

Rank-optimal assignments in uniform markets

AFSHIN NIKZAD

Department of Economics, USC

We prove that in a market where agents rank objects independently and uniformly at random, there exists an assignment of objects to agents with a constant average rank (i.e., an average rank independent of the market size). The proof builds on techniques from random graph theory and the FKG inequality (Fortuin et al. (1971)). When the agents' rankings are their private information, no Dominant Strategy Incentive Compatible mechanism can implement the assignment with the smallest average rank; however, we show that there exists a Bayesian Incentive Compatible mechanism that does so. Together with the fact that the average rank under the Random Serial Dictatorship (RSD) mechanism grows infinitely large with the market size, our findings indicate that the average rank under RSD can take a heavy toll compared to the first-best, and highlight the possibility of using other assignment methods in scenarios where average rank is a relevant objective.

KEYWORDS. Matching, average rank, random serial dictatorship, FKG inequality.

JEL CLASSIFICATION. C78, D82.

1. INTRODUCTION

In assignment problems involving ordinal preferences, *rank* is a commonly considered notion for comparing different assignments. School districts consider the number of students assigned to their first choice, second choice, and so on, to report their assignment results or compare different assignments (Abdulkadiroglu et al. (2005), De Haan et al. (2015), Featherstone (2020)). Similarly, some institutions such as Teach For America choose assignments explicitly based on such rank distributions (Featherstone (2020)).¹

The notion of rank is also used in evaluating assignments in the National Residency Matching Program (NRMP (2020)).² In evaluating course assignment mechanisms as well some educational institutions focus on rank; e.g., Harvard Business School (HBS) emphasizes the average over the ranks of the ten courses that a student allocates on her reported rank-order list. Budish and Cantillon (2012) consider average rank, as well

Afshin Nikzad: afshin.nikzad@usc.edu

¹Teach For America assigns college graduates to teach in low-performing public and charter schools in the United States. (The online appendix can be found in the full version of the paper available at https://papers.ssrn.com/abstract_id=3930653.)

²This program places U.S. medical school students into residency training programs through a central mechanism that elicits rank-order lists from them. There, *match success* for an applicant is defined as being matched to the specialty of her first-ranked program (NRMP (2020)).

as other rank-based notions, to compare different course assignment mechanisms with the one used at HBS.

The literature on matching markets has thus considered rank as a natural summary statistic for quality of a match. A closely related notion that we will focus on is *average rank*, loosely speaking defined as the average, over all agents, of the rank of the object assigned to each agent on her list.³ Average rank has been used, e.g., to compare stable assignments in two-sided markets where every agent ranks the agents on the other side independently and uniformly at random (Pittel (1989), Ashlagi et al. (2017)) or to evaluate core allocations in one-sided markets where agents rank objects uniformly at random (Knuth (1996)).

We similarly consider markets where agents on one side rank objects on the other side independently and uniformly at random. Our main finding is that, in such markets, there exists an assignment of objects to agents with a constant average rank. That is, there exists a constant R , independent of the market size, such that the object assigned to each agent has an average position of at most R on her list (Theorem 3.1). In addition, there exists a Bayesian Incentive Compatible mechanism that implements such an assignment (Section 5).

The proof of our main finding builds on techniques from random graph theory and probabilistic analysis. We use 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.⁴ This inequality is used to keep track of the correlations between the objects' ranks on an agent's preference list.

Knuth (1996) is among the first to analyze the average rank in the market that we consider. He computes the exact value of the average rank under the *Random Serial Dictatorship* mechanism (RSD), in which agents are ordered randomly and choose objects one by one in that order (Abdulkadirouglu and Sönmez (1998)). His result implies that, in markets with an equal number of objects and agents, the average rank under RSD grows logarithmically in the market size.

Our result, together with that of Knuth (1996), indicates that the average rank can take a heavy toll under the RSD mechanism compared to the first-best. We further show that the gap between the first-best average rank and the average rank under RSD can persist when objects have capacities larger than one.^{5,6} These findings highlight the possibility of using assignment methods other than RSD in scenarios where the average rank is a relevant objective. This is also aligned with the findings of Budish and Cantillon (2012) who, using field data from HBS, show that switching from the “draft” mechanism

³The formal definition that we use is in Section 3.

⁴Section 3 recalls this inequality.

⁵That is when a bounded number of agents may be matched to the same object.

⁶More precisely, we prove that the first-best average rank does not grow infinitely large as the number of agents approaches infinity, whereas the average rank under RSD does, as long as objects' capacities grow sublogarithmically in the number of agents.

used at HBS to RSD increases the average rank of the ten courses in a student's schedule by a "large magnitude" (from 7.99 to 8.74).⁷

In Section 5, we consider the scenario where the agents' preference orders over objects are their private information, and we ask whether there exists a Dominant Strategy Incentive Compatible (DSIC) mechanism that implements the assignment with the first-best average rank. While the answer to this question turns out to be negative, we show that the answer is positive if the notion of DSIC is relaxed to Bayesian Incentive Compatibility (BIC). Together with Theorem 3.1, this implies the existence of BIC mechanisms in our main setup that implement assignments with a constant average rank of at most R .

Finally, we highlight the work of Che and Tercieux (2018) on the payoff equivalence of Pareto efficient assignments⁸ and its connection to our findings. They show that when the agents' values for objects are independently and identically drawn (iid) from a distribution with bounded support, all Pareto efficient assignments achieve the utilitarian upper bound in the large market limit, and thus are payoff equivalent.^{9,10} At first glance, this result may seem contradictory to the large gap between the average ranks of the RSD and the first-best assignments, as both assignments are Pareto efficient. However, the payoff equivalence result does not (need to) hold in our setup because the conditions required for it—the bounded support and the iid assumptions—do not hold here.¹¹

2. RELATED LITERATURE

In the literature on matching markets, rank has been used as a standard notion of evaluating assignments. Knuth (1996) is one of the earliest to consider the notion of average rank in one-sided markets that are drawn uniformly at random. He analyzes the assignment model of Shapley and Scarf (1974) and provides an exact analysis for the average rank and its variance in that model. He also shows that these statistics are respectively equal to the average rank and its variance under the RSD mechanism.

The existence of assignments with a constant average rank in uniform markets has been shown in the literature when there is a surplus of objects that grows linearly in the market size. For instance, this is a direct corollary of the work of Frieze and Melsted (2009) and Fountoulakis and Panagiotou (2012) on cuckoo hashing. Furthermore, it has been noted in the literature that in markets with a linear surplus of objects, the average rank of the assignment generated by RSD is constant.¹² Our results complement these

⁷Remarkably, while the draft mechanism is not strategy-proof, they compute the average rank according to the students' *true* preferences, taken from a survey data conducted by the HBS administration.

⁸An assignment is Pareto efficient if there is no other assignment in which every agent is weakly better off with at least one agent being strictly better off.

⁹Their result is in fact more general: they show that the payoff equivalence result holds even in the presence of common-value components for objects, as long as the idiosyncratic components are drawn independently and uniformly at random from distributions with bounded support.

¹⁰Lee and Yariv (2018) establish a similar result in two-sided markets for stable assignments.

¹¹The bounded support assumption does not hold because the rank of an object does not have a bounded range as the market grows large. The iid assumption does not hold because, if an object is ranked k th on an agent's list, no other object can be ranked k th.

¹²For example, see footnote 13 in Ashlagi et al. (2017)

findings by showing that the existence of assignments with a constant average rank does not depend on the level of market imbalance.

Featherstone (2020) defines an assignment to be rank efficient if its rank distribution cannot be stochastically dominated. He considers rank efficiency as a common objective of policy makers and shows that it can be attained by looking for local improvements that increase a natural objective. He also shows that no rank-efficient mechanism is strategy-proof in general.¹³

The literature on *online assignment* problems has also considered rank as a notion of efficiency. Bruss (2000) develops optimal policies for the *hiring problem* (also known as the *secretary problem*) in which an employer needs to hire the top-ranked employee from a finite number of applicants, arriving in random order one by one. Chun and Sumichrast (2006) extend this setting to a matching setting for sequential assignment of jobs to machines.

In two-sided markets as well, the notions of average rank and rank distribution have been considered to evaluate assignment efficiency. Abdulkadirouglu et al. (2005) use data from the NYC high school match to empirically compare different tie-breaking rules for the Deferred Acceptance algorithm according to their rank distributions. Ashlagi et al. (2019) and Ashlagi and Nikzad (2020) compare the effect of tie-breaking rules on the students' rank distributions. Pittel (1992) and Ashlagi et al. (2017) analyze the average rank of agents in stable assignments in the context of marriage markets and characterize the stark effect of market imbalance on the average rank.¹⁴

Finally, we discuss relevant results in random graph theory and random matrix theory. From a technical perspective, we build on techniques from random graph theory and probabilistic analysis. 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.

We also highlight the Parisi conjecture in random matrix theory, proved by Linusson and Wästlund (2004) and Nair et al. (2005) independently. The Parisi conjecture states that in an $n \times n$ matrix whose elements are iid from the exponential distribution with mean 1, there exists a set of n elements, no two in the same row or column, that asymptotically sum up to $\pi^2/6$. While this is reminiscent of Theorem 3.1 in some aspects, there is no formal connection to the extent of our knowledge.¹⁵ Theorem 3.1 is concerned with the average rank, whereas the Parisi conjecture is concerned with the sum of iid real numbers. In particular, in the setting of the Parisi conjecture, the elements of the matrix are independent, whereas in our setting, the elements in a row (ranks of the objects) are correlated. For handling these correlations, our proof takes a completely different approach using the FKG inequality (Fortuin et al. (1971)). Our approach also provides an upper bound—in the sense of first-order stochastic dominance—on the distribution of the rank of the object assigned to each agent (Remark 4.11).

¹³We complement this finding in Section 5 by considering all market sizes and object capacities.

¹⁴Their results on the effect of market imbalance also extend to the number of stable partners.

¹⁵In Section ii of the Online Appendix, we discuss two natural but unsuccessful proof approaches based on the Parisi conjecture and the complexities that these approaches encounter. Making a formal connection to the Parisi conjecture remains an interesting theoretical direction.

3. SETUP

A *matching 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, we assume that $n = cm$, unless explicitly stated otherwise. (We will see that the main theorem also holds without the latter assumption.) An *assignment* is a function $\mu : A \rightarrow O$ that assigns each agent to an object without assigning more than c agents to any object.

Each agent has a strict¹⁶ preference order over the objects, which is also called the *preference list* of the agent. 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 it is denoted by $r_a(o)$. The *average rank* of an assignment μ is defined as

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

The set of preference lists of all agents is called a *preference profile*. When A, O are known from the context, we denote the set of all preference profiles by Π and a typical member of it by π . For a preference profile π , let the *rank-optimal* assignment, $r^*(\pi)$, be the assignment with the minimum average rank.

