequality.

These correlations are the reason behind the previously discussed gap between the rank-optimal assignment and the RSD assignment (after Proposition 4.1). In particular, when

[11]Recall that an independent set in a graph is a subset of its nodes which are pairwise non-adjacent.

agents' utilities for objects are iid (for all agent-object pairs), then these correlations vanish, and all Pareto-efficient assignments, including the rank-optimal and RSD assignments, become payoff-equivalent as shown by [Che and Tercieux 2018].

Finally, we highlight the similarities between Proposition 4.1 and the Parisi conjecture, proved by Linusson and Wästlund [2004]. The Parisi conjecture states that in an $n \times n$ matrix, whose elements are iid from the exponential distribution with mean 1, there exist a set of n elements, no two in the same row or column, which asymptotically sum up to $\pi^2/6$. While this is reminiscent of Proposition 4.1, we have not found a proof for Proposition 4.1 using the Parisi conjecture. Making a formal connection to Parisi conjecture is a theoretical direction of independent interest. The main technical difference between the setting of Proposition 4.1 and the Parisi conjecture is that, in there latter, the elements of the matrix are independent, whereas in our setting, the elements in a row (ranks of the objects) are correlated. To handle these correlations our proof takes a completely different proof approach using the FKG inequality [Fortuin et al. 1971]. Our approach also provides an "upper-bound distribution" for the rank of the object assigned to each agent, rather than merely providing a bound on the average rank as in the Parisi conjecture. (See Section B.3 in the appendix.)

5 Abundance of competent choices

In this section, we first show that when the idiosyncratic components are not iid, different assignments can have very different payoff distributions. (This will also explain the discussed gap in Section 4 between average ranks of different Pareto efficient assignments.) Then, we will show that under the iid assumption, a large family of assignments, including Pareto efficient assignments as well as assignments with large Pareto inefficiencies, become payoff-equivalent.

The next proposition compares average utilities of the utility-optimal and the RSD assignments when agents' preferences are *dichotomous* . This means that for each agent there are two utility levels: the agent derives utility v_1 from her *top objects*, and utility $v_2 < v_1$, from any other object [Bogomolnaia et al. 2005].

Proposition 5.1. *Consider an assignment market with unit capacities. Agents have dichotomous preferences, where each agent derives utility v_1 from any of her* top *objects and utility $v_2 < v_1$ from any other object. The set of top objects of each agent contains 3 distinct objects, and is drawn independently and uniformly at random. Then, the difference between average utilities of the utility-optimal and the RSD assignment is at least $\frac{v_1 - v_2}{20}$, whp.*

There exists a non-vanishing gap between the utility-maximizing and the RSD assignments in the market considered above. On the other hand, Pareto-efficient assignments are asymptotically payoff-equivalent under the assumptions of [Che and Tercieux 2018]. To explain this difference we note that, in the market considered in Proposition 5.1, the idiosyncratic utility components are *negatively correlated* (as explained below), whereas the payoff equivalence result of [Che and Tercieux 2018] requires the independence of idiosyncratic components. When the idiosyncratic components are iid, the gap between the utility-optimal assignment and the RSD assignment vanishes: not only these assignments but a large family of assignments, including all Pareto efficient assignments as well as assignments with large Pareto inefficiencies, become asymptotically payoff-equivalent, as shown in Section 5.1.

To clarify in what sense the idiosyncratic components are negatively correlated in Proposition 5.1, first we note that, for each agent a, the joint distribution of the idiosyncratic components $v_1^a, \ldots, v_n^a$ is the uniform distribution over the binary vectors $(v_1^a, \ldots, v_n^a)$ such that $v_i^a \in \{0, 1\}$ and $\sum_{i=1}^n v_i^a = 3$. This means that $\mathbb{P}\left[v_i^a = 1 | v_j^a = 1\right] < \mathbb{P}\left[v_i^a = 1\right]$ for any $i \neq j$; that is, conditioning on $v_j^a = 1$ makes it less likely that $v_i^a = 1$.[12]

To provide some intuition on why the iid assumption causes such a difference, first we consider the case of no common-value components: suppose that $v_o = 0$ for all objects o and that the random variables v_o^a are iid with support $[\underline{u}, \overline{u}]$. In this case, Proposition 5.3 will show that all agents attain the maximum conceivable utility, $\overline{u}$, asymptotically. Loosely speaking, the reason is that there are many objects that provide a high utility for each agent. We call this effect the *abundance of competent choices*. The next proposition isolates this effect in a simple environment.

Proposition 5.2. *Let $u_1, \ldots, u_n$ be a sequence of iid random variables with support $[\underline{u}, \overline{u}]$. Let u denote the k-th greatest element of the sequence. Then, as n approaches infinity, u converges in probability to the degenerate distribution centered at $\overline{u}$ as long as $k = o(n)$.*

One might assume that the (heterogeneity of) common-value components might eliminate the abundance of competent choices. While this is true in the setting of the above proposition, this is not the case in our market, as discussed in Section 5.2.

We emphasize that the driving force of the abundance of competent choices is the iid assumption. This effect does not occur when the idiosyncratic components have a sufficiently large degree of negative correlation between them (e.g., as in Proposition 5.1). Strict ordinal preferences indeed encapsulate a notion of negatively correlated idiosyncrasy as well: when

[12]The negative correlation property holds in a more general sense: it holds when we condition on the subsets of variables that are equal to 1; e.g., $\mathbb{P}\left[v_i^a = 1 | v_j^a = 1, v_k^a = 1\right] < \mathbb{P}\left[v_i^a = 1\right]$ holds for distinct i, j, k.

an object is ranked first, any other object must be ranked second or worse. This negative correlation leads to the large gap discussed in Section 4 between average ranks of the RSD and the rank-optimal assignments.

5.1 Failure of the large market limit in *ranking* the assignments

We will show that when the idiosyncratic components are iid, an inefficient version of the RSD mechanism attains the maximum utilitarian upper bound together with many other Pareto-efficient or -inefficient assignment mechanisms (all of which therefore would be payoff-equivalent). First, we present the result for the case of zero common-value components and then we make a similar observation for the case of positive iid common-value components.

Consider an assignment market with n agents and n objects each with unit capacity. Also, let $k(n) : \mathbb{N} \to \mathbb{N}$ be any increasing function. For notational simplicity, we drop the argument n when it is clearly known from the context.

The *Inefficient Random Serial Dictatorship* (Inefficient-RSD) mechanism assigns agents to objects as follows. The agents are ordered randomly and, in that order, choose objects one by one. The choosing agent chooses an object with rank k on her preference list, if such an object is available. Otherwise, she chooses an object with rank $k-1$, if available. Otherwise, she chooses an object with rank $k-2$, if available, and so on. At the end, if an object with rank 1 is also not available, the agent chooses an object with rank $k+1$, if available. Otherwise, she chooses an object with rank $k+2$, if available, and so on.

We say an assignment $\mu : A \to O$ contains a *Pareto-improving cycle of length* l when there exist agents $a_0, \ldots, a_{l-1}$ such that, for all $i \in \{0, 1 \ldots, l-1\}$, agent a_i prefers $\mu(a_{i+1})$ to $\mu(a_i)$, where $i+1$ is computed modulo l.

Proposition 5.3. *Suppose that the utility of any agent $a \in A$ from any object $o \in O$ is drawn independently from a continuous CDF G with support $[\underline{u}, \overline{u}]$. Then:*

i. (Abundance of competent choices) *When $k = o(n)$, the utility of any agent in the Inefficient-RSD mechanism converges in probability to the degenerate distribution centered at $\overline{u}$, as n approaches infinity.*

ii. When $k \geq k_0$, where k_0 is a sufficiently large constant,[13] *the Inefficient-RSD assignment is not Pareto-efficient, whp: it contains a Pareto-improving cycle of length $\Theta(n)$.*

[13] Our simulations show that $k_0 = 3$ suffices. (Section D.4 in the appendix)

iii. For all k, the average rank in the Inefficient-RSD assignment is $\Omega(\max\{k, \ln n\})$, whereas in the RSD assignment it is $O(\ln n)$.

The above observation is not specific to the Inefficient-RSD mechanism; there are other families of assignments with similar properties.[14] Complementing our findings in Section 4, this proposition suggests that (i) in our market, abundance of competent choices is a likely consequence of iid idiosyncratic components, and (ii) under this assumption, cardinal models can fail to capture rank or Pareto inefficiencies, asymptotically; however, ordinal models or cardinal models with negatively correlated idiosyncratic components can be sharper in distinguishing between the assignments and capturing the assignment inefficiencies.

In Section 5.2 we provide a counterpart for the above proposition in the presence of non-zero iid common-value components. A brief intuition for why a similar observation holds in this case is that, in a market with n agents and n objects, common-value components always contribute a fixed amount to any assignment. For example, to maximize the assignment's average utility, it suffices to maximize the sum of idiosyncratic components. This intuition is the rationale behind dubbing the effect "abundance of competent choices", rather than abundance of top choices. Even when the heterogeneity of common-value components makes "top" choices scarce, the independence of idiosyncratic components provides for each agent many objects with (almost) the highest possible idiosyncratic utility component. This effect is consequential in that it leads to payoff equivalence results (between assignments or assignment mechanisms) which do not generally hold otherwise, as discussed next.

5.2 The case of heterogeneous common-value components

Here we provide a counterpart for Proposition 5.3 with heterogeneous common-value components. Consider an assignment market where the common-value components are drawn iid from a CDF F and the idiosyncratic utility components are iid from a CDF G, where F, G have support $[\underline{u}, \overline{u}]$. For any assignment μ the *utility distribution of* μ is a CDF D_μ, where $D_\mu(u)$ denotes the fraction of agents who attain utility at most u in μ. We say a utility distribution D_{μ_1} *stochastically dominates* D_{μ_2} if $D_{\mu_1}(u) \leq D_{\mu_2}(u)$ holds for all $u \geq 0$.

Fact 5.4. *Let $M_1, M_2, \ldots$ be a sequence of markets where $|A| = |O| = n$ in M_n. Let μ_n be an assignment of objects to agents in M_n. Then, when $\lim_{n\to\infty} D_{\mu_n}$ exists, it is stochastically dominated by D^*, where the CDF D^* is defined by $D^*(x) = F(x - \overline{u})$ for all $x \in [\underline{u} + \overline{u}, 2\overline{u}]$.*

[14] For instance, randomly permute the first k objects on all agents' lists and then run RSD.

Hence, D^* is an asymptotic "upper bound" on the utility distribution of any assignment; i.e., no assignment mechanism can, asymptotically, attain a "better" utility distribution than the distribution D^*. The proof is a straight-forward application of the law of large numbers; we omit the proof.

Proposition 5.3 captured the Pareto inefficiency pointed out in its part ii through the notion Pareto-improving cycles. Here, we use the notion of *Pareto-improving pairs*, i.e. cycles of length 2, to capture the Pareto inefficiencies. This is done by showing that there exists a set of pairwise-disjoint Pareto-improving pairs that contain almost all agents. The same holds for Proposition 5.3, as the case of zero common-value components is a special case of the next proposition.

To provide the counterpart for Proposition 5.3, we need to generalize the the definition of the Inefficient-RSD mechanism to the setting with heterogeneous common-value components. The generalization that we consider simply ignores the common-value components, as follows.

The Inefficient-RSD mechanism. For any agent a let $\pi'(a)$ be a list of the objects in which object o appears before object o' if $v^a_o > v^a_{o'}$. Given an increasing function $k(n) : \mathbb{N} \to \mathbb{N}$, run the Inefficient-RSD mechanism, as defined in Section 5.1, with $\{\pi'(a)\}_{a \in N}$ given to that mechanism as the agents' preference lists. The resulting assignment is called the *Inefficient-RSD assignment* in the rest of this section.

Proposition 5.5. *Consider an assignment market where the common-value components are iid from a CDF F and the idiosyncratic utility components are iid from a CDF G, where F, G have finite-valued PDFs with support $[\underline{u}, \overline{u}]$.*

i. (Abundance of competent choices) *When $k = o(n)$, the utility distribution of the Inefficient-RSD assignment converges in probability to D^*.*

ii. When $k = \omega(n^{1/2})$, any fixed agent is involved in at least one Pareto-improving pair in the Inefficient-RSD assignment, whp. Furthermore, at least $n - o(n)$ agents are members of pairwise-disjoint Pareto-improving pairs, whp.

