Rank Guarantees in Matching Markets

Yiying Tan*

March 9, 2025

Abstract

Centralized matching markets involve many participants operating under limited information about others' preferences and priority rankings. This paper adopts a robust approach to imperfect information, analyzing worst-case outcomes for applicants in the deferred acceptance mechanism. In a setting where applicants' preferences are IID uniform, I characterize the probability that an applicant can secure a guaranteed outcome, independent of positions' priority over applicants. The probability that an applicant can guarantee his k^{th} favorite position declines exponentially with both k and market size, leading to the asymptotic failure of percentile guarantees. These results highlight potential limits on applicants' participation incentives. The proof connects weak guarantees to the prevalence of perfect matchings in random markets. I further demonstrate that correlated applicant preference rankings hamper guarantees, while anti-correlated rankings improve guarantees. Finally, I explore how constraining positions' priority structures may enhance outcomes.

*Princeton University. I am especially grateful to my advisor, Professor Leeat Yariv, for her invaluable guidance on this project. I am also thankful for helpful comments from Fedor Sandomirskiy, Mira Frick, Andrew Ferdowsian, Clayton Thompson, participants of the 35^{th} Stony Brook International Conference on Game Theory and participants of the 2023/24 Microeconomic Theory Student Workshop at Princeton University. All errors are my own.

1 Introduction

Many large-scale centralized clearinghouses, such as the National Resident Matching Program (NRMP) and numerous school choice systems across cities like New York, Chicago, and Denver, employ the applicant-proposing deferred acceptance (DA) algorithm to assign individuals to institutions. Given the scale and complexity of these markets, it is unrealistic to assume that participants possess complete knowledge of everyone else's preferences or the specific priorities used to assign spots. However, much of the matching literature has thus far operated under a complete information benchmark, where agents are assumed to know all relevant details about the market.

This paper adopts a robust approach to imperfect information, asking the following: When does DA guarantee participants good outcomes despite uncertainty? Specifically, how do worst-case outcomes vary depending on how applicants' preferences are distributed? To analyze this, I consider an environment where applicants—such as junior doctors selecting hospitals or students ranking schools—know their own preferences but face uncertainty about both other applicants' preferences and the priorities assigned to spots. I begin with a model in which applicant preferences are independently and uniformly distributed, while priorities remain adversarially determined.

The central result in Section 3 characterizes the likelihood that any given applicant can guarantee their k^{th} favorite outcome, irrespective of how priorities are assigned. The implications are stark. The probability of securing a guaranteed k^{th} choice declines exponentially with k, meaning that meaningful guarantees are exceedingly rare. Furthermore, asymptotically, no applicant can be assured an outcome above a fixed percentile in the ranking distribution. This suggests that applicants' own preference rankings provide strong guarantees in only a vanishingly small fraction of markets.

I also demonstrate that the likelihood of securing one's k^{th} favorite outcome closely mirrors the probability that each of the lower-ranked outcomes (ranked $(k+1)^{th}$ and beyond) is someone else's top choice. In doing so, I highlight a link between the probability of any rank guarantee and Stirling Numbers of the Second Kind.

Correlated Preferences and the Erosion of Guarantees: In many real-world applications, applicant preferences are not independent. Doctors may rely on widely known hospital rankings, and students often prefer schools ranked highly in public reports. In Section 4, I extend the analysis to correlated preferences and find that the negative results

from the independent case are further exacerbated. When applicants broadly agree on which spots are least desirable, the lowest-priority applicants face significantly worse outcomes. Applicants with low priority are more likely to be left with their least-preferred spots, as others with higher priority systematically occupy the more desirable ones.

However, not all correlations are harmful. In Section 5, I introduce a formal notion of "anti-correlation", where preferences are sufficiently dispersed to avoid extreme concentration at the top spots. I identify a class of preferences where all applicants can be guaranteed good outcomes, showing that a particular degree of anti-correlation is necessary for every applicant to secure at least their second favorite option. Such guarantees are necessary for DA to be obviously strategy-proof (OSP; Li, 2017) so that participants can easily recognize their dominant strategies.