A *uniform market* is a matching market in which the preference order of each agent is drawn independently and uniformly at random from the set of all strict orderings of objects.

We next state some preliminaries required for the main findings.

Asymptotic notions We say a sequence of events $E_1, E_2, \dots$ occurs *with high probability* (whp) if $\mathbb{P}[E_k]$ approaches 1 as k approaches infinity.

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, $V(\mathcal{L})$ be the set of its elements, and $\mu : V(\mathcal{L}) \rightarrow \mathbb{R}_+$ be a nonnegative function 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 $f, g : V(\mathcal{L}) \rightarrow \mathbb{R}_+$ are respectively increasing and decreasing functions on the lattice $\mathcal{L}$, it holds that

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

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

¹⁶Strictness simplifies notation, but it is not a necessary assumption.

3.1 Average rank of rank-optimal assignments

We next show that, in a uniform market, the expected average rank of the rank-optimal assignment is bounded from above by a constant independent of the market size.

THEOREM 3.1. *There exists a constant R independent of c , n such that the expected average rank of the rank-optimal assignment in a uniform market is at most R .*

We discuss this result and the proof approach in Section 3.2. There, using numerical simulations, we illustrate that the constant R is small (less than 2) for any capacity parameter c . The full proof of the theorem is presented in Section 4.

From this theorem, it also follows that rank-optimal assignments have a constant expected average rank in uniform markets where there is a surplus or shortage of objects. (In case of a shortage of objects, the average rank is defined over the agents who are assigned an object.) This is formally proved in the Online Appendix, Section i.

3.2 Discussion of the result and the proof approach

To prove Theorem 3.1, we show that the expected average rank of the rank-optimal assignment is smaller in a uniform market with n agents and n/c objects each with capacity c than in a uniform market with n agents and n objects each with capacity 1 (Lemma A.1). Thus, by this lemma, 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. Our proof shows that $R < 7\frac{3}{4}$ for sufficiently large n , whereas our simulations demonstrate that $R < 2$ (Figure 1).

A natural proof approach for the case of unit capacities would be constructing a bipartite graph with agents on one side and objects on the other side, where every agent is connected to her k top-ranked objects, for some constant $k \geq 1$. The existence of a perfect matching¹⁷ in such a graph would prove the claim. It turns out, however, for this




FIGURE 1. For each $n \in \{10, 20, 30, \dots, 400\}$, we report the average over 1000 independent uniform markets with unit capacities. The largest reported average is below 1.83. In addition, for $n = 1000$ and $n = 10,000$ we take the average over 10 independent draws due to computational limitations; the simulations report average ranks 1.831 and 1.827, respectively.

¹⁷We recall that a matching in a graph is a set of edges with no two edges having a node in common. A perfect matching is one that covers every node of the graph.

graph to contain a perfect matching, k must grow logarithmically with the number of agents (Frieze and Karonski (2012)). This construction thus can only prove the existence of an assignment with an average rank of the order of $\log n$.^{18,19}

Our proof starts with a different graph, which is inspired from a result of Walkup (1980). He 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. Inspired by this result, we construct a bipartite graph G with partitions A and O . In this graph, connect each node $a \in A$ to the three top-ranked objects on the preference list of a . We call these objects the *right-neighbors* of a . Also, connect every object o to every agent who ranks o worse than 3 and no worse than r_o , where r_o is defined as follows. For object $o \in O$, let r_o be the smallest integer larger than 3 such that there exists at least 3 agents for whom o has a rank (strictly) worse than 3 and no worse than r_o . If there is no such integer, then define $r_o = n$. The agents who rank o worse than 3 and no worse than r_o are called the *left-neighbors* of o .²⁰

Intuitively, the graph that we construct G has similarities with the graph considered in Walkup (1980), in the sense that each agent is adjacent to three distinct right-neighbors chosen *independently* and *uniformly at random*. Unlike in Walkup (1980), however, the left-neighbors of objects are not chosen independently: the objects' left-neighbors may not be independent from each other or from the agents' right-neighbors, as we elaborate later.

In what follows, we describe how G is used to prove the theorem. The proof has two steps. In step A, we show that G has a perfect matching with high probability, and in step B, we show that the perfect matching has a constant average rank.

Step A This step shows that G has a perfect matching with probability at least $1 - n^{-3}$ for all $n \geq 50$. As in Walkup (1980), the proof approach is based on the König theorem. By this theorem, the size of the maximum matching plus the size of the maximum independent set in a bipartite graph equals the total number of nodes.²¹ Therefore, to prove that G has a matching of size n , it suffices to show that it has no independent set of size $n + 1$. We prove that the probability that G contains an independent set of size $n + 1$ is at most n^{-3} . This is done by bounding the probability that a fixed subset of nodes of size $n + 1$ forms an independent set, and then taking a union bound over all subsets of nodes of size $n + 1$.

¹⁸Remarkably, the average rank generated by RSD is also of the same order.

¹⁹When there is a surplus of objects that grows linearly in n , Fountoulakis and Panagiotou (2012) show that there exists a constant k such that the constructed graph contains a perfect matching with high probability.

²⁰In the bipartite graph that Walkup (1980) studies, each node is connected to d other nodes drawn independently and uniformly at random from the other side. This graph cannot be used to prove our main theorem, because every "left-neighbor" of an object in the graph would rank that object $\frac{m+1}{2}$ th on her list, on average. Thus, a perfect matching in this graph would not necessarily have a constant average rank.

²¹We recall that the size of a matching equals its number of edges and that an independent set in a graph is a subset of nodes that are pairwise nonadjacent.

Bounding the probability that a fixed subset of nodes of size $n + 1$ forms an independent set is relatively straightforward when the neighbors of a node are chosen independently, as in Walkup (1980). In the graph that we construct, however, the neighbors of nodes are *not* chosen independently. In particular, there are two types of correlations involved:

- i. *Correlations across agents and objects.* The left-neighbors of a node $o \in O$ are not independent of the right-neighbors of a node $a \in A$.
- ii. *Correlations between objects.* The left-neighbors of a node o_1 are not independent of the left-neighbors of a node o_2 .

To bound the probability that a fixed subset of nodes of size $n + 1$ forms an independent set, we need to handle the above correlations. This is the bulk of the analysis in this step, and involves applications of the FKG inequality, as outlined next.

Suppose $X \subseteq A$ and $Y \subseteq O$ are arbitrary subsets of nodes of G such that $|X| + |Y| = n + 1$. We will provide an upper bound on $p_{(X,Y)}$, the probability that $X \cup Y$ forms an independent set in G . Let R denote the event that there is no node in X which has a right-neighbor in Y . Also, let L denote the event that there is no node in Y which has a left-neighbor in X . Thus, $p_{(X,Y)} = \mathbb{P}[L \cap R]$. We will show that $\mathbb{P}[L \cap R] \leq \mathbb{P}[L]\mathbb{P}[R]$, and then will bound each of the terms on the right-hand side separately.

The inequality $\mathbb{P}[L \cap R] \leq \mathbb{P}[L]\mathbb{P}[R]$ holds intuitively because, by the construction of the bipartite graph, the event that an object o is not a right-neighbor of an agent a is negatively correlated with the event that the agent a is not a left-neighbor of the object o . We prove this negative correlation inequality formally using the FKG inequality (Claim 4.2). It remains to provide upper bounds on $\mathbb{P}[L]$ and $\mathbb{P}[R]$.

We can derive a closed-form expression for $\mathbb{P}[R]$, essentially because the right-neighbors of every node $a \in X$ are chosen independently of the right-neighbors of every other node $a' \in X$. To provide an upper bound on $\mathbb{P}[L]$, let L_y denote the event that an object $y \in Y$ has no left-neighbor in X . We first show that $\mathbb{P}[L] \leq \prod_{y \in Y} \mathbb{P}[L_y]$ (Claim 4.5). This holds intuitively because, by the construction of the bipartite graph, the event that an agent is a left-neighbor of the object y is negatively correlated with the event that the agent is a left-neighbor of another object $y' \neq y$. Finally, we provide an upper bound on $\mathbb{P}[L_y]$, which gives an upper bound on $\mathbb{P}[L]$, and thereby the promised upper bound on $p_{(X,Y)}$ (Claim 4.8).

Step B This step is relatively simpler. We show that a perfect matching in G has an expected average rank of at most $7\frac{3}{4}$ when n is sufficiently large. Define the *weight* of an edge (a, o) in G , denoted by $w_{a,o}$, to be the rank of object o on agent a 's preference list. To prove the theorem, we show that the expected rank of the maximum-weight edge adjacent to an object is bounded by $7\frac{3}{4}$. To this end, we first note that $w_{a,o} \leq 3$ if o is a right-neighbor of a and $w_{a,o} \geq 4$ if a is a left-neighbor of o . Thus, it suffices to show that

$$\mathbb{E}\left[\max_{a \in A}\{w_{a,o} : a \text{ is a left-neighbor of } o\} \mid o \text{ has a left-neighbor}\right] \leq 7\frac{3}{4}.$$

This bound is proved in Claim 4.10. To provide intuition for it, we consider the case where o has exactly 3 left-neighbors. These are the 3 agents who, among all agents that rank o worse than 3, rank o the highest. Let these 3 agents be a_1, a_2, a_3 , and let r_i denote the rank of o on agent a_i 's list. Without loss of generality, suppose $r_1 \leq r_2 \leq r_3$. The above bound then boils down to $\mathbb{E}[r_3] \leq 7\frac{3}{4}$, the proof for which is sketched below.

Since the market is uniform, the chance that no agent ranks o fourth is at most $(1 - \frac{1}{n})^n$; the chance that no agent ranks o fourth or fifth is at most $(1 - \frac{1}{n})^{2n}$; and similarly, the chance that no agent ranks o fourth, fifth, $\dots$, k th is at most $(1 - \frac{1}{n})^{(k-3)n}$. Therefore, $r_1 - 3$ has a distribution that is stochastically dominated by a geometric random variable with failure probability $q_1 = (1 - \frac{1}{n})^n$. A similar argument shows that $r_2 - r_1$ and $r_3 - r_2$, respectively, have distributions that are stochastically dominated by geometric random variables with failure probabilities $q_2 = (1 - \frac{1}{n})^{n-1}$ and $q_3 = (1 - \frac{1}{n})^{n-2}$. Therefore,

$$\mathbb{E}[r_3] \leq 3 + \frac{1}{1 - q_1} + \frac{1}{1 - q_2} + \frac{1}{1 - q_3}.$$

This concludes the proof since the right-hand side approaches $3(1 + \frac{1}{1-1/e}) \approx 7.745$ as n approaches infinity.²²

4. PROOF OF THEOREM 3.1

We first define the preliminary notions. Some of the notions that we use in the following analysis were already defined in the proof sketch (Section 3.2); however, we include all such definitions below to make this section independently readable.