This proposition shows that, with heterogeneous common-value components as well, abundance of competent choices occurs as a result of iid idiosyncratic components. A large family of assignments, including all Pareto-efficient assignments (as shown by [Che and Tercieux 2018]) and assignments with large Pareto inefficiencies (as shown above), attain the best possible payoff distribution, asymptotically. This does not occur when the idd assumption is replaced with negative correlation, as shown by Proposition 5.1.

6 Abundance of competent choices in other setups

In the absence of negative correlation, *abundance of competent choices* is a likely phenomenon in other environments as well. Next, we consider a one-sided environment known as the *hiring problem* [Bruss 2000] (also known as the best choice problem) and, then, stable assignments in marriage markets.

Example: The hiring problem

Consider the *hiring problem* in which an Employer needs to hire precisely one employee from a finite number of applicants, arriving in random order one by one. Employer only knows the number of applicants, n, and there is no recall: she either accepts the present applicant or rejects her, waiting to potentially hire one of the the future applicants.

We compare the solutions to Employer's problem in two cases. Case (i) is a well-studied case where Employer derives utility 1 from hiring the best applicant and utility 0 otherwise. In Case (ii), the utility that Employer derives from hiring each applicant is iid over the interval $[\underline{u}, \overline{u}]$.

[Bruss 2000] provides a short elegant proof that shows the following policy is asymptotically[15] optimal for Case (i): Employer rejects the first n/e applicants (*screens* them), and then hires the first applicant who is preferred to all of the screened applicants. Moreover, the analysis shows that this optimal policy hires the best candidate with probability $1/e$. In Case (ii), however, an asymptotically optimal policy only needs to screen $\omega(1)$ applicants. This guarantees an expected utility of $\overline{u}$ for Employer, asymptotically. The difference between these two cases is, again, due to the abundance of competent choices caused by the iid assumption in Case (ii).[16]

Example: Stable assignments

This section demonstrates that iid idiosyncratic components lead to abundance of competent choices in *marriage markets* as well: there, almost all agents attain their maximum conceivable utility in all *stable* assignments [Lee and Yariv 2018]. This, however, would not be the case when the idiosyncratic components are negatively correlated, or in ordinal util-

[15] In here, asymptotically means as n approaches infinity.

[16] Similar observations hold in other variations. E.g, suppose Employer can choose a subset of the applicants to interview with and hires the best of them. Interviewing $o(n)$ applicants means missing the best applicant with probability 1, asymptotically. However, when Employer's utility from hiring each applicant is drawn iid., interviewing only $\omega(1)$ applicants suffices for attaining utility $\overline{u}$ asymptotically.

ity models. These models capture meaningful gaps between different stable assignments, as discussed next.

A *marriage market* consists of n men and n women. Each man (woman) has a complete strict preference list over women (men). A stable assignment is a bijection μ from men to women with no *blocking pair*, where a blocking pair consists of a man and a woman who are not matched together but prefer each other to their current matches in μ. A *man-optimal* stable assignment in a marriage market is a stable assignment in which every man (weakly) prefers his match in that assignment to any other stable assignment. A *woman-optimal* stable assignment is defined similarly, but for women. Both man-optimal and woman-optimal assignments exist in any marriage market [Gale and Shapley 1962].

A *random marriage market* is a marriage market where the preference lists of all men (women) over women (men) are drawn independently and uniformly at random.

Proposition 6.1. *Consider a random marriage market with n men and n women. Each agent (man or woman) attains utility 1 from being assigned to one of their top k choices, and utility 0 otherwise. When $k \geq \log^2 n$ and $k = o(\frac{n}{\log n})$, in the man-optimal assignment any man and any woman respectively attain utilities $1, 0$, whp, while in the woman-optimal assignment any man and any woman respectively attain utilities $0, 1$, whp.*

This proposition shows that when agents have dichotomous preferences and the conditions on k hold, there is a large gap between the average utilities of men and women in both men-optimal and women-optimal assignments.[17] (While, for expositional simplicity, the focus is on the case where the utilities are 0 or 1, the gap exists for any two distinct utility levels v_1, v_2.) On the other hand, the results of [Lee and Yariv 2018] imply that in any stable assignment both men and women attain average utility approaching 1 (the maximum conceivable utility) when the utilities of men from being matched to women are iid from a distribution F and the utilities of women are defined similarly but drawn from a distribution G.[18] This is in contrast to the observation made in Proposition 6.1.

These findings follow the patterns in the previous sections: iid idiosyncratic components in cardinal utility models can lead to abundance of competent choices. Hence, almost all agents attain their maximum conceivable utility asymptotically. This suggests that models with (negatively) correlated idiosyncratic components can be sharper in distinguishing

[17] We also highlight a related result of [Pittel 1992] that shows in a random matching market, the *average rank* of partners for men and women in the man-optimal assignment is close to $\log n$, $\frac{n}{\log n}$, respectively.

[18] Their findings are more general, e.g. they also allow for (non-zero) common surplus components $c_{m,w}$ for any man m and woman w.

between the assignments and capturing the inefficiencies than models with iid idiosyncratic components, especially in markets without abundance of competent choices.

7 Conclusion

The assumption of iid idiosyncratic components is ubiquitous in the literature on assignment mechanisms and matching markets. We show that, under this assumption, a large family of assignments—including all Pareto-efficient assignments as well as assignments with large Pareto inefficiencies—attain the "best" possible payoff distribution, in the sense of first-order stochastic dominance. We introduce the *abundance of competent choices* as a driving force of such equivalence results under the iid assumption, and show that this phenomenon leads to similar equivalence results in other setups, such as Gale and Shapley's model of marriage markets [Gale and Shapley 1962].

On the other hand, we show that when the iid assumption is replaced with *negative correlation*, the abundance of competent choices does not emerge, and the subsequent payoff equivalence results do not generally hold. In particular, our findings suggest that when there is no abundance of competent choices in a marketplace, models with negatively correlated idiosyncratic components are likely to approximate that marketplace better than models with iid idiosyncratic components. We hope that these findings can provide more refined guidelines for researchers' modeling choices.

References

Atila Abdulkadiroglu, Nikhil Agarwal, and Parag Pathak. 2017. The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match. *American Economic Review* (2017). 2

Nick Arnosti and Peng Shi. 2017. How (Not) to Allocate Affordable Housing. *Working paper (EC2017)* (2017). 2, 4

Itai Ashlagi and Afshin Nikzad. 2018. What matters in school choice tie-breaking? How competition guides design. *Working paper* (2018). 48

Anna Bogomolnaia, Hervé Moulin, and Richard Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165 – 184. https://doi.org/10.1016/j.jet.2004.05.005 10

F. T. Bruss. 2000. Sum the odds to one and stop. *Annals of Probability* 28 (06 2000). https://doi.org/10.1214/aop/1019160340 3, 15

Yeon-Koo Che and Olivier Tercieux. 2018. Payoff equivalence of efficient mechanisms in large matching markets. *Theoretical Economics* 13, 1 (2018), 239–271. https://doi.org/10.3982/TE2793 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE2793 2, 3, 4, 8, 10, 11, 14

Yeon-Koo Che and Olivier Tercieux. 2019. Efficiency and Stability in Large Matching Markets. *Journal of Political Economy* (2019). https://doi.org/10.1086/701791 arXiv:https://doi.org/10.1086/701791 2, 4, 5

Olivier Compte and Philippe Jehiel. 2008. Voluntary Participation and Re-Assignment in Two-Sided Matching. *Working paper* (2008). 4

C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. 1971. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics* 22, 2 (01 Jun 1971), 89–103. https://doi.org/10.1007/BF01651330 4, 5, 10, 20

Alan Frieze and Michal Karonski. 2012. Introduction to Random Graphs. (2012). 37

Alan Frieze and Páll Melsted. 2009. Maximum matchings in random bipartite graphs and the space utilization of Cuckoo Hash tables. *Random Structures & Algorithms* 41, 3 (2009), 334–364. https://doi.org/10.1002/rsa.20427 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/rsa.20427 34

David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15. 1, 3, 16, 17

Yash Kanoria and Daniela Saban. 2017. Facilitating the Search for Partners on Matching Platforms. *Working paper* (2017). 4

Yash Kanoria, Daniela Saban, and Jay Sethuraman. 2015. The Size of the Core in Assignment Markets. *Operations Research* (2015). 2, 4, 5

Donald E. Knuth. 1996. An Exact Analysis of Stable Allocation. *Journal of Algorithms* 20, 2 (1996), 431 – 442. https://doi.org/10.1006/jagm.1996.0020 2, 5, 8, 33, 39

SangMok Lee. 2017. Incentive Compatibility of Large Centralized Matching Markets. *Review of Economic Studies* 84, 1 (2017), 444–463. https://ideas.repec.org/a/oup/restud/v84y2017i1p444-463..html 4

SangMok Lee and Leeat Yariv. 2018. On the Efficiency of Stable Matchings in Large Markets. *Working paper* (2018). https://arxiv.org/abs/1512.07276 2, 3, 4, 5, 15, 16

Svante Linusson and Johan Wästlund. 2004. A proof of Parisi's conjecture on the random assignment problem. *Probability Theory and Related Fields* 128, 3 (01 Mar 2004), 419–440. https://doi.org/10.1007/s00440-003-0308-9 6, 10

Qingmin Liu and Marek Pycia. 2016. Ordinal Efficiency, Fairness, and Incentives in Large Markets. *Working paper* (2016). 5

Boris Pittel. 1992. On likely solutions of a stable marriage problem. *The Annals of Applied Probability* (1992), 358–401. 16, 48, 49, 50

A. E. Roth and Marilda Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Oxford University Press. 48

David W. Walkup. 1980. Matchings in random regular bipartite digraphs. *Discrete Mathematics* 31, 1 (1980), 59 – 64. https://doi.org/10.1016/0012-365X(80)90172-7 5, 8, 9, 28

Douglas West. 1995. Introduction to Graph Theory. (1995). 38, 40

Wikipedia. 2018a. Depth-first search. https://en.wikipedia.org/wiki/Depth-first_search 38

Wikipedia. 2018b. Strongly connected component. https://en.wikipedia.org/wiki/Strongly_connected_component 39

A Preliminaries

This section defines a few notions which will be used in the appendices.

For a positive integer n, we use $[n]$ to denote $\{1, \ldots, n\}$.

Consider a market with set of agents A and objects O where each agent has a complete strict preference list over objects. A *preference profile* is a function $\pi : A \to O^{|A|}$ that determines the preference list of each agent. We use a more compact notation of π_a instead of $\pi(a)$ to denote the preference list of an agent a.

For any lattice $\mathcal{L}$, we denote the set of its element by $V(\mathcal{L})$. For any graph G, we denote the set of its nodes by $V(G)$ and the set of its edges by $E(G)$.

A.1 The FKG inequality

The Fortuin-Kasteleyn-Ginibre (FKG) inequality [Fortuin et al. 1971] is a correlation inequality. Informally, it says that an "increasing event" and a "decreasing event" are negatively correlated, while two "increasing events" are positively correlated.

Let $\mathcal{L}$ be a finite distributive lattice and $\mu : V(\mathcal{L}) \to \mathbb{R}_+$ be a non-negative function on it that satisfies log-supermodularity, i.e. for any two $x, y \in V(\mathcal{L})$,

$$\mu(x \wedge y)\mu(x \vee y) \geq \mu(x)\mu(y),$$

where $\wedge, \vee$ are the meet and join operators of the lattice, respectively.

By the FKG inequality, when $f, g : V(\mathcal{L}) \to \mathbb{R}_+$ are respectively increasing and decreasing functions on the lattice $\mathcal{L}$, it holds that

$$\left(\sum_{x\in X} f(x)g(x)\mu(x)\right)\left(\sum_{x\in X} \mu(x)\right) \leq \left(\sum_{x\in X} f(x)\mu(x)\right)\left(\sum_{x\in X} g(x)\mu(x)\right).$$

The direction of the inequality is reversed when both functions are increasing (or decreasing).

A.2 Chernoff bounds

In the following proofs the Chernoff concentration bounds are used as stated below.

Let $X_1, \ldots, X_n$ be a sequence of n independent random binary variables such that $X_i = 1$ with probability p_i and $X_i = 0$ with probability $1 - p_i$. Also, let $\mu = \sum_{i=1}^{n} \mathbb{E}[X_i]$. Then for any ϵ with $0 \leq \epsilon \leq 1$ we have:

$$\mathbb{P}\left[\sum_{i=1}^{n} X_i > (1+\epsilon)\mu\right] \leq e^{-\epsilon^2\mu/(2+\epsilon)}$$
$$\mathbb{P}\left[\sum_{i=1}^{n} X_i < (1-\epsilon)\mu\right] \leq e^{-\epsilon^2\mu/2}.$$

Furthermore, the former inequality holds for all $\epsilon \geq 0$.