Policy Implications and Participation Decisions: Beyond its theoretical interest, this paper's findings have direct implications for policy and market design. The central message highlights that in many applications, applicants are likely to have poor outcome guarantees. An interpretation of this result is that in many markets, low-priority applicants will do very poorly. To the extent that market designers are concerned with ex-post equality in applicant welfare, it is important to achieve good worst-case guarantees for all participants. The discussion in Section 5 suggests that where feasible, making appropriate priority restrictions advances this objective. If not, policy practitioners must focus on reducing underlying disparities in the quality of spots, which may reduce correlations in preference rankings and ameliorate the prospect of being matched with the lowest-ranked spots.

In addition, the prevalence of poor rank guarantees suggests that applicants with access to good outside options may opt out of participation in centralized mechanisms. For instance, families may apply directly to private schools or prospective applicants may elect to make deals outside the centralized mechanism. This is especially salient if prospective applicants are ambiguity averse.[1] When participation rates are of concern, it may be consequential for policy practitioners to reassure participants about priorities.[2]

[1]This is reasonable in many applications where DA is used, such as residency matching, school choice, or public housing assignment, where outcomes are consequential yet hard to change ex-post assignment. Gilboa and Marinacci (2016) highlight that ambiguity-averse behavior may be sensible when agents have to make weighty decisions in an environment where choosing a particular probability measure over events may seem arbitrary.

[2]I also provide an argument in the Appendix that managing ambiguity over priorities is more important than managing ambiguity about other applicants' preferences.

Altogether, my results highlight a fundamental tension in large-scale matching markets: even though DA is strategy-proof, it does not always ensure acceptable outcomes for all participants, particularly those with low priority. Understanding the distribution of worst-case guarantees is thus crucial for designing and implementing mechanisms that not only match applicants optimally but also maintain their confidence in the system.

Related Literature

A growing literature presents evidence that participants in strategy-proof mechanisms, such as DA, do not in fact behave truthfully. (Ashlagi and Gonczarowski, 2018; Chen and Sönmez, 2006; Hassidim et al., 2016; Rees-Jones, 2018). Ashlagi and Gonczarowski (2018) show that strategy-proof mechanisms are not OSP (Li, 2017), which may partially account for this pattern. This paper quantifies the extent to which DA is not OSP, by showing that applicants' preferences must satisfy very particular conditions for DA to be OSP. A connected strand of literature characterizes when mechanisms are OSP implementable (Ashlagi and Gonczarowski, 2018; Thomas, 2021; Troyan, 2019). DA is OSP implementable when priorities over applicants are acyclic,[3] via a dynamic version of DA (Ashlagi and Gonczarowski, 2018; Thomas, 2021). The results in this paper reinforce the importance of dynamic implementation for OSP: the probability that we have OSP without dynamic implementation is infinitesimal.

The insights in this paper provide a complementary perspective to work discussing matching market outcomes when participants have imperfect information (e.g., Chen and Hu, 2023; Fernandez, 2024; Fernandez et al., 2022; Liu et al., 2014). Much of this literature specifies applicant beliefs over preferences on either side of the market, that map onto beliefs over potential outcomes. The focus is often on characterizing agents' optimal strategies given beliefs. In contrast, the applicants in this paper are akin to ambiguity-averse agents who take a robust approach to imperfect information: Here, applicants' evaluation of their potential outcomes does not require the formation of a prior over possible priorities. This may be a practical approach in applications where it is hard for agents to form defined beliefs, such as in large and complex markets. The results here imply that applicants may be disinclined to participate in the first place, and may not find it optimal to report truthfully.

This paper also connects to the literature on the design of priorities and tie-breaking

[3] An example of acyclic priorities is aligned priorities.

rules (e.g., Abdulkadiroğlu et al., 2009, 2017; Ashlagi et al., 2019; Pycia and Troyan, 2023). The receiving side in DA mechanisms, such as schools, often have coarse preferences over applicants and rely on tie-breaking lotteries. The design of such lotteries is a choice of the distribution over spots' priorities. Section 5 is informative about how the support over priorities influences worst-case outcomes. While global lotteries may potentially allow more applicants to achieve their first choice, we know that worst-case outcomes under global lotteries may be inferior to systematically breaking ties in favor of different applicants at different schools.