For every finite set S , let $U(S)$ denote the uniform distribution over S . Also, let $\Delta(S)$ denote the probability simplex defined over S , i.e., the set of all probability distributions defined over S . For every positive integer i , let $[i]$ denote the set $\{1, \dots, i\}$.

We recall that the set of all possible preference profiles in a market is denoted by Π . The preference profile π of a uniform market is drawn from $U(\Pi)$. Thus, unless stated otherwise, it is assumed in the following analysis that the probability space is defined by the sample space Π , the event space 2^Π , and the probability function that assigns probability $\frac{|\Psi|}{|\Pi|}$ to an event $\Psi \subseteq \Pi$.

We next prove Theorem 3.1 for the case of unit capacities ($c = 1$). The proof for non-unit capacities ($c > 1$) is a corollary of the proof for unit capacities and is relegated to Appendix A.1.

Throughout the proof, let $d = 3$. First, from a given preference profile, we construct a bipartite graph G with A, O being the set of nodes on each side. We call A the left side of the graph and O the right side. The set of edges in the graph is $E = \widehat{E} \cup \widetilde{E}$, where $\widehat{E}$ and $\widetilde{E}$ are defined next. $\widehat{E}$ (the set of edges from left to right) contains an edge (a, o) for every agent a and every object o that is ranked d or better²³ on a 's preference list.

We define $\widetilde{E}$ (the set of edges from right to left) as follows. For every object $o \in O$, define r_o to be the smallest integer larger than d such that there exists at least d agents

²²The convergence to the limit is from above.

²³A rank d or better is a rank that belongs to the set $\{1, \dots, d\}$.

For the following claims, without loss of generality, we suppose that $X = \{1, \dots, k\}$ and $Y = \{1, \dots, l\}$, where $l = n - k + 1$.

CLAIM 4.1. For a k -tuple (X, Y) ,

$$p_{(X,Y)} = \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k \cdot \mathbb{P}[L|R]. \quad (4.1)$$

PROOF. We first note that $p_{(X,Y)} = \mathbb{P}[L \cap R] = \mathbb{P}[R] \mathbb{P}[L|R]$. Thus, to prove the claim it suffices to show that $\mathbb{P}[R] = \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k$. To this end, for every agent $x \in X$, define R_x as the event that agent x lists no object in Y as one of her top d objects. Thus, $R = \bigcup_{x \in X} R_x$. Since the preference list of each agent is distributed uniformly over the set of all strict orderings of objects, then

$$\mathbb{P}[R_x] = \frac{\binom{n-|Y|}{d}}{\binom{n}{d}} = \frac{\binom{k-1}{d}}{\binom{n}{d}}.$$

To complete the proof, we observe that since the agents' preference lists are iid, then

$$\mathbb{P}[R] = \prod_{x \in X} \mathbb{P}[R_x] = \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k. \quad \square$$

CLAIM 4.2 (Negative correlation of L, R). $\mathbb{P}[L|R] \leq \mathbb{P}[L]$.

PROOF. We will use the FKG inequality to prove the claim. To define the lattice required for applying the FKG inequality, we first define the set of its elements, V_l . Each element $x \in V_l$ is a sequence $(x^1, \dots, x^n)$, where x^i is a subset of $[n]$ with size l (see Figure 3).

To define the lattice order over V_l , we first need a few definitions. For any object i , let $h(i) = i$ if $i > d$ and let $h(i) = n + i$ otherwise. Define the total order $\leq$ on the set of objects $O = \{1, \dots, n\}$ by $i \leq j$ if $h(i) \leq h(j)$. 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 $u \leq v$ if $u_j \leq v_j$ holds for all j . By $\vec{u}$, we denote vector that is obtained by sorting the elements of u in an increasing order with respect to $\leq$.

For any $x, y \in V_l$, we define $x \leq_l y$ when

$$\begin{cases} \vec{y}^a \leq \vec{x}^a, & \forall a \in X \\ \vec{x}^a \leq \vec{y}^a, & \forall a \in A \setminus X. \end{cases} \quad (4.2)$$




FIGURE 3. For $n = 8$ and $l = 5$, we illustrate an element $x = (x^1, \dots, x^n)$ of the lattice where $x^i = \{1, 5, 6, 7, 8\}$ for $i \leq 4$ and $x^i = \{1, 2, 3, 4, 8\}$ for $i \geq 5$. The columns from left to right, respectively, correspond to objects $1, \dots, n$.

The next fact shows that $\leq_l$ indeed defines a lattice over V_l . Figure 4 provides a graphical representation of this lattice.

FACT 4.3. *The relation $\leq_l$ is a lattice over V_l .*

This fact is proved in Appendix A. The intuition is that $\leq_l$ is the Cartesian product of two lattices: one that orders tuples $(x^1, \dots, x^k)$ and one that orders tuples $(x^{k+1}, \dots, x^n)$.

We denote by $\mathcal{L}[l]$ the lattice defined by $\leq_l$. Since l is fixed in the proof of this claim, we drop the argument and denote the lattice by $\mathcal{L}$ throughout the proof of this claim.

We associate each preference profile π with an element of the lattice, which we denote by $\mathcal{L}_\pi = (\mathcal{L}_\pi^1, \dots, \mathcal{L}_\pi^n)$, where $\mathcal{L}_\pi^a$ is the set of ranks of the objects $1, \dots, l$ on agent




FIGURE 4. $n = 8$ and $l = 5$. Panels from left to right: the smallest element of the lattice, an element between the smallest and the largest element, and the largest element. The elements are represented in the same fashion as in Figure 3. Consider each black circle to be a *chip*. One can *move* from a smaller to a larger element in the lattice by moving a chip to an empty cell in the direction of the arrows. Since $k = 4$, the direction is the same for the first four rows, and is reversed for the last four rows.

a 's preference list. For every element x of the lattice, define

$$\Pi_x = \{\pi \in \Pi : x = \mathcal{L}_\pi\}.$$

Define the functions $f_L, f_R : \mathcal{L} \rightarrow [0, 1]$:

$$f_L(x) = \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L],$$

$$f_R(x) = \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in R].$$

In words, $f_L(x)$ is the probability that event L occurs at a preference profile π drawn uniformly at random from the set of all preference profiles whose associated element of the lattice is x . Similarly, $f_R(x)$ gives that probability but for the event R instead of L .

FACT 4.4. *For every $x \in V_l$, $f_R(x) \in \{0, 1\}$; also, for every $y \in V_l$ with $x \preceq_l y$, $f_R(x) \leq f_R(y)$, and $f_L(x) \geq f_L(y)$.*

The proof for this fact is in Appendix A. The intuition for the first part is that, for any preference profile π , its associated lattice element $\mathcal{L}_\pi$ has all the information needed (i.e., the ranks of objects in Y on agents' lists) to evaluate whether $\pi \in R$. The proof for the second part is technical. A loose intuition can be given through Figure 4: as one transforms x to y by moving the "chips" one by one along the direction of the arrows, (i) the total number of preference profiles in Π_x remains the same (i.e., $|\Pi_x| = |\Pi_y|$), (ii) the number of preference profiles in Π_x that are a member of R increases, because moving a chip along the direction of the arrows may eliminate existing edges in $\hat{E}$ but does not create any new ones, and (iii) the number of preference profiles that are a member of L decreases, because by moving a chip along the direction of the arrows one cannot transform the set $\hat{E} \cap (X \times Y)$ from being nonempty into being empty.

Given Fact 4.4, we are now ready to apply the FKG inequality on the lattice $\mathcal{L}$ with the functions f_L and f_R and the uniform probability distribution $U(V_l)$ defined over the lattice elements. By Fact 4.4, the FKG inequality implies the following *negative correlation* bound:

$$\mathbb{E}_{x \sim U(V_l)}[f_L(x)f_R(x)] \leq \mathbb{E}_{x \sim U(V_l)}[f_L(x)]\mathbb{E}_{x \sim U(V_l)}[f_R(x)].$$

We next observe that when $\pi \sim U(\Pi)$, then the lattice element $\mathcal{L}_\pi$ is distributed uniformly at random over V_l , by symmetry. Therefore, we can rewrite the above inequality as

$$\mathbb{E}_{\pi \sim U(\Pi)}[f_L(\mathcal{L}_\pi)f_R(\mathcal{L}_\pi)] \leq \mathbb{E}_{\pi \sim U(\Pi)}[f_L(\mathcal{L}_\pi)]\mathbb{E}_{\pi \sim U(\Pi)}[f_R(\mathcal{L}_\pi)]. \quad (4.3)$$

Since $f_R(\mathcal{L}_\pi) \in \{0, 1\}$ holds for all π by Fact 4.4, then the left-hand side of the above inequality equals $\mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L \cap R]$. On the right-hand side, we note that

$$\mathbb{E}_{\pi \sim U(\Pi)}[f_L(\mathcal{L}_\pi)] = \mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L], \quad \mathbb{E}_{\pi \sim U(\Pi)}[f_R(\mathcal{L}_\pi)] = \mathbb{P}_{\pi \sim U(\Pi)}[\pi \in R]$$

hold by the definition of f_L, f_R . Therefore, we can rewrite (4.3) as

$$\mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L \cap R] \leq \mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L]\mathbb{P}_{\pi \sim U(\Pi)}[\pi \in R].$$

This means $\frac{\mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L \cap R]}{\mathbb{P}_{\pi \sim U(\Pi)}[\pi \in R]} \leq \mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L]$, which is the promised claim. $\square$

Recall that L_i is the event that there is no edge in $\widehat{E}$ that connects $i \in Y$ to an element of X .

CLAIM 4.5 (Negative correlation of $\{L_i\}_{i \in Y}$). $\mathbb{P}[L] \leq \prod_{i=1}^l \mathbb{P}[L_i]$.

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

$$\mathbb{P}[L_1 \cap L_2 \cap \dots \cap L_j] \leq \prod_{i=1}^j \mathbb{P}[L_i].$$

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 $j = l' + 1$.

For the induction step, we need the following definitions. Recall (4.2), which defines the lattice $\mathcal{L}[l]$, parameterized by the parameter l . The lattice that we will use in the induction step is $\mathcal{L}[l']$. Thus, $V_{l'}$ is the set of the lattice elements; an element of the lattice is $(x^1, \dots, x^n)$, where x^i is a subset of $[n]$ with size l' ; and the lattice elements are ordered according to (4.2). For notational simplicity, we denote the lattice $\mathcal{L}[l']$ by $\mathcal{K}$.

We associate a preference profile π with an element of the lattice, denoted by $\mathcal{K}_\pi = (\mathcal{K}_\pi^1, \dots, \mathcal{K}_\pi^n)$, where $\mathcal{K}_\pi^a$ is the set of ranks of the objects $1, \dots, l'$ on agent a 's preference list. Let $\Pi_x = \{\pi \in \Pi : x = \mathcal{K}_\pi\}$.