A.3 Asymptotic notations

For two functions $f, g : \mathbb{N} \to \mathbb{R}_+$ we write $f = o(g)$ when $\lim_{n\to\infty} f(n)/g(n) = 0$, $f = \omega(g)$ when $g = o(f)$, $f = \Omega(g)$ when $\lim_{n\to\infty} f(n)/g(n) > 0$, $f = O(g)$ when $g = \Omega(f)$, and $f = \Theta(g)$ when $f = O(g)$ and $g = O(f)$.

We say a sequence of events $E_1, E_2, \ldots$ happen *with high probability* (whp) when $\mathbb{P}[E_n]$ approaches 1 as n approaches infinity. We say the sequence happens *with very high probability* (wvhp) if $n^c \cdot (1 - \mathbb{P}[E_n])$ approaches 0 as n approaches infinity for any constant $c > 0$.

B Proof of Proposition 4.1

First, we prove the claim for the case of unit capacities (i.e. $c = 1$). After that, we present Lemma B.5 which then proves the claim for the case of non-unit capacities (i.e. $c > 1$).

B.1 Proof for the case of unit capacities ($c = 1$)

We need a few definitions for the proof. We use Π to denote the set of all preference profiles in a market of n agents and n objects. Also, let U_Π denote the uniform probability measure over Π.

Let d be a constant independent of n (we will set d to 3). First, we construct a bipartite graph $G[A, O]$. We call A the left side of the graph and O the right side. The set of edges in the graph is $E = E_L \cup E_R$, where E_L and E_R are defined as follows. For any agent $a \in A$, E_R (the set of edges that we draw from left to right) contains d edges that connect a to the first d objects that a ranks on her preference list, i.e. her d most favorite objects.

Define E_L (the set of edges that we draw from right to left) as follows. For any object $o \in O$, let $a^o_1, \ldots, a^o_{n_o}$ denote all agents who have listed object o at position $d + 1$ or worse. Without loss of generality, suppose that $a^o_1, \ldots, a^o_{n_o}$ are ordered in the order that object o appears on their list: a^o_1 has the earliest appearance of o (the most favorable position) and $a^o_{n_o}$ has the latest appearance of o (the least favorable position). Let i_o denote the smallest index such that $i_o \geq d$ and that $a^o_{i_o}$ and $a^o_{i_o+1}$ assign different ranks to object o. If there is no such index i_o, then let $i_o = n_o$. E_L contains the edges $(o, a^o_1,), \ldots, (o, a^o_{i_o})$, for all objects $o \in O$.

The proof is done in two steps. In Step A we show that $G[A, O]$ contains a perfect matching whp, and then we use this fact in Step B to provide an upper bound on the expected average rank of the rank-optimal matching.

Step A: existence of a perfect matching with high probability

We use A to denote the set of agents and O to denote the set of objects. A k-tuple is a pair of sets (X, Y) with $X \subseteq A$ and $Y \subseteq O$ such that $|X| = k$ and $|Y| = n - k + 1$. A k-tuple is *independent* if $(X \times Y) \cap E = \varnothing$. We use p_k to denote the probability that there exists and independent k-tuple. Let $p(n) = \sum_{k=1}^{n} p_k$. We will show that $p(n) \leq n^{-3}$ holds for $n \geq 50$.

To provide an upper bound on p_k, first we fix a k-tuple (X, Y) and provide an upper bound on the chance that that (X, Y) is independent. Without loss of generality, suppose $X = \{1, \dots, k\}$ and $Y = \{1, \dots, l\}$, where $l = n - k + 1$. Let L, R respectively denote the events that $(X \times Y) \cap E_L = \varnothing$ and $(X \times Y) \cap E_R = \varnothing$. (Note that each of the events L, R is just a subset of Π; therefore, we treat L, R as sets, as well as events.)

Suppose $p_{(X,Y)}$ is the probability that (X, Y) is an independent tuple. First, observe that

$$p_{(X,Y)} = \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k \cdot \mathbb{P}\left[L|R\right]. \tag{B.1}$$

The first factor in the RHS of the above equation is just equal to $\mathbb{P}[R]$. We use the next two claims to simplify the RHS of the above equation.

Claim B.1. $\mathbb{P}[L|R] \leq \mathbb{P}[L]$.

Proof. The proof uses the FKG inequality to show that the events L and R are negatively correlated. To use the inequality, we first define a distributive lattice. To define the lattice, we need a few definitions. For any object i, let $f(i) = i$ if $i > d$ and let $f(i) = n+i$ otherwise. Define a total order $\preceq$ on the set of objects $O = \{1, \dots, n\}$ as $i \preceq j$ iff $f(i) < f(j)$. With slight abuse of notation, for any two vectors of the same size, namely $u = (u_1, \dots, u_s), v = (v_1, \dots, v_s)$, we write $u \preceq v$ iff $u_j \preceq v_j$ holds for all j.

The lattice that we define for applying the FKG inequality is denoted by $\mathcal{L}[l]$. Each element of the lattice, namely x, is an ordered list of n l-dimensional vectors, namely $x_1, \dots, x_n$. Since l remains fixed in the proof of this claim, we drop the argument and denote the lattice simply by $\mathcal{L}$. As we will see, x_a corresponds to agent a for any $a \in A$. Each vector, namely $x_a = (x_a^1, \dots, x_a^l)$, contains l distinct integers belonging to O. For $a \leq k$, the integers are ordered in decreasing order with respect to $\preceq$, and for $a > k$, they are ordered in increasing order with respect to $\preceq$. For any two elements of the lattice x, y, we have $x \preceq_{\mathcal{L}} y$ iff

$$\begin{cases} y_a \preceq x_a, \ \forall a \in \{1, \dots, k\} \\ x_a \preceq y_a, \ \forall a \in \{k+1, \dots, n\}. \end{cases}$$

Observe that the defined partial order $\preceq_{\mathcal{L}}$ over $V(\mathcal{L})$ is a lattice with the meet ($\wedge$) and join ($\vee$) operators being the component-wise minimum and maximum with respect to $\preceq$, respectively. Figure 2 provides a graphical representation of the lattice elements.

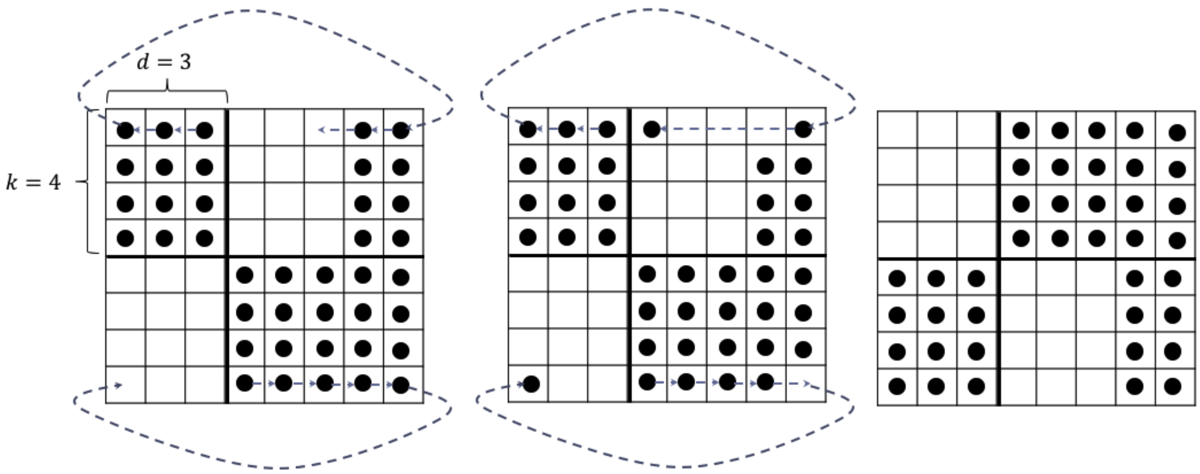

Figure 2: From left to right: the smallest element of the lattice, an element between the smallest and the largest element, and the largest element. We have $A = [8]$ and $O = [8]$. Rows and columns correspond to the elements of A, O, respectively. Each row $a \in A$ contains $l = 5$ marked cells which determine the l-dimensional vector x_a corresponding to that row. The dotted arrows determine the directions for "moving" in the lattice from a smaller to a larger element.

Next, we link the definition of the lattice to the set of all preference profiles π of the assignment problem. Each preference profile π is *represented* by an element of the lattice, which we denote by $\mathcal{L}(\pi)$, and define it as follows. Given π, for each agent a, construct an l-dimensional vector x_a that contains the positions (ranks) of objects $1, \dots, l$ on π_a. Order the elements of this vector in decreasing order with respect to $\preceq$ if $a \le d$, and otherwise, order them in increasing order with respect to $\preceq$. Let $\mathcal{L}(\pi) = \langle x_1, \dots, x_n \rangle$.

Observe that, given any preference profile π, all the information that we need to assess whether events R, L happen at π are coded in $\mathcal{L}(\pi)$. That is, $\mathcal{L}(\pi)$ is a sufficient statistic for detecting whether events L, R happen in π. This observation is used in the last step of the proof, as follows.

Define two functions $f_L, f_R : \mathcal{L} \to \{0,1\}$ such that $f_L(x) = 1$ iff the event L holds at element x of the lattice and $f_R(x) = 1$ iff the event R holds at element x. A straightforward coupling argument shows that the functions f_L, f_R, respectively, are increasing and decreasing with respect to the partial order $\preceq$ defined over $V(\mathcal{L})$. Also, observe that the probability distribution induced by U_Π over $\{\mathcal{L}(\pi) : \forall \pi\}$ is a uniform distribution itself, and therefore, it satisfies the log-supermodularity condition required for the FKG inequality. The inequality implies that

$$\mathbb{E}\left[f_L(x) f_R(x)\right] \leq \mathbb{E}\left[f_L(x)\right] \mathbb{E}\left[f_R(x)\right],$$

i.e. the events L and R are negatively correlated. □

Claim B.2. $\mathbb{P}[L] \leq \prod_{i=1}^{l} \mathbb{P}[L_i]$, *where* L_i *is the event that there is no edge in* E_L *that connects a node* $i \in Y$ *to a node in* X.

Proof. The proof is by induction. For any $j \leq l$, we will show that

$$\mathbb{P}\left[L_1 \wedge L_2 \wedge \ldots \wedge L_j\right] \leq \prod_{i=1}^{j} \mathbb{P}\left[L_i\right].$$

The induction basis for $j = 1$ is trivial. The induction step supposes that, for some $l' < l$, the claim holds for all $j \leq l'$, and then proves the claim for $j = l' + 1$.

The proof of the induction step is by an application of the FKG inequality. First, we need a few definitions.

Definitions. The lattice which will be used in the application of the FKG inequality is $\mathcal{L}[l']$. (Recall the definition from the proof of Claim B.1.) For brevity, we drop the argument l' and denote the lattice simply by $\mathcal{L}$ throughout this proof.

Define a *t-subprofile* to be a partially filled preference profile in the following sense: each row of the preference profile has $n - t$ empty positions. The rest of the positions in the row are filled with numbers $1, \ldots, t$, with each number appearing precisely once. The set of *filled indicies* in a t-subprofile is simply the set of all pairs (a, i) where position i of agent a's preference list is filled. Note the correspondence between the notions of an l'-subprofile and the lattice $\mathcal{L}$: each element of the lattice corresponds to the set of filled indices of precisely $(l'!)^n$ l'-subprofiles.

An *Extension* of an element of the lattice, namely x, to an l'-subprofile is defined as follows. Recall that $x = (x_1, \ldots, x_n)$ where for each $a \in A$, $x_a = (x_a^1, \ldots, x_a^{l'})$ is a vector that contains l' distinct integers belonging to O. Suppose that π, which is initially empty, denotes the l'-subprofile which is to be constructed from x. To extend x to an l'-subprofile, assign a number from $[l']$ to each preference list position $\pi_a(x_a^i)$ for all $a \in A$ and $i \in [l']$, such that all the numbers assigned to any preference list π_a are distinct (i.e. such that for any $a \in A$, the set $\{\pi_a(x_a^i) : i \in [l']\}$ has size l').

An *Extension* of a j-subprofile to a complete preference profile is defined in the natural way: by filling out the empty positions in the j-subprofile in a way that it creates a valid preference profile.