2 Model

A *market* is a quadruple $\mathcal{M} = (\mathcal{A}, \mathcal{S}, \succ_{\mathcal{A}}, \succ_{\mathcal{S}})$ where $\mathcal{A}$ is a finite set of *applicants*, and $\mathcal{S}$ is a finite set of *spots*. Each applicant a has a complete ordinal ranking of the set of spots $\mathcal{S}$, assumed to be strict.[4] I denote this ranking by $\succ_a$. Let $\mathcal{P}_{\mathcal{S}}$ denote the set of all possible ordinal rankings over spots, so, $\succ_a \in \mathcal{P}_{\mathcal{S}}$. Likewise, each spot s has a complete strict ordinal ranking of the set of applicants $\mathcal{A}$, denoted $\succ_s$, where $\succ_s \in \mathcal{P}_{\mathcal{A}}$, the set of all possible ordinal rankings over applicants. $\succ_{\mathcal{A}} = \{\succ_a\}_{a\in\mathcal{A}} \in \prod_{a\in\mathcal{A}} \mathcal{P}_{\mathcal{S}}$ and $\succ_{\mathcal{S}} = \{\succ_s\}_{s\in\mathcal{S}} \in \prod_{s\in\mathcal{S}} \mathcal{P}_{\mathcal{S}}$ denote the profile of applicant *preference* rankings and profile of spot *priority* rankings respectively. I assume throughout the paper that we have a balanced market and all participants are mutually acceptable.

Let $\Delta(\prod_{a\in\mathcal{A}} \mathcal{P}_{\mathcal{S}})$ denote the set of joint probability distributions defined over the joint profile of all applicants' ordinal preference rankings. The distribution of applicants' ordinal preference rankings is denoted by $\psi_{\mathcal{A}} \in \Delta(\prod_{a\in\mathcal{A}} \mathcal{P}_{\mathcal{S}})$, that is, $\succ_{\mathcal{A}} \overset{d}{\sim} \psi_{\mathcal{A}}$. In this paper, we are interested in characterizing the distribution of worst-case outcomes under applicant-proposing deferred acceptance (DA), given fixed $\psi_{\mathcal{A}}$, while being robust against $\succ_{\mathcal{S}} \in \prod_{s\in\mathcal{S}} \mathcal{P}_{\mathcal{S}}$.

We start by focusing on a setting where applicants' ordinal preference rankings are independently and identically distributed (IID) according to the uniform distribution over $\mathcal{P}_{\mathcal{S}}$ (e.g., in Ashlagi et al., 2017, 2019; Immorlica and Mahdian, 2015; Kojima and Pathak, 2009; and Che and Tercieux, 2019)). That is, $\psi_{\mathcal{A}} = \prod_{a\in\mathcal{A}} \frac{1}{|\mathcal{P}_{\mathcal{S}}|}$. IID uniform ordinal rankings

[4] In practical implementations of centralized matching, preference rankings are de-facto strict (e.g., application forms generally do not allow indifference to be expressed, and ties in priority status are arbitrarily broken).

imply that applicants, as well as spots, are ex-ante identical.[5] The probability of a good outcome guarantee, obtained in an IID uniform setting, can be interpreted literally as the proportion of markets where applicants' preference rankings generate good outcome guarantees. Hence, it provides a logical starting point and sheds light on the prevalence of markets that yield good guarantees. In Section 4, we allow for correlated preferences among applicants and show that results become even stronger: rank guarantees are worse.

Before proceeding to the main analysis, I provide some standard definitions that I will refer to throughout the paper.