The induction step The induction step will apply the FKG inequality on the lattice $\mathcal{K}$ and the function $f, g : V(\mathcal{K}) \rightarrow \mathbb{R}_+$, defined as follows:

$$\begin{aligned} f(x) &= \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_1 \cap \dots \cap L_{l'}], \\ g(x) &= \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_{l'+1}]. \end{aligned}$$

FACT 4.6. For every $x, y \in V(\mathcal{K})$ with $x \leq_{l'} y$, $f(x) \geq f(y)$ and $g(x) \leq g(y)$.

The proof is in Appendix A. The intuition for the first inequality is similar to the intuition for $f_L(x) \geq f_L(y)$ from Fact 4.4. A loose intuition for the second inequality can be given through Figure 4: as one transforms x to y by moving the “chips” one by one along the direction of the arrows, the total number of preference profiles in Π_x remains the same (i.e., $|\Pi_x| = |\Pi_y|$) but the number of preference profiles that are a member of $L_{l'+1}$ increases, because by moving a chip along the direction of the arrows one cannot transform the set $\widehat{E} \cap (X \times \{l' + 1\})$ from being empty into being nonempty.

By Fact 4.6, the FKG inequality implies the following negative correlation inequality:

$$\mathbb{E}_{x \sim U(V(\mathcal{K}))}[f(x)g(x)] \leq \mathbb{E}_{x \sim U(V(\mathcal{K}))}[f(x)]\mathbb{E}_{x \sim U(V(\mathcal{K}))}[g(x)]. \quad (4.4)$$

FACT 4.7. For every $x \in V(\mathcal{K})$,

$$\mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_1 \cap \dots \cap L_{l'+1}] = \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_1 \cap \dots \cap L_{l'}] \cdot \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_{l'+1}].$$

The above fact is proved in Appendix A. The intuition is that conditioning on the agents' ranks for the object $l' + 1$ in π does not affect the probability of the event $\pi \in L_1 \cap \dots \cap L_{l'}$, since $\pi \sim U(\Pi_x)$ and $x \in V(\mathcal{K})$.

By the definition of f , g ,

$$\begin{aligned} & \mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L_1 \cap \dots \cap L_{l'+1}] \\ &= \mathbb{E}_{x \sim U(V(\mathcal{K}))}[\mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_1 \cap \dots \cap L_{l'+1}]] \\ &= \mathbb{E}_{x \sim U(V(\mathcal{K}))}[\mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_1 \cap \dots \cap L_{l'}] \cdot \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_{l'+1}]] \\ &\leq \mathbb{E}_{x \sim U(V(\mathcal{K}))}[\mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_1 \cap \dots \cap L_{l'}]] \cdot \mathbb{E}_{x \sim U(V(\mathcal{K}))}[\mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_{l'+1}]] \\ &= \mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L_1 \cap \dots \cap L_{l'}] \cdot \mathbb{P}_{\pi \sim U(\Pi)}[\pi \in L_{l'+1}], \end{aligned}$$

where in the first equality we have used the fact that the distribution imposed on π when $x \sim U(V(\mathcal{K}))$ and $\pi \sim U(\Pi_x)$ is just the uniform distribution over Π . The second equality is by Fact 4.7, and the inequality is by (4.4). The above bound completes the induction step and proves the claim. $\square$

The next claim gives a closed-form upper bound for $p_{(X,Y)}$ by bounding the right-hand side of (4.1) from above. We then use this claim to provide a closed-form upper bound on p_k , and prove that $\sum_{k=1}^n p_k \leq n^{-3}$.

CLAIM 4.8.

$$p_{(X,Y)} \leq \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k (\omega + \delta)^l, \quad (4.5)$$

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

PROOF. By (4.1) and Claims 4.2 and 4.5, it suffices to show that $\mathbb{P}[L_i] \leq \omega + \delta$ for $i \in Y$. 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

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

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

(i) $\mathbb{P}[\overline{D}] \leq \delta$ Let $\overline{D}_j$ be the event in which object i is ranked worse than d by exactly j agents. Observe that

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

Therefore, we can write

$$\mathbb{P}[\bar{D}] = \sum_{j=0}^{d-1} \mathbb{P}[\bar{D}_j] \leq d \binom{n}{d} \left(\frac{d}{n}\right)^{n-d} \leq \frac{d^{n-2d+1} e^d}{n^{n-2d}} = \delta, \quad (4.6)$$

where the first inequality holds since $d = 3 \leq \frac{n}{2}$, and the second one holds since $\binom{n}{d} \leq \left(\frac{ne}{d}\right)^d$ (Das (2016)).

(ii) $\mathbb{P}[L_i|D] \leq \omega$ Let A_j be the set of agents who assign rank j to object i , and let $A^h = \bigcup_{j=d+1}^h A_j$ be such that h is the smallest integer for which $|A^h| \geq d$. Such h exists when the event D holds. Observe that the set A^h is a random variable whose distribution, conditional on its size being s , is the uniform distribution over the set of all subsets of A with size s . (This holds because the agents' preference lists are iid from the uniform distribution.) Therefore, for every integer $g \geq 1$, we can write

$$\mathbb{P}[L_i|g = |A^h|] = \frac{\binom{n-k}{g}}{\binom{n}{g}} = \prod_{j=1}^g \frac{n-k-j+1}{n-j+1},$$

where in the middle term the numerator is the number of subsets of $A \setminus X$ of size g , and the denominator is the number of subsets of A of size g . Recall that event D holds if and only if h exists (and if h exists, then $|A^h| \geq d$ would also hold). Since the right-hand side of the above equation is decreasing in g , and since $|A^h| \geq d$, then

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

which is the promised bound. $\square$

We now use Claim 4.8 and a union bound over all k -tuples (X, Y) with $|X| = k$ to write

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

where recall that $l = n - k + 1$, and the last inequality holds since the term $(\omega + \delta)^l$ has 2^l summands when expanded and $\omega, \delta \leq 1$. We therefore can bound $p(n)$ as follows:

$$\begin{aligned}
 p(n) &= \sum_{k=1}^n p_k \leq \sum_{k=1}^n \binom{n}{k} \binom{n}{l} \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k (\omega^l + 2^l \delta) \\
 &= \sum_{k=1}^n \binom{n}{k} \binom{n}{k-1} \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k \cdot \omega^l \\
 &\quad + \sum_{k=1}^n \binom{n}{k} \binom{n}{k-1} \left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k \cdot 2^l \delta.
 \end{aligned} \tag{4.7}$$

Let $S_1(n)$ and $S_2(n)$, respectively, denote the first and the second summand in (4.7). Walkup (1980) shows that²⁵

$$S_1(n) \leq \frac{1}{122} \left(\frac{d}{n} \right)^{(d+1)(d-2)}. \tag{4.8}$$

We complete the proof by providing an upper bound on $S_2(n)$. Since $\binom{n}{j} \leq 2^n$ for all positive integers j , and $\left[\frac{\binom{k-1}{d}}{\binom{n}{d}} \right]^k \leq 1$, then

$$S_2(n) \leq 2^{3n} \delta = 2^{3n} \cdot \frac{d^{n-2d+1} e^d}{n^{n-2d}}. \tag{4.9}$$

The following fact uses (4.8) and (4.9) to bound $S_1(n) + S_2(n)$ from above.

FACT 4.9. For $d = 3$, $S_1(n) + S_2(n) \leq n^{-3}$ for $n \geq 50$.

PROOF. When $d = 3$ and $n \geq 50$, (4.8) implies that

$$S_1(n) \leq \frac{81}{122n} \left(\frac{1}{n} \right)^3 \leq \frac{1}{2n^3}. \tag{4.10}$$

Also, (4.9) implies that

$$S_2(n) \leq 2^{3n} \cdot \frac{3^{n-5} e^3}{n^{n-6}}. \tag{4.11}$$

²⁵This is the last bound in the proof of Theorem 1 in Walkup (1980).

Let $g(n)$ denote the right-hand side of the above inequality. Define $h(n) = 2n^3 g(n)$. We observe that

$$h'(n) = e^3 2^{3n+1} 3^{n-5} n^{8-n} (n(\log(24) - 1) - n \log(n) + 9).$$

Since

$$n(\log(24) - 1) - n \log(n) + 9 < 0$$

for $n \geq 50$, then $h'(n) < 0$ for such n . Therefore, $h(n) \leq h(50) < 1$. Recall that $h(n) = 2n^3 g(n)$. Thus, $g(n) \leq \frac{1}{2n^3}$ for $n \geq 50$. Together with (4.11), this implies that $S_2(n) \leq \frac{1}{2n^3}$. This bound, together with (4.10), implies that $S_1(n) + S_2(n) \leq n^{-3}$ for $n \geq 50$. $\square$

By (4.7), $p(n) \leq S_1(n) + S_2(n)$. By Fact 4.9, the right-hand side is at most n^{-3} when $n \geq 50$. This proves the promised bound on $p(n)$ and completes step A.

4.2 Step B: Bounding the average rank in the perfect matching

Define the *weight* of an edge (a, o) in the graph G to be the rank of object o on agent a 's preference list. We showed that G has a perfect matching with probability at least $1 - n^{-3}$ when $n \geq 50$. To bound the expected average rank 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 adjacent to o .)

For any object o , we provide an upper bound on $\mathbb{E}[w_o]$, where the expectation is taken over $\pi \sim U(\Pi)$. Since all of the expectation and probability operators are with respect to $\pi \sim U(\Pi)$ in the rest of the proof, we drop this subscript from the notation.

Let $w = \sum_{o \in O} w_o$, and $\mathbf{m}$ denote the event that G has a perfect matching. Since $\mathbb{P}[\mathbf{m}] \mathbb{E}[w|\mathbf{m}] + \mathbb{P}[\bar{\mathbf{m}}] \mathbb{E}[w|\bar{\mathbf{m}}] = \mathbb{E}[w]$, then $\mathbb{P}[\mathbf{m}] \mathbb{E}[w|\mathbf{m}] \leq \mathbb{E}[w]$. Because $\mathbb{P}[\mathbf{m}] \geq 1 - n^{-3}$ as shown in Step A, then

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

Therefore, to complete Step B, it suffices to bound $\mathbb{E}[w]$ from above.