Let K_i^j denote the set of j-subprofiles π' such that there is an extension of π' to a complete preference profile, namely π'', with $\pi'' \in L_i$.[19] (Observe that if there exists one such extension of π' satisfying this condition, then *any* of its extensions to a complete preference profile also satisfy this condition, so long as $j \geq i$.) Define

$$K^j = K_1^j \cap \ldots \cap K_{l'}^j.$$

For an element x of the lattice, define $f(x)$ to denote the number of extensions of x to an l'-subprofile belonging to $K^{l'}$. Also, let $g(x)$ denote the number of extensions of x to an $(l'+1)$-subprofile belonging to $K_{l'+1}^{l'+1}$ divided by $(l'!)^n$. Intuitively, $g(x)$ is the number of extensions of x to a $(l'+1)$-subprofile belonging to $K_{l'+1}^{l'+1}$ in which each of the positions that contain an element of $[l']$ is filled with an '*' instead. (Note that only the position of the object $l'+1$ in the subprofile determines whether it belongs to $K_{l'+1}^{l'+1}$.)

A straight-forward coupling argument shows that the functions $f, g : V(\mathcal{L}) \to \mathbb{R}_+$ are decreasing and increasing functions, respectively.

The induction step. Let U, V denote the uniform measures induced over the set of all preference profiles and over $V(\mathcal{L})$, respectively. Since the functions f, g are respectively decreasing and increasing, and since the measure U is uniform (and therefore, log-supermodular), then the FKG inequality implies that

$$\mathbb{E}_U\left[f(x)g(x)\right] \leq \mathbb{E}_U\left[f(x)\right]\mathbb{E}_U\left[g(x)\right]. \tag{B.2}$$

[19] Recall the equivalence between events and subsets of Π.

We will complete the proof by rewriting the above inequality. To this end, let

$$\alpha = |V(\mathcal{L})| \cdot (l'!)^n,$$
$$\beta = (n - l')^n,$$
$$\gamma = (n - l' - 1)^n,$$

and define the events

$$A = L_1 \wedge \ldots \wedge L_{l'},$$
$$B = L_{l'+1}.$$

Observe that

$$\mathbb{P}_U[A] = \frac{\sum_x \beta f(x)}{\alpha\beta} = \frac{\mathbb{E}_V[f(x)]}{(l'!)^n},$$
$$\mathbb{P}_U[B] = \frac{\sum_x \gamma \cdot (l')^n \cdot g(x)}{\alpha\beta} = \frac{\gamma \cdot \mathbb{E}_V[g(x)]}{(n-l')^n},$$
$$\mathbb{P}_U[A \wedge B] = \frac{\sum_x \gamma f(x) g(x)}{\alpha\beta} = \frac{\gamma \cdot \mathbb{E}_V[f(x)g(x)]}{(l'!)^n (n-l')^n},$$

where all the sums are taken over $x \in V(\mathcal{L})$. The above equalities together with (B.2) imply that

$$\mathbb{P}_U[A \wedge B] \leq \mathbb{P}_U[A] \cdot \mathbb{P}_U[B], \tag{B.3}$$

which completes the induction step and proves the claim. □

We can provide an upper bound on the RHS of (B.1) using the above two claims.

Claim B.3.

$$p_{(X,Y)} \leq \left[\frac{\binom{k-1}{d}}{\binom{n}{d}}\right]^k (\omega + \delta)^l, \tag{B.4}$$

where $\omega = \left[\binom{n-k}{d} / \binom{n}{d}\right]^{n-k+1}$ *and* $\delta = \frac{d^{n-2d+1} e^d}{n^{n-2d}}$.

Proof. By Equation (B.1) and Claims B.1 and B.2, it suffices to show that $\mathbb{P}[L_i] \leq \omega + \delta$ for any $i \in Y$. To this end, let D be the event in which object i has a rank worse than d in at least d of the agents' preference lists. Observe that

$$\mathbb{P}[L_i] = \mathbb{P}[L_i|D] \times \mathbb{P}[D] + \mathbb{P}\left[L_i|\overline{D}\right] \times \mathbb{P}\left[\overline{D}\right].$$

We will show that (i) $\mathbb{P}\left[\overline{D}\right] \leq \delta$ and (ii) $\mathbb{P}[L_i|D] \leq \omega$, which would prove that $\mathbb{P}[L_i] \leq \omega+\delta$.

Step (i) $\mathbb{P}\left[\overline{D}\right] \leq \delta$. Let $\overline{D}_j$ be the event in which object i a rank worse than d in precisely j of the agents' preference lists. Observe that

$$\mathbb{P}\left[\overline{D}_j\right] = \binom{n}{j}\left(\frac{d}{n}\right)^{n-j}\left(1-\frac{d}{n}\right)^{j}.$$

Therefore, we can write

$$\mathbb{P}\left[\overline{D}\right] = \sum_{j=0}^{d-1} \mathbb{P}\left[\overline{D}_j\right] \leq d\binom{n}{d}\left(\frac{d}{n}\right)^{n-d} \leq \frac{d^{n-2d+1}e^{d}}{n^{n-2d}} = \delta, \tag{B.5}$$

where in the last inequality we have used the bound $\binom{n}{d} \leq \left(\frac{ne}{d}\right)^d$.

Step (ii) $\mathbb{P}[L_i|D] \leq \omega$. let A_j be the set of agents who rank object i on the j-th position of their preference list, and let $A^h = \cup_{j=d+1}^{h} A_j$ be such that h is the smallest number for which $|A^h| \geq d$. Observe that A^h is a random variable whose distribution, conditional on its size being equal to x, is the uniform distribution over the set of all subsets of A with size x. (This holds by symmetry.) Therefore, we can write

$$\mathbb{P}\left[L_i \middle| g = |A^h|\right] \leq \frac{\binom{n-k}{g}}{\binom{n}{g}}.$$

Since the RHS of the above equality is decreasing in g, and since $|A_h| \geq d$, therefore we can write

$$\mathbb{P}[L_i|D] \leq \frac{\binom{n-k}{d}}{\binom{n}{d}} = \omega$$

□

We now use Claim B.3 and a union bound to write

$$p_k \le \binom{n}{k}\binom{n}{l}\left[\frac{\binom{k-1}{d}}{\binom{n}{d}}\right]^k (\omega+\delta)^l$$
$$\le \binom{n}{k}\binom{n}{l}\left[\frac{\binom{k-1}{d}}{\binom{n}{d}}\right]^k \left(\omega^l + 2^l\delta\right),$$

where recall that $l = n-k+1$, and the last inequality holds since $\omega, \delta \le 1$. Using the above inequality, we can bound $p(n)$ as follows:

$$p(n) = \sum_{k=1}^{n} p_k \le \sum_{k=1}^{n}\binom{n}{k}\binom{n}{l}\left[\frac{\binom{k-1}{d}}{\binom{n}{d}}\right]^k \left(\omega^l + 2^l\delta\right)$$
$$= \sum_{k=1}^{n}\binom{n}{k}\binom{n}{k-1}\left[\frac{\binom{k-1}{d}}{\binom{n}{d}}\right]^k \cdot \omega^l + \sum_{k=1}^{n}\binom{n}{k}\binom{n}{k-1}\left[\frac{\binom{k-1}{d}}{\binom{n}{d}}\right]^k \cdot 2^l\delta \quad \text{(B.6)}$$

Observe that (B.6) has two summands. Let S_1, S_2 denote the first and the second summand, respectively. [Walkup 1980] shows that

$$S_1 \le \frac{1}{122}\left(\frac{d}{n}\right)^{(d+1)(d-2)}.$$

We finish the proof by providing an upper bound on S_2. Observe that

$$S_2 \le 2^{3n}\delta = 2^{3n} \cdot \frac{d^{n-2d+1}e^d}{n^{n-2d}} \quad \text{(B.7)}$$

In particular, letting $d = 3$ and summing up S_1 and S_2 implies that for sufficiently large n, we have $p(n) \le n^{-3}$. (While the crude bounds we have used provide that $n \ge 50$ is sufficiently large, the right-hand side 50 can be reduced with a more careful analysis.)

Step B: Bounding the average rank in the perfect matching

Recall that $G[A, O]$ denotes the bipartite graph constructed in Step i. We showed that, for sufficiently large n, $G[A, O]$ has a perfect matching with probability at least $1 - n^{-3}$. To bound the expected sum of the ranks in the matching, we provide an upper bound on $\sum_{o \in O} w_o$, where w_o is the weight of the maximum-weight edge adjacent to $o \in O$. (Define $w_o = 0$ if there are no edges incident to o.)

Recall that U_Π denotes the uniform distribution over Π. We can show that, for any o, $\mathbb{E}_{U_\Pi}[w_o]$ is bounded by $d + dz$, where $z = e/(e-1)$. This will imply that, for $d = 3$, the expected average rank is at most

$$(1 - n^{-3}) \cdot \frac{n(d + 3e/(e-1))}{n} + n^{-3} \cdot n^2 < 7.75$$

for $n \geq 50$.

The formal proof is presented below. For notational simplicity, we drop the subscript U_Π from the notations $\mathbb{E}_{U_\Pi}[\cdot]$ and $\mathbb{P}_{U_\Pi}[\cdot]$. Also, all the inequalities that we write below hold for sufficiently large n (i.e. $n \geq 50$). Let $w = \sum_{o \in O} w_o$. Let $\mathfrak{m}$ denote the event that $G[A, O]$ has a perfect matching. By our analysis in Step A, we have

$$\mathbb{E}[w|\mathfrak{m}] \leq \frac{\mathbb{E}[w]}{1 - n^{-3}}.$$

Therefore, to complete Step B, it suffices to bound $\mathbb{E}[w]$. The following claim implies that $\mathbb{E}[w] \leq d + 3e/(e-1)$, which would imply that

$$\mathbb{E}[w|\mathfrak{m}] \leq \frac{d + 3e/(e-1)}{1 - n^{-3}} < 7.75$$

for $d = 3$ and $n \geq 50$.

Claim B.4. *For any object $o \in O$, $\mathbb{E}[w_o] \leq d + 3e/(e-1)$.*

Proof. Let $\mathcal{M}$ be an $|A| \times |O|$ matrix that contains π_a as its a-th row for each agent $a \in A$. For any integer $i > d$, let H_i^j denote the event that object o appears exactly j times in columns $d+1, \ldots, i$ of the matrix $\mathcal{M}$.

Observe that

$$\mathbb{P}\left[H_{d+1}^0\right] = (1 - 1/n)^n \leq 1/e.$$

Also, observe that for all $i \geq d + 1$,

$$\mathbb{P}\left[H_{i+1}^0 | H_i^0\right] < (1 - \frac{1}{n})^n \leq 1/e,$$
$$\mathbb{P}\left[H_{i+1}^2 | H_i^1\right] < (1 - \frac{1}{n})^n \leq 1/e,$$

and that

$$\mathbb{P}\left[H_{i+1}^3 | H_i^2\right] \leq \begin{cases} (1-\frac{1}{n-2})^{n-2} \leq 1/e, & \text{for } i > d+1 \\ (1-\frac{1}{n-1})^{n-2} \leq (1/e)\cdot\frac{n-1}{n-2}, & \text{for } i = d+1. \end{cases}$$

Define the random variable H as follows: let $H = i$ where i is the smallest index for which the event $H_i^3 \vee \ldots \vee H_i^n$ holds. If such i does not exist, then let $H = 0$. Observe that when $H > 0$, $H = w_o$, by definition. The above four inequalities imply that $H - d$ is stochastically dominated by the sum of three independent geometric random variables with means $\frac{1}{1-1/e}$, $\frac{1}{1-1/e}$, and $\frac{1}{1-\frac{n-1}{n-2}/e}$. In fact, with a straight-forward coupling argument it is possible to slightly refine this bound and eliminate the coefficient $\frac{n-1}{n-2}$, i.e. it is possible to show that $H - d$ is first-order stochastically dominated by sum of three independent geometric random variables each with mean $\frac{1}{1-1/e}$. Therefore, $\mathbb{E}[w_o] = \mathbb{E}[H] \leq d + \frac{3e}{e-1}$. □

B.2 Proof for the case of non-unit capacities ($c > 1$)

First, we need a few definitions. The *rank distribution* of an assignment $\mu : A \to O$ is a function $\mathcal{R} : \mathbb{N} \to \mathbb{R}_+$ where $\mathcal{R}(i)$ denotes the fraction of agents who are assigned to their i-th favorite object (i.e. the object that they rank i-th). Given a random market M, the *expected rank distribution* of the rank optimal assignment in M is a function $\mathcal{R} : \mathbb{N} \to \mathbb{R}_+$ where $\mathcal{R}(i)$ denotes the expected fraction of agents who are assigned to their i-th favorite object in the rank-optimal assignment.