The next claim will prove that $\mathbb{E}[w] \leq 3n(1 + \frac{1}{1 - (1 - \frac{1}{n})^{n-2}})$. This would complete the proof because this inequality together with the above bound implies that, for $n \geq 50$,

$$\frac{1}{n} \mathbb{E}[w|\mathbf{m}] \leq \frac{3 \left(1 + \frac{1}{1 - \left(1 - \frac{1}{n} \right)^{n-2}} \right)}{1 - n^{-3}} < 7.9,$$

where the last inequality holds because the middle term is decreasing in n (since $1 - (1 - \frac{1}{n})^{n-2}$ is increasing in n), and the value of the middle term at $n = 50$ is less than 7.9. Recall that the probability that G has no perfect matching $\mathbb{P}[\bar{\mathbf{m}}]$ is at most n^{-3} . When G has no perfect matching, the rank-optimal assignment has average rank at most n . Thus, the expected average rank of the rank-optimal assignment is at most $7.9 + n^{-2}$. (We

remark that this bound can be improved for sufficiently large n , as in the latter inequality, the middle term approaches $3 + 3\frac{e}{e-1}$ as n approaches infinity. Thus, for sufficiently large n , the bound R on the expected average rank improves to any constant larger than $3 + 3\frac{e}{e-1} \approx 7.746$.)

To prove the theorem, it remains to prove the next claim.

CLAIM 4.10. $\mathbb{E}[w] \leq 3n(1 + \frac{1}{1-(1-\frac{1}{n})^{n-2}})$.

PROOF. Equivalently, we will show that for an object $o \in O$, $\mathbb{E}[w_o] \leq 3(1 + \frac{1}{1-(1-\frac{1}{n})^{n-2}})$. Let $\mathcal{M}$ be an $n \times n$ matrix such that, for every $a \in A$, $\mathcal{M}_{a,1}, \dots, \mathcal{M}_{a,n}$ is the preference list of agent a where objects are ordered from the most to the least preferred. We will use the principle of deferred decisions for the proof. The idea is that, rather than drawing the preference profile π (associated with the matrix $\mathcal{M}$) at once, we determine the elements of the matrix $\mathcal{M}$ one by one through a stochastic process, namely $\mathcal{P}$.

The process $\mathcal{P}$ starts from an empty matrix $\mathcal{M}$ (with the values of its entries undetermined) and determines the values of the entries of $\mathcal{M}$ one by one. At round i of the process, the values of the entries in column $i+3$ (i.e., $\mathcal{M}_{1,i+3}, \dots, \mathcal{M}_{n,i+3}$) will be determined one by one. The process stops at the moment that object o appears $d=3$ times in the matrix, or at the end of round $n-3$, whichever occurs first. Let s denote the round at which the process stops. We will complete the proof in two steps by (i) showing that $s+3 \geq w_o$ and (ii) providing an upper bound on $\mathbb{E}[s]$.

Step (i): $s+3 \geq w_o$ Every edge in $\hat{E}$ that is adjacent to o has weight of at most 3. To see why, recall that $\hat{E}$ contains an edge (a, o) for every agent a that ranks o third or better on her preference list.

Also, every edge in $\hat{E}$ that is adjacent to o has weight of at most $s+3$. To see why, we recall the definition of $\hat{E}$. We defined r_o to be the smallest integer greater than $d=3$ such that there exists at least 3 agents who rank o (strictly) worse than 3 and no worse than r_o . If there is no such integer r_o , then we defined $r_o = n$. $\hat{E}$ contains an edge (o, a) for every agent a that ranks o better than 3 and no better than r_o . Thus, process $\mathcal{P}$ stops in round $r_o - 3$. That is, $r_o = s+3$. By definition, the weight of every edge in $\hat{E}$ that is adjacent to o is at most r_o . Therefore, the weight of every edge in $\hat{E}$ that is adjacent to o is at most $s+3$. This bound, together with the fact that the weights of the edges in $\hat{E}$ are at most 3, concludes step (i).

Step (ii): An upper bound on $\mathbb{E}[s]$ Consider a round $t \geq 1$ in process $\mathcal{P}$ and an agent a . If agent a has not listed object o on her preference list in the previous rounds (i.e., if $o \notin \{\mathcal{M}_{a,4}, \dots, \mathcal{M}_{a,t+2}\}$), then $\mathcal{P}$ sets $\mathcal{M}_{a,t+3} = o$ with probability $\frac{1}{n-t+1}$, which is at least $\frac{1}{n}$. Conditional on process $\mathcal{P}$ not having stopped, there are at most 2 agents who have listed object o on their list in previous rounds. Thus, the probability that during round t no agent lists object o on her list is at most $(1 - \frac{1}{n})^{n-2}$. Process $\mathcal{P}$ stops at a round s when at least 3 agents list object o , or by the end of round $s = n-3$ if that event never occurs. Since the chance that no agent lists o during a round is at most $(1 - \frac{1}{n})^{n-2}$, then

s is stochastically dominated by the sum of three iid geometric random variables with failure probabilities $(1 - \frac{1}{n})^{n-2}$. Thus, $\mathbb{E}[s] \leq \frac{3}{1 - (1 - \frac{1}{n})^{n-2}}$.

Finally, from steps (i) and (ii) it follows that

$$\mathbb{E}[w_o] \leq \mathbb{E}[s + 3] \leq 3 + \frac{3}{1 - \left(1 - \frac{1}{n}\right)^{n-2}},$$

which concludes the proof. $\square$

REMARK 4.11. The above analysis can also provide an asymptotic upper bound—in the sense of first-order stochastic dominance—on the distribution of the rank of the object assigned to an agent, as described next. Let μ denote a perfect matching of G chosen uniformly at random; if G has no perfect matching, which can occur with probability at most n^{-3} , then draw an assignment of agents to objects uniformly at random and denote it by μ . Define the random variable p_a as the rank of the object assigned to agent a in μ . Also, define the random variable q_o as the rank that agent $\mu(o)$ assigns to o on her list. Observe that $\mathbb{P}[p_a = i] = \mathbb{P}[q_o = i]$ for every $a \in A$, $o \in O$, and rank $i \in [n]$ (This holds because the expected number of agents who are assigned to their i th choice in μ is equal to $\sum_{a' \in A} \mathbb{P}[p_{a'} = i] = n\mathbb{P}[p_a = i]$, and also equal to $\sum_{o' \in O} \mathbb{P}[q_{o'} = i] = n\mathbb{P}[q_o = i]$). Therefore, showing that q_o is first-order stochastically dominated by some distribution D also implies that p_a is first-order stochastically dominated by D . When G has a perfect matching, $q_o \leq w_o$. Step (ii) in the proof of Claim 4.10 proves that w_o is stochastically dominated by a random variable $Z = 3 + \sum_{j=1}^3 Z_j$ where Z_j 's are iid geometric random variables with failure probabilities $(1 - \frac{1}{n})^{n-2}$. Hence, $\mathbb{P}[q_o > i] \leq n^{-3} + \mathbb{P}[Z > i]$; which also means that $\mathbb{P}[p_a > i] \leq n^{-3} + \mathbb{P}[Z > i]$. This is the promised asymptotic upper bound.

5. MECHANISMS FOR ELICITING PRIVATE PREFERENCES

In this section, we take a mechanism design perspective by considering the setting where every agent's preference list is her private information. After noting that there exists no *Dominant-Strategy Incentive-Compatible* (DSIC) mechanism that implements a rank-optimal assignment, we show that such mechanisms exist if the notion of incentive compatibility is relaxed to *Bayesian Incentive Compatibility* (BIC).

Consider a market, as defined in our main setup in Section 3, with m objects each with capacity c and $n = mc$ agents. Given a preference profile π in this market, we use π_i to denote the preference list of agent i and the vector $\pi_{-i} = (\pi_j)_{j \in A, j \neq i}$ to denote the preference lists of the rest of the agents. Let Π_{-i} denote the set of all such possible π_{-i} . Given a preference list σ , we sometimes use $[\sigma; \pi_{-i}]$ to denote a preference profile π where $\pi_i = \sigma$ and π_{-i} determines the preference lists of the agents in $A \setminus \{i\}$. Finally, let M denote the set of all assignments in the market.

5.1 DSIC mechanisms

A mechanism is a function $\mathcal{M} : \Pi \rightarrow M$. It elicits a preference list π_i from every agent i , constructs the preference profile $\pi = (\pi_1, \dots, \pi_n)$, and then assigns objects to agents according to $\mathcal{M}(\pi)$. For every agent i , let $\mathcal{M}_i(\pi)$ denote the object assigned to i in the assignment $\mathcal{M}(\pi)$. A mechanism is DSIC if for every agent i , every preference order σ over O , and every π_{-i} , agent i weakly prefers the object $\mathcal{M}_i(\pi)$ to the object $\mathcal{M}_i([\sigma; \pi_{-i}])$.

A mechanism $\mathcal{M}$ is *rank-optimal* if, for every preference profile π , $\mathcal{M}(\pi)$ is a rank-optimal assignment (with respect to the objects' ranks in π). From a result of Featherstone (2020), it follows that no DSIC rank-optimal mechanism exists in a market with 4 agents and 4 objects each with capacity 1.²⁶ The proof is by constructing a specific market. Our next proposition shows that rank-optimal DSIC mechanisms exist in no parametrization of our setup (i.e., choices of c , m) when $m \geq 4$. The proof is by constructing a (different) market for every such c , m .

PROPOSITION 5.1. *No rank-optimal DSIC mechanism exists when $m \geq 4$.*

A well-known DSIC mechanism that is applicable to our setup is the Random Serial Dictatorship (RSD) mechanism (Abdulkadiroglu and Sönmez (1998)). RSD orders agents uniformly at random and then allows each agent to choose her most preferred object from the set of remaining objects, one agent at a time in that order. We call the assignment generated by RSD the *RSD assignment*. Remarkably, there is a quite large gap between the average ranks of the RSD and rank-optimal assignments. When $c = 1$, Knuth (1996) shows that the expected average rank of the RSD assignment almost equals $\ln n$, and thus approaches infinity as n does, in contrast to the average rank of the rank-optimal assignment.²⁷ We next show that the gap between the average ranks of these two assignments persists even when the object capacities are greater than one—in particular, when $c = o(\log n)$.²⁸

PROPOSITION 5.2. *The expected average rank of the RSD assignment is at least $\frac{\ln m - 1}{c}$.*

It remains an open question whether there exist DSIC mechanisms with a constant average rank (i.e., an average rank independent of the market size). We remark that the answer is negative if the mechanism is required to be *Pareto efficient*²⁹ and *nonbossy*³⁰ as well. This is because any such mechanism is equivalent to RSD, as shown by Bade (2020).

We next relax the notion of DSIC to BIC and ask whether rank-optimal BIC mechanisms exist in our main setup.

²⁶His result is in fact stronger, in that he shows that no DSIC mechanism can be *rank efficient* (in the sense of first-order stochastic dominance) in the market that he constructs.

²⁷We remark that if there is a linear excess of objects, i.e., $m = (1 + \epsilon)n$ for a constant $\epsilon > 0$, then the assignment generated by RSD has a constant average rank (e.g., see Ashlagi and Nikzad (2020, Lemma D.15)).

²⁸We recall that for two functions $f, g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, the notation $f = o(g)$ means $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$.

²⁹An assignment is Pareto efficient if no other assignment exists in which all agents are weakly better off with at least one agent being strictly better off.

³⁰Nonbossyness essentially means that an agent can only change someone else's match if she also changes her own match. See Bade (2020) for the formal definitions and result.

5.2 BIC mechanisms in uniform markets

To formally define BIC mechanisms, we allow a mechanism to be a function $\mathcal{M} : \Pi \rightarrow \Delta(M)$; that is, $\mathcal{M}(\pi)$ is a probability distribution over assignments for every $\pi \in \Pi$.

For an assignment μ and an agent i , define $\mathbb{1}_i^\mu \in \{0, 1\}^m$ to be a vector v such that $v_o = 1$ if and only if $o = \mu(i)$. Let $\mathcal{M}_i(\pi) = \mathbb{E}_{\mu \sim \mathcal{M}(\pi)}[\mathbb{1}_i^\mu]$. The *interim allocation rule* η of a mechanism $\mathcal{M}$ is a sequence of functions $\eta_1, \dots, \eta_n$ where, for every preference list σ and every agent i , $\eta_i(\sigma) = \mathbb{E}_{\pi_{-i} \sim U(\Pi_{-i})}[\mathcal{M}_i([\sigma; \pi_{-i}])]$.

A mechanism $\mathcal{M}$ with an interim allocation rule η is BIC if, for every agent $i \in A$ with preference order π_i , and every strict preference order σ over the set of objects, the following holds. Let $v = \eta_i(\pi_i)$ and $v' = \eta_i(\sigma)$; also, let $\preceq_i = \pi_i$,³¹ then, for every $o \in O$,

$$\sum_{q \in O: o \preceq_i q} v_q \geq \sum_{q \in O: o \preceq_i q} v'_q. \quad (5.1)$$

Let $\mathcal{M}^*$ be the mechanism in which, for every $\pi \in \Pi$, $\mathcal{M}^*(\pi)$ is the uniform distribution over all rank-optimal assignments. (Thus, if the rank-optimal assignment is unique for a given preference profile π , then $\mathcal{M}^*(\pi)$ is simply the unique rank-optimal assignment.)

PROPOSITION 5.3. *The mechanism $\mathcal{M}^*$ is BIC.*

This shows the existence of rank-optimal BIC assignments in uniform markets (i.e., where the agents' preference orders are drawn iid and uniformly at random). The existence of rank-optimal BIC mechanisms in markets where the agents' preference orders are correlated or have a nonuniform distribution remains an interesting open question.

6. CONCLUSION

We show that in markets where agents rank objects uniformly at random, the average rank of the rank-optimal assignment (i.e., the first-best average rank) is bounded from above by a constant independent of the number of agents, and thus does not grow infinitely large as the number of agents does. At the same time, the average rank under the RSD mechanism can grow infinitely large when the object capacities grow at a sufficiently slow rate. Thus, the average rank can take a heavy toll under the RSD mechanism compared to the first-best. Furthermore, while no DSIC mechanism can implement the rank-optimal assignment when the agents' rankings are their private information, we show that there exists a BIC mechanism that does so. These findings highlight the possibility of using assignment methods other than RSD in scenarios where the average rank is a relevant objective.

APPENDIX A: PROOFS FROM SECTION 4

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 $\pi : A \rightarrow O^{|A|}$ that

³¹Hence, $o_1 \preceq_i o_2$ means that object o_1 is ranked weakly lower than o_2 on the preference list π_i .

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 .

PROOF OF FACT 4.3. We will show that the lattice defined by the relation $\preceq_l$ is the Cartesian product of two lattices, namely $\mathcal{L}^1$, $\mathcal{L}^2$. The claim then would be proved since the Cartesian product of two lattices is a lattice itself.

The set of the elements of the lattice $\mathcal{L}^1$ is denoted by V_l^1 . An element $w \in V_l^1$ is a sequence $(w^1, \dots, w^k)$ where w^i is a subset of $[n]$ with size l . Similarly, the set of the elements of $\mathcal{L}^2$ is V_l^2 . An element $w \in V_l^2$ is a sequence $(w^{k+1}, \dots, w^n)$ where w^i is a subset of $[n]$ with size l .

Recall that $h(i) = i$ if $i > d$ and let $h(i) = n + i$ otherwise. Also, recall that for integers α, β , we defined $\alpha \preceq \beta$ if $h(\alpha) \leq h(\beta)$. Similarly, for two vectors of the same length u, v , $u \preceq v$ if $u_j \preceq v_j$ for every coordinate j .

For $w, z \in V_l^1$, we define $w \preceq_l^1 z$ when $\vec{w}^a \preceq \vec{z}^a$ for all $a \in X$. We observe that $\preceq_l^1$ is a lattice, because it has a well-defined meet operator corresponding to coordinatewise maximum with respect to $\preceq$, and a well-defined join operator corresponding to coordinatewise minimum with respect to $\preceq$.

For $w, z \in V_l^2$, define $w \preceq_l^2 z$ if $\vec{w}^a \preceq \vec{z}^a$ for all $a \in A \setminus X$. The relation $\preceq_l^2$ is a lattice, because it has a well-defined meet operator corresponding to coordinatewise minimum with respect to $\preceq$, and a well-defined join operator corresponding to coordinatewise maximum with respect to $\preceq$.

Consider $x, y \in V^l$ with $x = (x^1, \dots, x^n)$ and $y = (y^1, \dots, y^n)$. By the definition of $\preceq_l$, we have $x \preceq_l y$ if

$$(x^1, \dots, x^k) \preceq_l^1 (y^1, \dots, y^k), \quad (\text{A.1})$$

$$(x^{k+1}, \dots, x^n) \preceq_l^2 (y^{k+1}, \dots, y^n). \quad (\text{A.2})$$

That is, the relation $\preceq_l$ is the Cartesian product of the lattices $\preceq_l^1$ and $\preceq_l^2$. Thus, $\preceq_l$ itself is a lattice. $\square$

PROOF OF FACT 4.4. Recall that R denotes the event $(X \times Y) \cap \widehat{E} = \emptyset$ where $\widehat{E}$ contains an edge (a, o) for every agent a and every object o that is ranked d or better on a 's preference list. Consider an arbitrary preference profile $\pi \in \Pi$ and recall that $\mathcal{L}_\pi = (\mathcal{L}_\pi^1, \dots, \mathcal{L}_\pi^n)$ where, for every $a \in A$, $\mathcal{L}_\pi^a$ is the set of ranks of objects $1, \dots, l$ on agent a 's preference list. We observe that $\pi \in R$ if and only if $\mathcal{L}_\pi^a \cap \{1, \dots, d\} = \emptyset$ for every $a \in X$. Therefore, $f_R(\mathcal{L}_\pi) = 1$ if the latter condition holds; otherwise, $f_R(\mathcal{L}_\pi) = 0$. Thus, $f_R(\mathcal{L}_\pi) \in \{0, 1\}$ for every $\pi \in \Pi$. This proves the first claim since, for every $x \in V_l$, there exists $\pi \in \Pi$ such that $\mathcal{L}_\pi = x$.

To prove the second claim, consider preference profiles π, π' such that $\mathcal{L}_\pi \preceq_l \mathcal{L}_{\pi'}$. For every agent $a \in X$, $\mathcal{L}_\pi^a \cap \{1, \dots, d\} = \emptyset$ implies that $\mathcal{L}_{\pi'}^a \cap \{1, \dots, d\} = \emptyset$, because $\mathcal{L}_\pi \preceq_l \mathcal{L}_{\pi'}$. Thus, $f_R(\mathcal{L}_\pi) = 1$ implies that $f_R(\mathcal{L}_{\pi'}) = 1$, which proves that $f_R(x) \leq f_R(y)$.

It remains to prove that $f_L(x) \geq f_L(y)$. We need a few definitions first. An *extension* of an element $x = (x^1, \dots, x^n)$ of the lattice, namely χ , is defined as follows. Recall that, for each $a \in A$, x^a is subset of $[n]$ of size l . Then, we define $\chi = (\chi^1, \dots, \chi^n)$, where

χ^a is an l -dimensional vector that contains the elements of x^a in some order. (Thus, every element x of the lattice has $(l!)^n$ distinct extensions.) The interpretation of χ is as follows. Suppose $\chi^a = (\chi^{a,1}, \dots, \chi^{a,l})$. Then $\chi^{a,j}$ is the rank of object j on the agent a 's preference list.

We say a preference profile π is *aligned* with the extension χ if, for every agent $a \in A$ and object $j \in Y$, $\chi^{a,j}$ equals the rank of object j on agent a 's preference list in π . (Thus, are $((n-l)!)^n$ distinct preference profiles that are aligned with χ , as the objects in $O \setminus Y$ can be ranked in $(n-l)!$ different ways on each agent's preference list.)

An extension χ is *consistent* with event L if there exists a preference profile π aligned with χ such that $\pi \in L$. (We observe that if there exists one such preference profile π , then for any preference profile π' aligned with χ , $\pi' \in L$. This holds because only the ranks of the objects in $Y = \{1, \dots, l\}$ on the agents' lists determine whether a preference profile belongs to L , and these ranks are determined by χ .)

Let Ξ_x denote the number of extensions of x consistent with L . Therefore,

$$f_L(x) = \frac{\Xi_x((n-l)!)^n}{(l!)^n((n-l)!)^n} = \frac{\Xi_x}{(l!)^n}.$$

To see why this equality holds, observe that in the middle term (i) the numerator is the total number of preference profiles belonging to $\Pi_x \cap L$, and (ii) the denominator is the total number of preference profiles in Π_x .

By the above equation, to prove $f_L(x) \geq f_L(y)$, it suffices to show that $\Xi_x \geq \Xi_y$. We will prove this by constructing a bijection B from the set of extensions of y to the set of extensions of x such that for every extension ψ of y , $B(\psi)$ is consistent with L if ψ is consistent with L .