We say a rank distribution $\mathcal{R}' : \mathbb{N} \to \mathbb{R}_+$ stochastically dominates another rank-distribution $\mathcal{R} : \mathbb{N} \to \mathbb{R}_+$ when for any $r \in \mathbb{N}$,

$$\sum_{i=1}^{r} \mathcal{R}'(i) \geq \sum_{i=1}^{r} \mathcal{R}(i).$$

The next lemma completes the proof of Proposition 4.1, as it shows that an upper bound on the expected average rank for the case of $c = 1$ is a valid upper bound for the general case ($c \geq 1$) as well.

Lemma B.5. *Let M, M' be two random markets both with n agents. Suppose that M has n objects each with capacity 1 and M' has n/c objects each with capacity $c \geq 1$, respectively. Let $\mathcal{R}, \mathcal{R}'$ denote the expected rank distributions of the rank-optimal assignments in M, M', respectively. Then, $\mathcal{R}'$ stochastically dominates $\mathcal{R}$.*

Proof. For any preference profile, π, let r_π denote the average rank in the rank-optimal assignment when the preference profile is π. Let Π, Π' denote the set of possible preference profiles in M, M'. The proof works by defining a function $f : \Pi \to \Pi'$ that maps every $\frac{|\Pi|}{|\Pi'|}$ elements of Π to precisely one element of Π. Moreover, this function is defined such that, for any $\pi \in \Pi$, the rank distribution of the rank-optimal assignment for π is stochastically dominated by the rank distribution of the rank-optimal assignment for $f(\pi)$. The existence of such a function will prove the claim.

In the rest of the proof, we define the function. Let n be the number of agents in both markets. In the market M, relabel the objects as $O = \{o_t^j : j \in [c], t \in [n/c]\}$. We say an object o is of type t if $o = o_t^j$ for some j, t. With slight abuse of notation, we use $t(o)$ to denote the type of an object o. Given a preference list σ over O, define $\overline{\sigma}$ to be a list in which $\overline{\sigma}(i) = t(\sigma(i))$. (That is, each object is replaced with its type.)

Let $g(\sigma)$ be the preference profile over $[n/c]$ defined as follows: in its i-th position, $g(\sigma)$ contains the i-th distinct number that appears in $\overline{\sigma}$, for $i \in [n/c]$. In other words, the function g removes the second, third, and the higher appearances of a number in $\overline{\sigma}$ and outputs the resulting list.

For each preference profile $\pi \in \Pi$, define $f(\pi)$ to be the preference profile π' where $\pi'_a = g(\overline{\pi_a})$ for all $a \in A$. Observe that, by symmetry, $|f^{-1}(\pi')|$ does not depend on π', i.e. $|f^{-1}(\pi')| = \frac{|\Pi|}{|\Pi'|}$. To complete the proof, it remains to show that the rank distribution of the rank-optimal assignment for π is stochastically dominated by the rank distribution of the rank-optimal assignment for $f(\pi)$, for all preference profiles $\pi \in \Pi$. Let μ be the rank-optimal assignment for the preference profile π. We define the assignment μ' in the market M' for the preference profile π' as follows: for each agent a, let $\mu'(a) = t(\mu(a))$. Observe that μ' is a feasible assignment, and that $\mu'(a)$ does not have a worse rank in π'_a than $\mu(a)$ has in π_a. Therefore, the rank distribution of μ' stochastically dominates the rank distribution of μ. This finishes the proof.

□

This completes the proof of Proposition 4.1. The above lemma together with our analysis in Section B.1 immediately imply the following corollary.

Corollary B.6 (of Proposition 4.1)**.** *The expected average rank in the rank-optimal assignment is at most* $7\frac{3}{4}$ *for all* $n \geq 50$*.*

B.3 Distribution of the rank of an agent's object

This section shows that the "upper bound" distribution for the rank of an agent's object in the assignment found by Proposition 4.1 is $3 + X + Y + Z$, where X, Y, Z are iid geometric random variables with mean $\frac{e}{e-1}$. We remark that the mean of this upper bound distribution, and thus the average rank of the assignment, is less than $7\frac{3}{4}$, as noted in Section 4.1.

Let μ_π denote the assignment that we find for a given preference profile π in the proof of Proposition 4.1. Also, let μ be a random variable that denotes μ_π, where $\pi \in \Pi$ is drawn uniformly at random. For any fixed agent $a \in A$, our analysis provides an "upper bound" for the distribution of $r_a(\mu(a))$, as follows.

Formally, the *rank distribution of an agent* a is a function $\mathcal{R} : \mathbb{N} \to \mathbb{R}_+$ where $\mathcal{R}(i)$ denotes the fraction preference profiles π for which $r_a(\mu_\pi(a)) = i$. We say a rank distribution $\mathcal{R}' : \mathbb{N} \to \mathbb{R}_+$ stochastically dominates another rank-distribution $\mathcal{R} : \mathbb{N} \to \mathbb{R}_+$ when for any $r \in \mathbb{N}$,

$$\sum_{i=1}^{r} \mathcal{R}'(i) \geq \sum_{i=1}^{r} \mathcal{R}(i).$$

We emphasize the difference with the definition of first-order stochastically dominance, where the direction of the inequality is reversed. The reason is that, here, we are concerned with ranks (rather than utilities) where lower numbers are preferred to higher ones.

Claim B.4 finds a distribution that stochastically dominates the rank distribution of a, asymptotically. It shows that this rank distribution stochastically dominates (in the sense defined above) the distribution of $3 + X + Y + Z$, where X, Y, Z are iid geometric random variables with mean $\frac{e}{e-1}$. For completeness, we review that proof here.

The key point is that, by the construction of the bipartite graph G, the edge $(a, \mu(a))$ belongs to $E_L \cup E_R$. If it belongs to E_R, then $\mu(a) \leq 3$. Otherwise, if it must belong to E_L. In the latter case, we show that $r_a(\mu(a))$ stochastically dominates $3 + X + Y + Z$, which would prove the claim. Before continuing the argument, we first recall the definition of E_L.

The set of edges E_L is defined as follows. For any object $o \in O$, let $a_1^o, \ldots, a_{n_o}^o$ denote all agents who have listed object o at position $d + 1$ or worse. (We recall that d is set to 3 at the end of the proof.) Without loss of generality, suppose that $a_1^o, \ldots, a_{n_o}^o$ are ordered in the order that object o appears on their list: a_1^o has to earliest appearance of o (the most favorable position) and $a_{n_o}^o$ has the latest appearance of o (the least favorable position). Let i_o denote the smallest index such that $i_o \geq d$ and that $a_{i_o}^o$ and $a_{i_o+1}^o$ assign different ranks to object o. If there is no such index i_o, then let $i_o = n_o$. E_L contains the edges $(o, a_1^o,), \ldots, (o, a_{i_o}^o)$, for all objects $o \in O$.

We also recall the following definitions from Claim B.4. $\mathcal{M}$ is an $|A| \times |O|$ matrix that contains π_a as its a-th row for each agent $a \in A$. For any integer $i > d$, let H_i^j denote the event that object o appears exactly j times in columns $d+1, \ldots, i$ of the matrix $\mathcal{M}$. Define the random variable H as follows: let $H = i$ where i is the smallest index for which the event $H_i^3 \vee \ldots \vee H_i^n$ holds. If such i does not exist, then let $H = 0$. (In that case, we have $\mu(a) \leq 3$ when μ exists, and hence we can set $H = 0$.) Observe that when $H > 0$, H is the rank of the least favorite object of a that is connected to a in G. By Claim B.4, $H - d$ is first-order stochastically dominated by the sum of three independent geometric random variables with mean $\frac{e}{e-1}$. Hence, $r_a(\mu(a))$ stochastically dominates $3+X+Y+Z$, as claimed above. (Recall from above our distinct definitions for stochastic dominance and first-order stochastic dominance.)

C Proof of Proposition 4.2

Consider the last m agents who choose (i.e. the m agents with lowest priority numbers). Let them be indexed by $a_0, \ldots, a_{m-1}$, ordered with respect to their priority numbers with a_0 having the best priority number and a_{m-1} having the worst. Also, let R_i denote the average rank of agent a_i, and $R = \frac{1}{m} \cdot \sum_{i=0}^{m-1} R_i$.

To provide a lower bound on R_i, we define an auxiliary problem instance, which is just running RSD on a market with m agents, namely $a'_0, \ldots, a'_{m-1}$, and m objects with unit capacities. Suppose the agents rank objects independently and uniformly at random, and that agents choose objects in the same order as their indices: agent a'_0 chooses the first object. Let R'_i denote the average rank of agent a'_i, and $R' = \frac{1}{m} \cdot \sum_{i=0}^{m-1} R'_i$. [Knuth 1996] shows that $R' \geq \ln m - 1$. The next claim states that $R_i \geq R'_i$, which would imply $R \geq R'$. That would complete the proof, as it shows that the expected average rank in the original instance is at least $R' \cdot \frac{m}{n} \geq \frac{\ln m - 1}{c}$.

Claim C.1. *For any $i \in \{0, \ldots, m-1\}$, $R_i \geq R'_i$.*

Proof. When agent a'_i is choosing in the auxiliary problem, there are exactly i objects allocated by the agents before her, and therefore $m - i$ possible choices remain. In the original problem, when agent a_i is choosing, at most i objects have a positive number of copies remaining; therefore, agent a_i has at most $m - i$ possible choices. This implies that the rank distribution for agent a'_i stochastically dominates the rank distribution for agent a_i, which implies that $R_i \geq R'_i$. □

D Proofs from Section 5

D.1 Proof of Proposition 5.1

We prove that the *expected* difference between the average utility in the utility-optimal assignment and the average utility in the RSD assignment is at least $\frac{v_1 - v_2}{20}$. The high probability result follows from a straight-forward application of Chernoff bounds.[20]

First, we compute an upper bound for average utility under the RSD mechanism. For any $i \geq 3$, the $i+1$-th agent who chooses in the RSD mechanism attains utility v_2 with probability

$$\frac{i}{n} \cdot \frac{i-1}{n-1} \cdot \frac{i-2}{n-2} \geq (\frac{i-2}{n-2})^3.$$

Therefore, the expected number of agents who attain utility v_2 is at least

$$\sum_{i=1}^{n-2} (\frac{i}{n-2})^3 = (n-2)^{-3} \cdot \left(\frac{(n-2)(n-1)}{2} \right)^2 \geq \frac{n}{4}.$$

This shows that the expected average utility under the RSD mechanism is at most $\frac{3}{4} \cdot v_1 + \frac{1}{4} \cdot v_2$.

Next, we show that, whp, there exists a Pareto-optimal assignment with average utility at least $\frac{4}{5} \cdot v_1 + \frac{1}{5} \cdot v_2$.[21] To this end, construct the bipartite graph $G[A', O]$ where $A' = [\lceil \frac{4n}{5} \rceil]$. For any agent $a \in A'$, add 3 edges to $G[A', O]$ that connect a to the first, second, and the third object that a lists on her preference list. Since the preference lists of agents are drawn independently and uniformly at random, then Theorem 3 of [Frieze and Melsted 2009] implies that $G[A', O]$ has a matching of size $|A'|$, whp. Therefore, whp, there exists an assignment in the original random market with average utility at least $\frac{4}{5} \cdot v_1 + \frac{1}{5} \cdot v_2$. This completes the proof.

D.2 Proof of Proposition 5.2

Let F denote the CDF from which the random variables are drawn. Fix an arbitrary small $\epsilon > 0$, and let I denote the interval $[\overline{u} - \epsilon, \overline{u}]$. Note that since F has support $[\underline{u}, \overline{u}]$, $\delta > 0$.

[20] The presented analysis for the utility-optimal assignment is a high-probability statement. The application of Chernoff bounds is required for the RSD analysis.

[21] In fact, we can derive a tighter bound of $(\frac{4}{5} + \epsilon) \cdot v_1 + (\frac{1}{5} - \epsilon) \cdot v_2$ for some $\epsilon > 0.01$.