Recall that $x = (x^1, \dots, x^n)$ and $y = (y^1, \dots, y^n)$ where x^a and y^a are subsets of $[n]$ with size l . For every a , let $\vec{x}^a = (x^{a,1}, \dots, x^{a,l})$ and $\vec{y}^a = (y^{a,1}, \dots, y^{a,l})$. Let C^a be a bijection from y^a to x^a such that $C^a(y^{a,j}) = x^{a,j}$ for every $j \in [l]$. For an extension $\psi = (\psi^1, \dots, \psi^n)$ of y with $\psi^a = (\psi^{a,1}, \dots, \psi^{a,l})$ for every $a \in A$, define $B(\psi)$ to be the extension $\chi = (\chi^1, \dots, \chi^n)$ with

$$\chi^a = (C^a(\psi^{a,1}), \dots, C^a(\psi^{a,l})).$$

Since $x \preceq y$, then $y^{a,j} \preceq C^a(y^{a,j})$ for $a \in X$, and $C^a(y^{a,j}) \preceq y^{a,j}$ for $a \notin X$. Therefore, if ψ is consistent with L , $B(\psi)$ is also consistent with L . This proves the promised claim. $\square$

PROOF OF FACT 4.6. The proof for f follows from Fact 4.4, as described next. In Fact 4.4, we showed that $f_L(x) \geq f_L(y)$ for $x \preceq_l y$, where

$$f_L(x) = \mathbb{P}_{\pi \sim U(\Pi_x)}[\pi \in L_1 \cap \dots \cap L_l].$$

The function f here is the same as the function f_L , with the difference that f is defined over the lattice $\mathcal{L}[l']$ whereas f_L is defined over $\mathcal{L}[l]$. Thus, the proof for $f(x) \geq f(y)$ if $x \preceq_{l'} y$ is the same as the proof for $f_L(x) \geq f_L(y)$ if $x \preceq_l y$ but for l replaced with l' .

To prove the claim for g , we recall the lattice $\mathcal{K} = \mathcal{L}[l']$ defined over the set $V_{l'}$. Also, recall that we say a preference profile π is associated with an element of the lattice, denoted by $\mathcal{K}_\pi = (\mathcal{K}_\pi^1, \dots, \mathcal{K}_\pi^n)$ where $\mathcal{K}_\pi^a$ is the set of ranks of the objects $1, \dots, l'$ on the agent a 's preference list. We let $\Pi_x = \{\pi \in \Pi : x = \mathcal{K}_\pi\}$.

Recall that $g(z) = \mathbb{P}_{\pi \sim U(\Pi_z)}[\pi \in L_{l'+1}]$ for every $z \in V_{l'}$. Thus, to prove that $g(x) \leq g(y)$, it suffices to prove the following: there exists a bijection B from Π_x to Π_y such that if $\pi \in L_{l'+1}$ for a preference profile $\pi \in \Pi_x$, then $B(\pi) \in L_{l'+1}$. We construct B as follows.

Let $x = (x^1, \dots, x^n)$ and $y = (y^1, \dots, y^n)$ be the lattice elements. Consider $\pi \in \Pi_x$. Let $r_a^\pi(o)$ denote the rank of object o on a 's preference list in the preference profile π . We define $B(\pi)$ to be the unique preference profile π' in which, for every agent $a \in A$ it holds that³²

- i. for all objects $o_1, o_2 \in [l']$, $r_a^\pi(o_1) \leq r_a^\pi(o_2)$ if and only if $r_a^{\pi'}(o_1) \leq r_a^{\pi'}(o_2)$ and
- ii. for all objects $o_1, o_2 \notin [l']$, $r_a^\pi(o_1) \leq r_a^\pi(o_2)$ if and only if $r_a^{\pi'}(o_1) \leq r_a^{\pi'}(o_2)$.

We note that the set of ranks $\{r_a^\pi(o) : o \in [l']\}$ and $\{r_a^{\pi'}(o) : o \in [l']\}$ are determined by x^a and y^a , respectively. Thus, the above two conditions uniquely determine the preference list of every agent a in $\pi' = B(\pi)$. Figure 5 demonstrates the construction of π' by an example.

By the construction of π' , $r_{\pi'}^a(l' + 1) \leq r_\pi^a(l' + 1)$ for $a \in [k]$ and $r_\pi^a(l' + 1) \leq r_{\pi'}^a(l' + 1)$ for $a \notin [k]$ hold since $x \leq_{l'} y$. This implies that if $\pi \in L_{l'+1}$, then $\pi' \in L_{l'+1}$. That is, if $\pi \in L_{l'+1}$, then $B(\pi) \in L_{l'+1}$, which is the promised claim. $\square$




FIGURE 5. $n = 8$, $k = 4$, and $l' = 3$. The left and right panels correspond to $\pi \in \Pi_x$ and $\pi' \in \Pi_y$, respectively. The rows correspond to the preference lists of agents 1 to n , from top to bottom, respectively. On each list the objects are ordered from left to right in the order of the most favorite to the least favorite. We have $x^1 = \{1, 3, 8\}$ (these positions are marked with a circle in the first row of π), $x^2 = \{1, 6, 8\}$, $x^3 = \{1, 3, 8\}$, $x^4 = \{2, 4, 7\}$, $x^5 = x^6 = x^7 = x^8 = \{4, 5, 6\}$. Also, $y^1 = y^2 = \{4, 5, 6\}$, $y^3 = \{5, 7, 8\}$, $y^4 = \{1, 4, 5\}$, $y^5 = y^6 = \{1, 3, 8\}$, and $y^7 = y^8 = \{1, 3, 6\}$.

³²For the following definition, we recall that for integers α, β , we defined $\alpha \leq \beta$ if $h(\alpha) \leq h(\beta)$, where $h(i) = i$ if $i > d$ and $h(i) = n + i$ otherwise.

PROOF OF FACT 4.7. Recall that x is a sequence $(x^1, \dots, x^n)$ where, for every $a \in A$, x^a is the set of ranks of objects $1, \dots, l$ on agent a 's preference list. For a subset $Q \subseteq O$, a *partial preference profile* $r^Q : A \times Q \rightarrow [n]$ is a function such that $r^Q(a, o) \neq r^Q(a, o')$ for every distinct $o, o' \in Q$, and $a \in A$. Given a preference profile π and a subset $Q \subseteq O$, we let π^Q denote a partial preference profile such that $\pi^Q(a, o)$ denotes the rank of an object $o \in Q$ on an agent a 's preference list.

Let $Y' = \{1, \dots, l'\}$. Since $\pi \in \Pi_x$, then, an object in Y' has a rank belonging to the set x^a , and object $l' + 1$ has a rank belonging to the set $[n] \setminus x^a$ on the preference list of an agent a . Thus, when $\pi \sim U(\Pi_x)$, conditioning on $\pi^{\{l'+1\}}$ does not change the distribution of $\pi^{Y'}$ (compared to the unconditional distribution of $\pi^{Y'}$). Since $\pi^{Y'}$ and $\pi^{\{l'+1\}}$ are respectively sufficient statistics for determining whether the events $\pi \in L_1 \cap \dots \cap L_{l'}$ and $\pi \in L_{l'+1}$ hold, then these events are independent when $\pi \sim U(\Pi_x)$. $\square$

A.1 Proof for the case of nonunit capacities

The proof for nonunit capacities ($c > 1$) is a corollary of the proof for unit capacities. We first need a few definitions. The *rank distribution* of an assignment $\mu : A \rightarrow O$ is a vector $\mathcal{R}_\mu \in \mathbb{R}_+^{|O|}$ where its i th entry $\mathcal{R}_\mu(i)$ denotes the fraction of agents who are assigned by μ to their i th favorite object (i.e., the object that they rank i).

Consider a uniform market M with its preference profile π drawn uniformly at random from the set of all preference profiles possible in M . Let μ_π denote the rank-optimal assignment in M . The *expected* rank distribution of the rank-optimal assignment in M is the vector $\mathbb{E}_\pi[\mathcal{R}_{\mu_\pi}]$.

We say a vector $\mathcal{R}' \in \mathbb{R}_+^{n_1}$ with $\sum_{i=1}^{n_1} \mathcal{R}'(i) = 1$ *rankwise* dominates another vector $\mathcal{R} \in \mathbb{R}_+^{n_2}$ with $\sum_{i=1}^{n_2} \mathcal{R}(i) = 1$ if for all positive integers $r \leq \min\{n_1, n_2\}$,

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

We will use the next lemma to prove Theorem 3.1 for the case of $c > 1$.

LEMMA A.1. *Let M, M' be two uniform 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 $\mathcal{R}'$ rankwise dominates $\mathcal{R}$.*

PROOF. Let r_π denote the average rank of the rank-optimal assignment in a market with preference profile π . Let Π, Π' denote the set of possible preference profiles in M, M' . The proof works by defining a function $f : \Pi \rightarrow \Pi'$ 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 rankwise dominated by the rank distribution of the rank-optimal assignment for $f(\pi)$. To prove the lemma, it suffices to show that such a function exists as follows.

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 $\bar{\sigma}$ to be a list in which $\bar{\sigma}(i) = t(\sigma(i))$. (That is, each object is replaced with its type.)

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

For a preference profile $\pi \in \Pi$, define $f(\pi)$ to be the preference profile π' where $\pi'_a = g(\bar{\pi}_a)$ 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 rankwise dominated by the rank distribution of the rank-optimal assignment for π' . 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 on π'_a than $\mu(a)$ has on π_a . Therefore, the rank distribution of μ' rankwise dominates the rank distribution of μ . This completes the proof. $\square$

By the rankwise dominance relation established by the above lemma, the expected average rank of the rank-optimal assignment is weakly higher in M than in M' . Thus, the upper bound on the expected average rank of the rank-optimal assignment for the case of $c = 1$ is a valid upper bound for the general case ($c \geq 1$) as well.

APPENDIX B: PROOFS FROM SECTION 5

PROOF OF PROPOSITION 5.1. For a preference profile π and an assignment μ , let $R_\pi(\mu)$ denote the sum of the agents' ranks in the assignment μ , where the ranks are defined with respect to π .

We first construct the following market:

- i. There are n agents, namely agents $1, \dots, n$. Also, there are m objects, namely $o_1, \dots, o_m$. All of the objects $o_1, \dots, o_m$ have capacity c . The preference order of agent i is denoted by $\prec_i$.
- ii. The preference orders of agents $1, \dots, 4$ are defined by

$$o_1 \prec_1 o_2 \prec_1 o_3 \prec_1 o_4 \prec_1 o_5 \prec_1 \dots \prec_1 o_m,$$

$$o_2 \prec_2 o_3 \prec_2 o_1 \prec_2 o_4 \prec_2 o_5 \prec_2 \dots \prec_2 o_m,$$

$$o_3 \prec_3 o_2 \prec_3 o_4 \prec_3 o_1 \prec_3 o_5 \prec_3 \dots \prec_3 o_m,$$

$$o_3 \prec_4 o_4 \prec_4 o_1 \prec_4 o_2 \prec_4 o_5 \prec_4 \dots \prec_4 o_m.$$

- iii. Agents $S_1 = \{5, \dots, c+3\}$ rank object o_1 first and object o_4 last.³³

³³The agents' preferences are otherwise arbitrary. For example, $o_1, o_2, o_3, o_5, \dots, o_m, o_4$ could be such a preference order, listing objects from the most to the least preferred.

- iv. Agents in the set $S_2 = \{c + 4, \dots, 2c + 2\}$ rank object o_2 first and object o_4 last.
- v. Agents in the set $S_3 = \{2c + 3, \dots, 3c + 1\}$ rank object o_3 first and object o_4 last.
- vi. Agents in the set $S_4 = \{3c + 2, \dots, 4c\}$ rank object o_4 first.
- vii. For every integer $i \in [5, m]$, every agent in the set $S_i = \{c(i - 1) + 1, \dots, ci\}$ ranks object o_i first and ranks the rest of the objects in an arbitrary way in the remaining positions on her preference list such that objects o_1, o_2, o_3, o_4 are ranked as her least favorite objects, with o_4 being ranked last.

Denote the resulting preference profile by π . We first observe that there exists an assignment μ^* with $R_\pi(\mu^*) = n + 1$: μ^* is defined by $\mu^*(j) = o_j$ for $j \in [4]$, and $\mu^*(a) = o_k$ for every agent $a \in S_k$ and $k \in [m]$. In fact, μ^* is the unique assignment satisfying $R_\pi(\mu^*) \leq n + 1$. This holds because o_4 has capacity c , but only $c - 1$ agents (i.e., those in S_4) rank o_4 first. On the other hand, the only agent not in S_4 who ranks o_4 second is agent 4. This uniquely determines μ^* as specified above.

We next suppose that agent 4 deviates by reporting

$$o_3 \prec_4 o_2 \prec_4 o_1 \prec_4 o_4 \prec_4 o_5 \prec_4 \dots \prec_4 o_n.$$

We will show that, under this deviation, agent 4 will be assigned to object o_3 , her top-ranked object.

Let the preference profile resulting from the deviation of agent 4 be denoted by $\bar{\pi}$. There is no assignment μ' with $R_{\bar{\pi}}(\mu') < n + 2$. This holds because precisely c agents should be assigned to o_4 . Thus, at least one agent not in S_4 would be assigned to o_4 . Note, however, that any such agent would rank o_4 third or worse in $\bar{\pi}$. Therefore, $R_{\bar{\pi}}(\mu') \geq n + 2$ for every assignment μ' .

We next show that there exists a unique assignment μ with $R_{\bar{\pi}}(\mu) = n + 2$. Recall that at least one agent not in S_4 would be assigned to o_4 . The only agent not in S_4 who ranks o_4 third or better is agent 3, who ranks o_4 third. Therefore, $\mu(3) = o_4$ must hold if $R_{\bar{\pi}}(\mu) = n + 2$. In addition, every agent other than 3 should be assigned to her top choice in such μ . This uniquely determines the assignment μ as follows: $\mu(1) = o_1$, $\mu(2) = o_2$, $\mu(3) = o_4$, $\mu(4) = o_3$, and $\mu(a) = o_k$ for every agent $a \in S_k$ and $k \in [m]$. We thus have shown that μ is the unique rank-optimal assignment in $\bar{\pi}$.

Finally, we note that $\mu^*(4) \prec_4 \mu(4)$. Therefore, agent 4 can strictly improve her assigned object by misreporting her preference list. This completes the proof. $\square$

PROOF OF PROPOSITION 5.3. Consider an arbitrary agent $a \in A$. For every strict preference order σ over O and every $r \in [m]$, let $p_\sigma(r)$ denote the probability that agent a is assigned to the object ranked r th on σ if she reports σ to $\mathcal{M}^*$. (We note this probability is computed by taking into account the randomization used by the mechanism as well as the randomization of π_{-a} .) Since the agents' preference lists are drawn iid and uniformly at random, then $p_\sigma(r) = p_{\sigma'}(r)$ for every preference list σ' . We therefore use $p(r)$ to denote the latter quantity, with slight abuse of notation.

We next show that $p(r)$ is decreasing in r . Suppose this is not the case. Then there exists s such that $p(s) < p(s + 1)$. Consider the mechanism $\mathcal{M}'$ that, after receiving the

agents' preferences, (i) for every agent it swaps the objects at positions s and $s + 1$ on her preference list and then (ii) runs $\mathcal{M}^*$ on the resulting preference profile. Let the random variables μ^* and μ' , respectively, denote the assignments generated by $\mathcal{M}^*$ and $\mathcal{M}'$. Since $p(s) < p(s + 1)$, then the average rank of μ' is smaller than the average rank of μ , in expectation. This, however, contradicts the fact that $\mathcal{M}^*$ selects only rank-optimal assignments. Thus, $p(s)$ is decreasing in s . This implies that no agent can benefit from misreporting her preference list, in the sense of (5.1). $\square$

PROOF OF PROPOSITION 5.2. Consider the last m agents who choose (i.e., the m agents with lowest priority numbers). Let them be indexed by $a_0, \dots, a_{m-1}$, 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 expected 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 running RSD on a market with m agents, namely $a'_0, \dots, a'_{m-1}$ 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 first. 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 implies that $R \geq R'$. That would complete the proof as it shows that the expected average rank in the original problem is at least $R' \cdot \frac{m}{n} \geq \frac{\ln m - 1}{c}$.

CLAIM B.1. For any $i \in \{0, \dots, m - 1\}$, $R_i \geq R'_i$.

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 not run out of capacity; therefore, agent a_i has at most $m - i$ possible choices. This implies that the rank distribution for agent a'_i rankwise dominates the rank distribution for agent a_i , which implies that $R_i \geq R'_i$. $\square$

$\square$

REFERENCES

- Abdulkadirouglu, Atila, Parag A. Pathak, and Alvin E. Roth (2005), "The New York City High School match." *American Economic Review*, 95, 364–367. [25, 28]
- Abdulkadirouglu, Atila and Tayfun Sönmez (1998), "Random serial dictatorship and the core from random endowments in house allocation problems." *Econometrica*, 66, 689–701. [26, 45]
- Ashlagi, Itai, Yash Kanoria, and Jacob D. Leshno (2017), "Unbalanced random matching markets: The stark effect of competition." *Journal of Political Economy*, 125, 69–98. [26, 27, 28]
- Ashlagi, Itai, Afshin Nikzad, and Assaf Romm (2019), "Assigning more students to their top choices: A comparison of tie-breaking rules." *Games and Economic Behavior*, 115, 167–187. <https://doi.org/10.1016/j.geb.2019.02.015>. [28]

- Ashlagi, Itai and Afshin Nikzad (2020), “What matters in school choice tie-breaking? How competition guides design.” *Journal of Economic Theory*, 190, 105120. <https://doi.org/10.1016/j.jet.2020.105120>. [28, 45]
- Bade, Sophie (2020), “Random serial dictatorship: The one and only.” *Mathematics of Operations Research*, 45, 353–368. <https://doi.org/10.1287/moor.2019.0987>. [45]
- Bruss, F. T. (2000), “Sum the odds to one and stop.” *Annals of Probability*, 28, 1384–1391. <https://doi.org/10.1214/aop/1019160340>. [28]
- Budish, Eric and Estelle Cantillon (2012), “The multi-unit assignment problem: Theory and evidence from course allocation at Harvard.” *American Economic Review*, 102, 2237–2271. [25, 26]
- Che, Yeon-Koo and Olivier Tercieux (2018), “Payoff equivalence of efficient mechanisms in large matching markets.” *Theoretical Economics*, 13, 239–271. <https://doi.org/10.3982/TE2793>. [27]
- Chun, Young H. and Robert T. Sumichrast (2006), “A rank-based approach to the sequential selection and assignment problem.” *European Journal of Operational Research*, 174, 1338–1344. <https://doi.org/10.1016/j.ejor.2005.03.040>. [28]
- Das, Shagnik (2016), “A brief note on estimates of binomial coefficients.” <http://page.mi.fu-berlin.de/shagnik/notes/binomials.pdf>. [40]
- De Haan, Monique, Pieter A. Gautier, Hessel Oosterbeek, and Bas Van der Klaauw (2015), “The performance of school assignment mechanisms in practice.” CEPR Discussion Paper No. DP10656. [25]
- Featherstone, Clayton R. (2020), “Rank efficiency: Modeling a common policymaker objective.” Unpublished paper, The Wharton School, University of Pennsylvania. [25, 28, 45]
- Fortuin, C. M., P. W. Kasteleyn, and J. Ginibre (1971), “Correlation inequalities on some partially ordered sets.” *Communications in Mathematical Physics*, 22, 89–103. <https://doi.org/10.1007/BF01651330>. [25, 26, 28, 29]
- Fountoulakis, Nikolaos and Konstantinos Panagiotou (2012), “Sharp load thresholds for cuckoo hashing.” *Random Structures & Algorithms*, 41, 306–333. [27, 31]
- Frieze, Alan and Michal Karonski (2012), *Introduction to Random Graphs*. <https://www.math.cmu.edu/~af1p/BOOK.pdf>. [31]
- Frieze, Alan and Páll Melsted (2009), “Maximum matchings in random bipartite graphs and the space utilization of Cuckoo Hash tables.” *Random Structures & Algorithms*, 41, 334–364. <https://doi.org/10.1002/rsa.20427>. [27]
- Knuth, Donald E. (1996), “An exact analysis of stable allocation.” *Journal of Algorithms*, 20, 431–442. <https://doi.org/10.1006/jagm.1996.0020>. [26, 27, 45, 53]
- Lee, SangMok and Leeat Yariv (2018), “On the efficiency of stable matchings in large markets.” Working paper. <https://arxiv.org/abs/1512.07276>. [27]

Linusson, Svante and Johan Wästlund (2004), “A proof of Parisi’s conjecture on the random assignment problem.” *Probability Theory and Related Fields*, 128, 419–440. <https://doi.org/10.1007/s00440-003-0308-9>. [28]

Nair, Chandra, Balaji Prabhakar, and Mayank Sharma (2005), “Proofs of the Parisi and Coppersmith–Sorkin random assignment conjectures.” *Random Structures & Algorithms*, 27, 413–444. <https://doi.org/10.1002/rsa.20084>. [28]

NRMP (2020), “Charting outcomes in the match: Senior students of U.S. MD medical schools.” <https://therideshareguy.com/what-its-like-to-be-part-of-lyfts-new-prime-time-experiment/>. [25]

Pittel, Boris (1989), “The average number of stable matchings.” *SIAM Journal on Discrete Mathematics*, 2, 530–549. [26]

Pittel, Boris (1992), “On likely solutions of a stable marriage problem.” *The Annals of Applied Probability*, 2, 358–401. [28]

Shapley, Lloyd and Herbert Scarf (1974), “On cores and indivisibility.” *Journal of mathematical economics*, 1, 23–37. [27]

Walkup, David W. (1980), “Matchings in random regular bipartite digraphs.” *Discrete Mathematics*, 31, 59–64. [https://doi.org/10.1016/0012-365X\(80\)90172-7](https://doi.org/10.1016/0012-365X(80)90172-7). [26, 28, 31, 32, 41]

West, Douglas B. (2000), *Introduction to Graph Theory*, second edition. Prentice Hall. [34]

Co-editor Federico Echenique handled this manuscript.

Manuscript received 24 February, 2020; final version accepted 13 March, 2021; available online 20 April, 2021.