Also, let $\delta = 1 - F(\overline{u} - \epsilon)$. Observe that

$$\begin{aligned}
\mathbb{P}\left[v \notin I\right] &= (1-\delta)^n + \sum_{i=1}^{k-1} \binom{n}{i} \cdot \delta^i (1-\delta)^{n-i} \\
&\leq (1-\delta)^n + \sum_{i=1}^{k-1} \delta^i \cdot e^{i \log(\frac{ne}{i})} (1-\delta)^{n-i} \\
&\leq (1-\delta)^n + \sum_{i=1}^{k-1} \delta^i \cdot e^{i \log(\frac{ne}{i}) - (n-i)\log(\frac{1}{1-\delta})} \\
&= (1-\delta)^n + \sum_{i=1}^{k-1} \delta^i \cdot s_i
\end{aligned} \tag{D.1}$$

where $s_i = i \log(\frac{ne}{i}) - (n-i)\log(\frac{1}{1-\delta})$. (In the first inequality above, we have used the fact that $\binom{n}{i} \leq (\frac{en}{i})^i$.)

Claim D.1. *For any fixed $i > 0$ (i.e. i can depend on ϵ but not on n),* $\lim_{n\to\infty} s_i = 0$.

Proof. It suffices to show that

$$\lim_{n\to\infty} i \log(\frac{en}{i}) - (n-i) \cdot \log(\frac{1}{1-\delta}) = -\infty.$$

To see why the above equation holds, note that $i \leq k$ and $k = o(n)$, which imply that

$$\begin{aligned}
i \log(\frac{en}{i}) &= o(n), \\
(n-i) \cdot \log(\frac{1}{1-\delta}) &= \Theta(n),
\end{aligned}$$

where the first equation holds by L'Hospital's rule. □

Claim D.1 together with (D.1) imply that, for any sufficiently large fixed $j > 0$,

$$\lim_{n\to\infty} \mathbb{P}\left[v \notin I\right] \leq \delta^j.$$

Since j can be any sufficiently large fixed number, the above inequality implies that

$$\lim_{n\to\infty} \mathbb{P}\left[v \notin I\right] = 0. \tag{D.2}$$

Observe that the above equation holds for any arbitrary small ϵ (because we imposed no restrictions on ϵ and the CDF F is strictly increasing and continuous). Therefore, by (D.2), we have shown that v converges in probability to $\overline{u}$, i.e. for all $\epsilon > 0$ we have

$$\lim_{n\to\infty} \mathbb{P}\left[|v - \overline{u}| > \epsilon\right] = 0.$$

D.3 Proof of Proposition 5.3

Let a_i denote the i-th agent who chooses an object. Suppose π_i is the ordinal preference list of a_i, i.e. the list of the objects ordered in a decreasing order with respect to their utilities for agent a_i. Let r_i be a random variable denoting the rank of the object assigned to a_i on π_i.

Lemma D.2. *In expectation, at least $n/2$ agents are assigned to their k-th ranked object. Furthermore, for any fixed $\epsilon > 0$, at least $n(1-\epsilon)/2$ and at most $n(1-\epsilon)/2$ agents are assigned to their k-th ranked object, wvhp.*

Proof. Observe that

$$\mathbb{P}\left[r_{i+1} \neq k\right] = \frac{i}{n},$$

which implies that the expected number of students not assigned to their k-th ranked object is $\frac{n-1}{2}$. This proves the first part of the claim. A standard application of Chernoff concentration bounds proves the second part. □

D.3.1 Proof of Proposition 5.3, part (i)

Let $r^* = \max\{k(n), \sqrt{n}\}$, where we have suppressed the argument of $r^*(n)$ for notational brevity. Observe that $r^* = o(n)$. Let p_i denote the probability that $r_i > r^*$. First, we provide an upper bound on p_i. Observe that for $i \leq r^*$, $p_i = 0$. For $i \geq r^*$,

$$p_{i+1} = \frac{i}{n} \cdot \frac{i-1}{n-1} \cdots \frac{i-r^*}{n-r^*} \leq \left(\frac{i}{n}\right)^{r^*} \leq e^{-\frac{r^* j}{n}}, \tag{D.3}$$

where $j = n - i$. Observe that, for any j such that $j = \omega(n/r^*)$, we have $p_{i+1} = o(1)$. Choose such arbitrary j, namely j^*. Therefore, (D.3) implies that for all $i \leq n - j^*$, we have $p_{i+1} = o(1)$.

This fact, together with Proposition 5.2 imply that that the expected utility of any agent under the Inefficient-RSD mechanism is at least $\frac{n-j^*}{n} \cdot (1 - o(1)) \cdot \overline{u}$, which approaches $\overline{u}$ as n approaches infinity.

D.3.2 Proof of Proposition 5.3, part (ii)

Let $\mu : A \times O$ denote the assignment generated by the Inefficient-RSD mechanism. Let A' denote the subset of agents who are assigned to the object ranked k-th on their list. Without loss of generality, suppose that $A' = \{a_1, \ldots, a_N\}$. Recall that, by Lemma D.2, $N \geq \frac{n(1-\epsilon)}{2}$ holds whp, for any $\epsilon > 0$. Let $o_i = \mu(a_i)$ for all $i \in [N]$, and let $O' = \{o_1, \ldots, o_N\}$.

Construct a directed graph G as follows. Let the set of nodes of G be O'. For any $o_i \in O'$, add a directed edge to G from o_i to o where o is any object belonging to O' that a_i prefers to o_i. Observe that any node in G has out-degree $k-1$ if it belongs to O', and otherwise it has out-degree 0. In the rest of the proof, we show that there exists a sufficiently large constant k_0 such that, when $k \geq k_0$, G contains a *large* cycle whp, i.e. a cycle of length $\Theta(n)$. This will complete the proof since any cycle in G corresponds to a Pareto-improving cycle of the same length in μ.

The existence result is based on the probabilistic method, and is inspired from the techniques used in [Frieze and Karonski 2012]. The proof has two steps. In the first step, we show that G does not contain a certain subgraph, whp. (Lemma D.4) In the second step, we use this result to complete the proof. In the rest of the proof, suppose k_0 is a sufficiently large constant which will be fixed at the end.

Definition D.3. *An* independent pair (X, Y) *is a pair of subsets* $X, Y \in V(G)$ *such that* $(X \times Y) \cap E(G) = \varnothing$. *Size of an* independent pair *is* $\min\{|X|, |Y|\}$.

Lemma D.4. *Let* $\beta \in (0, 1)$ *be a constant such that* $\left(\frac{1}{1-\beta}\right)^{k_0-1} \geq \frac{e^2}{4\beta^2}$. *Then, there is no independent pair* (X, Y) *with size at least* βn*, whp.*

Proof. It suffices to show that there exists no independent pair with size exactly $\lfloor \beta n \rfloor$. To avoid notational complexity, we suppose $\beta n = \lfloor \beta n \rfloor$. This will not change any step of the proof.

Fix two subsets $X, Y \subseteq V(G)$ of size βn. Let $p_{X,Y}$ denote the chance that (X, Y) is an independent pair. First, First, we compute an upper bound on $p_{X,Y}$, and then we use a union bound over all such subsets X, Y to prove the lemma. In the rest of the proof we suppose that $n(1-\epsilon)/2 \leq N \leq n(1+\epsilon)/2$ holds for any arbitrarily small constant $\epsilon > 0$, since this condition holds whp by Lemma D.2.

Observe that

$$p_{X,Y} = \left(\frac{\binom{n-\beta n}{k_0-1}}{\binom{n}{k_0-1}} \right)^{|X|} \leq (1-\beta)^{n\beta(k_0-1)}.$$

A union bound implies that

$$\sum_{X,Y} p_{X,Y} \leq \binom{N}{n\beta}^2 (1-\beta)^{n\beta(k_0-1)}$$
$$\leq \binom{N}{n\beta}^2 (1-\beta)^{n\beta(k_0-1)}$$
$$\leq \left(\frac{e^2(1+\epsilon)^2}{(2\beta)^2}\right)^{n\beta} \left((1-\beta)^{k_0-1}\right)^{n\beta}.$$

Therefore, so long as

$$\left(\frac{1}{1-\beta}\right)^{k_0-1} \geq \frac{e^2}{(2\beta)^2},$$

we have $\sum_{X,Y} p_{X,Y} = o(1)$ (since we can take ϵ to be any arbitrarily small constant). This completes the proof.

□

Set $\alpha = 1/25$ and choose a sufficiently large constant k_0 such that $\left(\frac{1}{1-\alpha}\right)^{k_0-1} \geq \frac{e^2}{4\alpha^2}$. (Choosing $k_0 = 175$ satisfies this condition. Our simulations in Section D.4 show that $k_0 = 3$ suffices for the main claim to hold.)

Corollary D.5 (Corollary of Lemma D.4)**.** *There is no independent pair of size at least αn.*

Lemma D.6. *Suppose H is a strongly connected induced subgraph of G. Then, H contains a cycle of length at least $|V(H)| - 4\alpha n$.*

Proof. The proof uses the Depth-First Search (DFS) algorithm. (The reader may recall the definition of DFS from [West 1995] or [Wikipedia 2018a], among many other sources.)

Run DFS on the graph H starting from an arbitrary node. We keep track of two sets of vertices in the course of DFS: the set of nodes that the algorithm has not visited yet, namely S, and the set of nodes that the algorithm has finished visiting them, namely T. (Recall that when DFS visits a node v, it runs DFS recursively from all of its unvisited neighbors. The DFS at node v is finished when all the recursive calls are finished.)

When DFS starts, $|S| = |V(H)|$ and $|T| = 0$. When DFS is finished, $|S| = 0$ and $T = |V(H)|$. On the other hand, at each step (recursive call) of DFS, either $|S|$ decreases by 1 and $|T|$ is unchanged, or $|S|$ is unchanged and $|T|$ increases by 1. Therefore, at some point in the course of the algorithm, we must have $|S| = |T|$. Fix the values of S, T to be their values at that point, and let $P = V(H) \backslash (S \cup T)$. Note that P must form a directed path by the definition of DFS, and that $(S \times T) \cap E(H) = \varnothing$. The proof is almost complete using these two facts: the latter one implies that (S, T) is an independent pair. Corollary D.5 implies that the size of (S, T) is at most αn, and therefore, $|S|, |T| \leq \alpha n$. This means $|P| \geq |V(H)| - 2\alpha n$.

Suppose that $|P| \geq 2\alpha n$, otherwise the claim is trivial. Let the P_1, P_2 be subsets denoting the first and the last αn nodes of P, respectively. (P_2, P_1) cannot be an independent pair by Corollary D.5. Therefore, there must be an edge from some node in P_1 to some node in P_2. This creates the promised cycle.

□

Lemma D.7. *G contains a strongly connected subgraph that contains at least $N/3$ nodes.*

Proof. The proof is by contradiction. Suppose the claim is false. Let $C_1, \ldots, C_l$ denote the maximal strongly connected components of G that partition its vertices. Without loss of generality, suppose G contains no edge from any node in C_i to any node in C_j, for any $i < j$. (This is without loss of generality since the condensation of G is an acyclic directed graph [Wikipedia 2018b].) Let i be the smallest integer for which $\sum_{j=1}^{i} |V(C_j)| \geq N/3$. Then, we must have that $\sum_{j=i+1}^{l} |V(C_j)| \geq N/3$. Let $X = \cup_{j=i+1}^{l} C_j$ and $Y = \cup_{j=1}^{i} C_j$. Observe that (X, Y) is an independent pair of size at least $N/3$. This contradicts Corollary D.5.

□

Lemmas D.6 and D.7 together imply that, whp, G contains a cycle of length at least $N/3 - 4\alpha n \geq n(\frac{4(1-\epsilon)}{24} - \frac{4}{25})$, for any arbitrary small constant $\epsilon > 0$. This completes the proof.

D.3.3 Proof of Proposition 5.3, part (iii)

First, observe that the rank distribution under the Inefficient-RSD mechanism is stochastically dominated by the rank distribution under RSD. Therefore, the average rank under the Inefficient-RSD mechanism is higher than the average rank under RSD, which is at least $\ln n - 1$, as shown by [Knuth 1996]. On the other hand, Lemma D.2 says that, in expectation, at least $n/2$ agents are assigned to their k-th ranked object, which means the expected average rank under the Inefficient-RSD mechanism is at least $k/2$. The two latter facts imply that the expected average rank under Inefficient-RSD is at least $\max\{\ln n - 1, k/2\}$.

D.4 Simulations for Proposition 5.3

We provide some computational experiments which show that setting $k_0 = 3$ makes k_0 sufficiently large for the purpose of Proposition 5.3, part ii.

In each experiments, after fixing n, we generate 1000 random markets, each market involving n agents and n objects. We run the Inefficient-RSD mechanism with $k = 3$ on each market. In Table 1, we report the fraction of random markets in which the Inefficient-RSD assignment contains a Pareto-improving cycle. We observe that the empirical probability of the existence of a Pareto-improving cycle is essentially 1.

Table 1 also reports another statistic in its last column, which is the average length of a "large" cycle that we can find (conditional on the existence of a cycle). Since the algorithmic problem of finding the largest cycle in a graph is NP-Complete, finding the largest Pareto-improving cycle in our problem is also quite difficult (computationally). We therefore use a heuristic algorithm using the Depth-First Search (DFS) method [West 1995], which does not necessarily find the largest possible cycle. Therefore, the statistics reported in the last column of the table is a lower bound on the length of the largest possible Pareto-improving cycle. Observe that the lower bound grows roughly proportional to n.

n	Probability	Average length
2×10^2	0.998	19.09
2×10^3	1	162.22
2×10^4	1	1584.07
2×10^5	1	15779.3

Table 1: The second column reports the empirical probability that a cycle exists and the third column reports the average length of the cycle found by our heuristic.

E Proof of Proposition 5.5

Proof of Proposition 5.5, Part i.

The proof is an immediate corollary of Proposition 5.3, part i. Since, in here, the mechanism assigns the objects solely based on the idiosyncratic components, Proposition 5.3 implies that each agent asymptotically attains utility approaching $\overline{u}$ from the idiosyncratic component. Since the common-value component of the object owned by any agent is drawn independently, its distribution is equal to F. This proves the claim.

Proof of Proposition 5.5, Part ii.

For expositional simplicity, we first provide the proof assuming that F, G are the uniform distribution over the unit interval. After that we prove the claim for the general case. The proof of the general case has only minor differences with this proof.

Let μ denote the Inefficient-RSD assignment. Let $a_1, \ldots, a_n$ denote the agents in the order they choose objects. Recall the notation $\pi'(a)$ from the definition of Inefficient-RSD mechanism: it is a list of the objects, ordered in decreasing order, according to the idiosyncratic utility components of agent a. Let r'_a denote the rank of the object assigned to agent a on $\pi'(a)$. Let $l = k/3$ and define the interval $I = (2l, 3l]$.

Claim E.1. *For any agent $a \in A$, $r'_a \in I$ holds, whp.*

Proof. Let p_i denote the probability that $r'_i \notin I$. First, we provide an upper bound on p_i. Observe that for $i \leq l$, $p_i = 0$. For $i \geq l$,

$$p_{i+1} = \frac{i}{n} \cdot \frac{i-1}{n-1} \cdots \frac{i-l}{n-l} \leq \left(\frac{i}{n}\right)^l \leq e^{-\frac{lj}{n}}, \tag{E.1}$$

where $j = n - i$. Observe that, for any j such that $j = \omega(n/l)$, we have $p_{i+1} = o(1)$. Choose such arbitrary j, namely j^*. Therefore, the above inequality implies that for all $i \leq n - j^*$, we have $p_{i+1} = o(1)$. This concludes the proof since the order of agents is chosen uniformly at random. $\square$

In the rest of the proof, we will show that there exists $n - o(n)$ pairwise-disjoint Pareto-improving pairs in the Inefficient-RSD assignment. This will also imply that each agent is involved in at least one Pareto-improving pair, whp. Therefore, this will prove part ii of the proposition. To this end, we need a few definitions.

Let $N = |\{a : r'_a \in I\}|$. Claim E.1 implies that $N = n - o(n)$. Let the set $B = \{b_1, \ldots, b_N\}$ denote the set of agents $\{a : r'_a \in I\}$. Also, let $o_i = \mu(b_i)$, for all $b_i \in B$. Denote the common-value component corresponding to objects o_i by v_i. Without loss of generality suppose that $v_1 \leq \ldots \leq v_N$. Furthermore, by relabeling the agents we can assume that $b_i = i$ for all $i \in [N]$.

Let x be a positive number which we will fix later. Partition the unit interval to segments of length x: the partition would be

$$(0, x], (x, 2x], (2x, 3x] \dots,$$

where x is chosen so that $x = \omega(\frac{\log n}{n})$ and $x = o(\frac{k}{n})$. (The last element of the partition may be an interval with length smaller than x, in which case we will safely ignore the last element in the analysis below.)

For any integer $i \geq 0$, let A_i be the set of indices j such that $v_j \in (ix, ix + x]$. (Recall that, since the object with common-value v_i is owned by the agent b_i, and since we supposed $b_i = i$ for all $i \in [N]$, we can also interpret A_i as a set of agents.)

Claim E.2. *For any constant $\epsilon > 0$ and any $i \geq 0$, wvhp*

$$|A_i| \in [(1 - \epsilon)xN, (1 + \epsilon)xN]$$

Proof. Recall that the common-value components of the objects are drawn iid from the uniform distribution over the unit interval. Therefore, the expected number of objects in B whose common-value component falls into a subinterval of the unit interval with length α is αN. Since $x = \omega(\frac{\log n}{n})$, the length of the subinterval defining A_i (i.e. the interval $(ix, ix+x]$) would be $\omega(\frac{\log n}{n})$, and therefore the expected size of A_i would be $\omega(\log n)$. The rest of the proof follows from a standard application of Chernoff concentration bounds. □

Next, we show that the set of agents A_i contains at least $|A_i|/2 - o(|A_i|)$ pairwise-disjoint Pareto-improving pairs, whp. This would imply that the Inefficient-RSD assignment contains at least $n/2 - o(n)$ pairwise-disjoint Pareto-improving pairs, in expectation, which is the claim. The next lemma completes the proof.

Lemma E.3. *For any i, the set of agents A_i contains at least $|A_i|/2 - o(|A_i|)$ pairwise-disjoint Pareto-improving pairs, whp.*

Proof. First, we construct an undirected graph G as follows. Let $V(G) = A_i$, and connect a node z to node z' iff the rank of z on $\pi'(z')$ is at most l and the rank of z' on $\pi(z)$ is at most l. Let t denote the number of nodes in G. The proof contains two steps: we will show that (i) any edge in G corresponds to a Pareto-improving pair, wvhp, and (ii) for any constant $\beta < 1$, G contains a matching that covers at least $\beta|V(G)|$ nodes, whp. These two steps, together with a union bound, conclude the lemma.

Step (i) We can assume that the positions $1, \ldots, l$ on the preference lists of all agents in B are unfilled (by the Principle of Deferred Decisions). Let $E_{b,b'}$ denote the event that $\mu(b)$ has rank l or better on $\pi'(b')$ and $\mu(b')$ has rank l or better on $\pi'(b)$. We will show that that if $E_{b,b'}$ holds, then b, b' must form a Pareto-improving pair, wvhp.

Claim E.4. *For any two agents $b, b' \in A_i$, conditioned on $E_{b,b'}$, agents b, b' form a Pareto-improving pair, wvhp.*

Proof. Recall that for any agent a and object o, the utility of a from o is

$$u_a(o) = v_o + v_o^a.$$

Observe that $|v_{\mu(b)} - v_{\mu(b')}| \leq x$, by the definition of A_i. To prove the claim it suffices to show that the two inequalities below hold wvhp.

$$v^b_{\mu(b)} + x < v^b_{\mu(b')}, \tag{E.2}$$
$$v^{b'}_{\mu(b')} + x < v^{b'}_{\mu(b)}. \tag{E.3}$$

To prove (E.2), we will show that for any arbitrary small constant $\epsilon > 0$

$$v^b_{\mu(b)} < 1 - 2l(1-\epsilon)/n,$$
$$v^b_{\mu(b')} > 1 - l(1+\epsilon)/n,$$

hold wvhp. ((E.3) is proved similarly.) The two equations above hold whvp because of Chernoff concentration bounds, as explained next. Fix an arbitrary small constant $\epsilon > 0$. Since the idiosyncratic components are drawn iid, then, for any agent there are at least l of the objects that give her and idiosyncratic utility component at least $1 - \frac{l(1+\epsilon)}{n}$, wvhp. This proves the latter inequality. Similarly, for any agent there are at most $2l$ of the objects that give her an idiosyncratic utility component at least $1 - \frac{2l(1-\epsilon)}{n}$, wvhp. This proves the former inequality.

□

Step(ii) In this step, we will show that for any constant $\beta < 1$, G contains a matching that covers at least $\beta|V(G)|$ nodes, whp. First, we need a definition. An independent set in G is a set of nodes which are pairwise non-adjacent. To prove the lemma, we show that, whp, G does not contain an independent set of size $\delta|V(G)|$, for any arbitrary small constant

$\delta > 0$. This completes this step because if the maximum size of a matching in G is M, then G must contain an independent of size $|V(G)| - M$.

The proof uses the Principle of Deferred Decisions. While running the Inefficient-RSD mechanism, suppose that the preference lists of agents are filled gradually in the course of the algorithm: whenever an agent goes to the next position on her list, a random draw from the remaining objects will fill out that position.

For any two distinct $b, b' \in B$, we let $p_{b,b'} = \mathbb{P}\left[E_{b,b'}\right]$. By the Principle of Deferred Decisions, the chance that $\mu(b')$ has rank l or better on $\pi'(b)$ list is at least l/n. Similarly, the chance that $\mu(b)$ has rank l or better on agent $\pi'(b')$ is at least l/n. Therefore, $p_{b,b'} \geq l^2/n^2$. It is straight-forward to verify that, for any subset $B' \subseteq B$

$$\begin{aligned}\mathbb{P}\left[\wedge_{b,b' \in B'} \overline{E_{b,b'}}\right] &\leq \prod_{b,b' \in B} (1 - p_{b,b'}) \\ &\leq (1 - l^2/n^2)^{\binom{|B'|}{2}}.\end{aligned}$$

When $|B'| > \delta|B|$ for a constant $\delta > 0$, the above inequality implies that

$$\begin{aligned}\mathbb{P}\left[\wedge_{b,b' \in B'} \overline{E_{b,b'}}\right] &\leq (1 - l^2/n^2)^{\frac{\delta^2 |B|^2}{2}} \\ &\leq e^{-\left(\frac{\delta l |B|}{\sqrt{2}n}\right)^2}\end{aligned}$$

A union bound then implies that G contains an independent set of size at least $\delta|B|$ with probability at most

$$2^{|B|} e^{-\left(\frac{\delta l |B|}{\sqrt{2}n}\right)^2} = o(1),$$

where the above equation holds because $|B| = n - o(n)$ and $l = \omega(n^{1/2})$. This completes the proof of the lemma.

□

Dismissing the uniformity assumption in part ii

We now dismiss the uniformity assumption made on F, G which was used in part ii of Proposition 5.5. For the proof, without loss of generality, we can suppose that the support of the distributions is the unit interval, i.e. $[\underline{u}, \overline{u}] = [0, 1]$. This is just a normalization that simplifies notation. By the assumption that F, G have finite-valued PDFs, there exists a constant $\Delta > 0$ that bounds the PDFs of F, G from above.

The proof follows the same steps as the proof for the case of uniform distributions. In the course of this proof, it is assumed that all of the parameters have the same definition as in the proof that we presented under the uniformity assumption, unless explicitly stated otherwise.

We start by partitioning the interval $[1, k]$ to 3 intervals by choosing two numbers l_1, l_2. The partition would be $[1, l_1]$, (l_1, l_2), $[l_2, k]$. In the previous proof, we chose $l_1 = k/3$ and $l_2 = \frac{2k}{3}$. In this proof, we choose $l_2 = \frac{2k}{3}$, but we choose a different l_1, as follows. Since the PDFs of F, G have support $[\underline{u}, \overline{u}]$ then, by the definition of the support of a function, there exists a positive constant $\delta > 0$ which bounds the PDFs from below over $[\underline{u}, \overline{u}]$. We set $l_1 = \frac{\delta k}{3\Delta}$. As in previous proof, we denote the interval $[l_2, k]$ by I. As in the previous proof, we also choose a parameter x such that $x = \omega(\frac{\log n}{n})$ and $x = o(\frac{k}{n})$

For this proof to work, we have chosen the parameters x, l_1, l_2 so that the crucial properties, which make the previous proof work under the uniformity assumption, hold here as well. In particular, the proofs for Claim E.1, Claim E.2 and Lemma E.3, which derive the previous proof, remain the same, *mutatis mutandis*. For completeness, we present the counterpart statements and their proofs.

Claim E.5 (Counterpart to Claim E.1)**.** *For any agent $a \in A$, $r'_a \in I$ holds, whp.*

Proof. The proof remains the same as the proof of Claim E.1, with a minor difference that the parameter l in that proof is replaced with the parameter $k - l_2$ for the proof in here. □

Claim E.6 (Counterpart to Claim E.2)**.** $|A_i| \in [x\delta(1-\epsilon)N, x\Delta(1-\epsilon)N]$.

Proof. Recall that the common-value components of the objects are drawn iid from the PDF F. Therefore, the expected number of objects in B whose common-value component falls into a subinterval of the unit interval with length α is at least $\alpha\delta N$ and at most $\alpha\Delta N$. Since $x = \omega(\frac{\log n}{n})$, the length of the subinterval defining A_i (i.e. the interval $(ix, ix + x]$) would be $\omega(\frac{\log n}{n})$, and therefore the expected size of A_i would be $\omega(\log n)$. The rest of the proof follows from a standard application of Chernoff concentration bounds.

□

Lemma E.7 (Counterpart to Lemma E.3)**.** *For any i, the set of agents A_i contains at least $|A_i|/2 - o(|A_i|)$ pairwise-disjoint Pareto-improving pairs, whp.*

Proof. The proof is similar to the proof of Lemma E.3. We present the full proof here for completeness.

First, we construct an undirected graph G as follows. Let $V(G) = A_i$, and connect a node z to node z' iff the rank of z on $\pi'(z')$ is at most l_1 and the rank of z' on $\pi(z)$ is at most

l_1. Let t denote the number of nodes in G. The proof contains two steps: we will show that (i) any edge in G corresponds to a Pareto-improving pair, wvhp, and (ii) for any constant $\beta < 1$, G contains a matching that covers at least $\beta|V(G)|$ nodes, whp. These two steps, together with a union bound, conclude the lemma.

Step (i) We can assume that the positions $1, \ldots, l_1$ on the preference lists of all agents in B are unfilled (by the Principle of Deferred Decisions). Let $E_{b,b'}$ denote the event that $\mu(b)$ has rank l_1 or better on $\pi'(b')$ and $\mu(b')$ has rank l_1 or better on $\pi'(b)$. We will show that that if $E_{b,b'}$ holds, then b, b' must form a Pareto-improving pair, wvhp.

Claim E.8. *For any two agents $b, b' \in A_i$, conditioned on $E_{b,b'}$, agents b, b' form a Pareto-improving pair, wvhp.*

Proof. Recall that for any agent a and object o, the utility of a from o is

$$u_a(o) = v_o + v_o^a.$$

Observe that $|v_{\mu(b)} - v_{\mu(b')}| \leq x$, by the definition of A_i. To prove the claim we will show that

$$v_{\mu(b)}^b + x < v_{\mu(b')}^b, \tag{E.4}$$
$$v_{\mu(b')}^{b'} + x < v_{\mu(b)}^{b'} \tag{E.5}$$

hold wvhp.

We will prove this for (E.4); the proof for (E.5) is similar. To this end, we first will show that

$$v_{\mu(b)}^b < 1 - \frac{l_2(1-\epsilon)}{n\Delta}, \tag{E.6}$$
$$v_{\mu(b')}^b > 1 - \frac{l_1(1+\epsilon)}{n\delta} \tag{E.7}$$

hold wvhp. If the above two inequalities are proved, then together with $l_1 = \frac{\delta l_2}{2\Delta}$ they will imply that

$$v_{\mu(b')}^b - v_{\mu(b)}^b > \frac{l_2(1-\epsilon)}{n\Delta} - \frac{l_2(1+\epsilon)}{2n\Delta} = \frac{l_2(1-2\epsilon)}{2n\Delta}.$$

Then, this inequality together with $l_2 = 2k/3$ and $x = o(k/n)$ would imply that $\frac{l_2(1-2\epsilon)}{2n\Delta} > x$

holds for sufficiently large n, which would imply that (E.4) holds. Hence, to prove that (E.4) holds wvhp, it suffices to show that (E.6) and (E.7) hold, wvhp. The two latter bounds hold because of Chernoff concentration bounds, as explained next.

Observe that the following event holds wvhp for any arbitrary small $\epsilon > 0$: for any agent there are at least l_1 of the objects that provide her an idiosyncratic utility component of at least $1 - \frac{l_1(1+\epsilon)}{n\delta}$. This follows from a standard application of Chernoff bounds, which hold because the idiosyncratic components are drawn iid from the PDF G. This proves that (E.6) holds wvhp. We prove that (E.6) holds wvhp similarly. By Chernoff bounds, the following event holds wvhp for any arbitrary small constant $\epsilon > 0$: for any agent there are at most l_2 of the objects that provide her an idiosyncratic utility component of at least $1 - \frac{l_2(1-\epsilon)}{n\Delta}$, wvhp. Hence, (E.7) holds wvhp.

□

Step(ii) In this step, we will show that for any constant $\beta < 1$, G contains a matching that covers at least $\beta|V(G)|$ nodes, whp. First, we need a definition. An independent set in G is a set of nodes which are pairwise non-adjacent. To prove the lemma, we show that, whp, G does not contain an independent set of size $\delta|V(G)|$, for any arbitrary small constant $\delta > 0$. This completes this step because if the maximum size of a matching in G is M, then G must contain an independent of size $|V(G)| - M$.

The proof uses the Principle of Deferred Decisions. While running the Inefficient-RSD mechanism, suppose that the preference lists of agents are filled gradually in the course of the algorithm: whenever an agent goes to the next position on her list, a random draw from the remaining objects will fill out that position.

For any two distinct $b, b' \in B$, we let $p_{b,b'} = \mathbb{P}[E_{b,b'}]$. By the Principle of Deferred Decisions, the chance that $\mu(b')$ has rank l_1 or better on $\pi'(b)$ list is at least l_1/n. Similarly, the chance that $\mu(b)$ has rank l_1 or better on agent $\pi'(b')$ is at least l_1/n. Therefore, $p_{b,b'} \geq l_1^2/n^2$. It is straight-forward to verify that, for any subset $B' \subseteq B$

$$\begin{aligned}\mathbb{P}\left[\wedge_{b,b' \in B'} \overline{E_{b,b'}}\right] &\leq \prod_{b,b' \in B} (1 - p_{b,b'}) \\ &\leq (1 - l_1^2/n^2)^{\binom{|B'|}{2}}.\end{aligned}$$

When $|B'| > \delta|B|$ for a constant $\delta > 0$, the above inequality implies that

$$\begin{aligned}\mathbb{P}\left[\wedge_{b,b'\in B'}\overline{E_{b,b'}}\right] &\leq (1 - l_1^2/n^2)^{\frac{\delta^2|B|^2}{2}} \\ &\leq e^{-\left(\frac{\delta l_1|B|}{\sqrt{2}n}\right)^2}\end{aligned}$$

A union bound then implies that G contains an independent set of size at least $\delta|B|$ with probability at most

$$2^{|B|}e^{-\left(\frac{\delta l_1|B|}{\sqrt{2}n}\right)^2} = o(1),$$

where the above equation holds because $|B| = n - o(n)$ and $l_1 = \omega(n^{1/2})$. This completes the proof of the lemma. □

F Proof of Proposition 6.1

In the proof we use the notions of the Deferred Acceptance (DA) algorithm, the man-proposing DA, and the man-optimal assignment. The reader may recall the definitions from [Roth and Sotomayor 1990], among other possible sources.

The proof directly follows from the next lemma together with a result of [Pittel 1992] that shows in the man-optimal matching, all men are matched to women that are ranked $\ln^2 n$ or better on the list of their match, whp.

Lemma F.1. *Consider a marriage market consisting n men and n women. Let $r(n) : \mathbb{N} \to \mathbb{R}_+$ be any positive function such that $r(n) = o(\frac{n}{\ln n})$. Then, there exists a stable matching in which each women, whp, is assigned to a man ranked $r(n)$ or worse on her list.*

Proof. The proof adapts the proof of of Lemma C.1. from [Ashlagi and Nikzad 2018]. We prove the lemma in two steps. In the Step (i), we show that under the man-proposing differed acceptance, each woman receives at most $(1 + \epsilon)\ln n$ proposals whp, where $\epsilon > 0$ is an arbitrary small constant independent of n. In Step (ii) we prove the claim of the lemma using the result of Step (i).

Step (i) The proof idea is defining another stochastic process that we denote by $\mathcal{B}$. Process $\mathcal{B}$ is defined by a sequence of binary random variables $X_1, \ldots, X_k$, where $k = (1 + \delta)n \ln n$ for some arbitrary small constant $\delta > 0$. Each random variable in this sequence takes the value 1 with probability $\frac{1}{n-\ln^2 n}$, and 0 otherwise. For convenience, we also refer to these random variables by *coins*, and the process that determines the value of a random variable by *coin-flip*.

Fix a woman w. Define $X = \sum_{i=1}^{k} X_i$. The goal is to show that X is a good upper bound on the number of proposals that are received by w. The high-level idea is based on two facts: First, the total number of proposals made by all men is stochastically dominated by the solution to the coupon-collector problem, and so, whp is at most k. Second, by [Pittel 1992], there is no man who makes more than $\ln^2 n$ proposals, whp. So, each proposal is made to w with probability at most $\frac{1}{n-\ln^2 n}$. Consequently, the number of proposals made to w cannot be more than $\frac{k(1+\delta')}{n-\ln^2 n}$ whp, for any constant $\delta' > 0$. (The latter fact is a direct consequence of the Chernoff bound which is applicable since the coin flips are independent).

To formalize the above argument, we couple the stochastic process governing DA with a new random process, $\mathcal{B}$, which is a simple coin-flipping process: it flips k coins independently, all with success probabilities $\frac{1}{n-\ln^2 n}$. The coupled process, $(\mathrm{DA}, \mathcal{B})$, has two components, one for each of the original random processes. For each of the components, the marginal distribution induced on its sample paths is identical to the distribution of the sample paths in the original process, but there is no restriction on the joint distribution of the sample paths in the coupled process. Next, we define a simple coupling in which in almost all sample paths (i.e. whp), the number of successful coin flips is an upper bound on the number of proposals made to w. Whenever a man m wants to make a proposal during DA, process $\mathcal{B}$ flips the next coin. Then:

1. If m has made a proposal to w before, ignore the coin flip, and let m make a proposal to a woman picked uniformly at random from the set of women to whom he has not proposed yet.

2. If m has not made a proposal to w before, then

   (a) Suppose m has made $d \leq \ln^2 n$ proposals so far. (Otherwise, ignore this sample path)

      i. If the coin flip was successful: with probability $\frac{n-\ln^2 n}{n-d}$ let m make a proposal to w, and otherwise (with probability $1 - \frac{n-\ln^2 n}{n-d}$) let m make a proposal to a woman picked uniformly at random from the set of women to whom he has not proposed yet.

      ii. If the coin flip was not successful: let m make a proposal to a woman picked uniformly at random from the set of women to whom he has not proposed yet, excluding w.

It is straight-forward to verify that this defines a valid coupling of DA, $\mathcal{B}$. Note that the total number of successful coin flips in $\mathcal{B}$ is an upper bound on the total number of proposals made to w in the coupled DA process, in almost all sample paths (i.e. whp). Therefore, as we explain next, we can apply the argument mentioned in the beginning of the proof to conclude the lemma.

[Pittel 1992] shows that, whp, no man makes more than $\ln^2 n$ proposals. Therefore, whp, a sample path is not ignored in line (2-a); that is, only a vanishing fraction of sample paths are ignored. On the other hand, we mentioned earlier that a standard application of Chernoff bounds implies that X is at most $\frac{k(1+\delta')}{n-\ln^2 n}$ whp, for any constant $\delta' > 0$. A union bound then implies that, whp, woman w receives at most

$$\frac{k(1+\delta')}{n-\ln^2 n} = \frac{(1+\delta)(1+\delta')n\ln n}{n-\ln^2 n}$$

proposals. Setting ϵ to be a constant larger than $\delta + \delta' + \delta\delta'$ completes the first step.

Step (ii) Fix a woman w. For any $k \leq (1+\epsilon)\ln n$, conditioned on w receiving k proposals, the chance that she is assigned to a man ranked $r(n)$ or worse on her list is

$$\left(1 - \frac{r(n)}{n}\right)^k \geq 1 - \frac{k \cdot r(n)}{n} = 1 - o(1).$$

Also, note that w receives no more than $(1+\epsilon)\ln n$ proposals, whp. Therefore, a union bound implies that w is assigned to a man ranked $r(n)$ or worse on her list, whp.

□