The *rank function* $\text{Rank}_{\succ_a}(s) = |\{s' \in \mathcal{S} \ : \ s' \succeq_a s\}|$ describes the rank of spot s according to the preference $\succ_a$ of applicant a. In particular, a spot s with a $\text{Rank}_{\succ_a}(s) = 1$ is applicant a's favorite spot, while a spot s' with $\text{Rank}_{\succ_a}(s) = 2$ is applicant a's second favorite spot, and so on. Symmetrically, let $\text{Rank}_{\succ_s}(a) = |\{a' \in \mathcal{A} \ : \ a' \succeq_s a\}|$ denote the rank of applicant a in the priority list of spot s.

A *matching* is a one-to-one function that (i) maps each applicant to either a spot or himself; (ii) maps each spot to either an applicant or itself; and (iii) is of order two. Formally, a matching is denoted by $\mu : \mathcal{A} \cup \mathcal{S} \mapsto \mathcal{A} \cup \mathcal{S}$, with the properties that for all $a \in \mathcal{A}$ and $s \in \mathcal{S}$, we have (i) $\mu(a) \in \mathcal{S} \cup \{a\}$; (ii) $\mu(s) \in \mathcal{A} \cup \{s\}$; and (iii) $\mu(a) = s$ if and only if $\mu(s) = a$. Let $\mathcal{M}$ denote the set of all possible (one-to-one) matchings.

We have a *perfect matching* when every applicant and spot is matched, that is, if $\mu(a) = s \in \mathcal{S}$ for all applicants $a \in \mathcal{A}$. We have an *individually rational perfect matching* if we have a perfect matching and for each participant $i \in \mathcal{A} \cup \mathcal{S}$, $\mu_i(i) \succ_i \{i\}$. For a given matching μ, a pair (a, s) is said to form a *blocking pair* if applicant a and spot s are not matched to one another at μ, but prefer each other, i.e., $a \succ_s \mu(s)$ and $s \succ_a \mu(a)$. A matching μ is *stable* for a market $(\mathcal{A}, \mathcal{S}, \succ)$ if there are *no blocking pairs*.[6]

A stable matching can be obtained for any one-to-one matching market by running the *deferred-acceptance algorithm*, as introduced in Gale and Shapley (1962). The following

[5] Ordinal rankings could be induced by random match utilities or surpluses that could be generated from any arbitrary (independent and identical) continuous distributions. For instance, we can generate IID uniform ordinal rankings when underlying cardinal utilities are exponentially distributed according to the CDF $G(u_{ij}) = 1 - \exp\{-u_{ij}$ where u_{ij} is the utility applicant i places on the spot j.

[6] In general, μ is *stable* if there are no blocking pairs *and* it is also *individually rational*: i.e., $\mu(a) = s$ only if $s \succeq_a a$ and $a \succeq_s s$. Any matching trivially satisfies individual rationality, since I assumed that all potential match partners are acceptable.

is an informal description of the applicant-proposing DA (DA):[7]

- Step 1: Every applicant applies to their most preferred spot. Every spot tentatively holds on to the application from the highest-ranked applicant among all applications that have been received (if any) and rejects the rest.
- Step t: Every applicant who has been rejected in step $t-1$ applies to their most preferred spot that has not yet rejected them. Every spot tentatively holds the application from the highest-ranked applicant among all applications that have been received up to period t (if any) and rejects the rest.
- Step T (Terminal Step): When no applicant has been rejected in step $T-1$, no new applications arrive and all tentatively held applications are permanently accepted.

A *centralized matching mechanism* is a mapping from the profile of reported ordinal rankings to a matching. DA induces a centralized matching mechanism. Formally, $DA : \mathcal{P} \mapsto \mathcal{M}$ as described above. Let $\mu_{DA}(\cdot \mid \succ)$ denote the matching under DA given reported rankings $\succ$. The stable matching found by running DA is the *applicant-optimal stable matching* (AOSM); that is, every applicant prefers their match outcome under DA to their outcome under any other *stable* matching.

Now, I provide some definitions I use throughout this paper that are less standard.

Definition (Worst-Case Rank). Suppose the profile of reported applicant preferences is $\succ_{\mathcal{A}}$. Applicant a's *worst-case rank*, denoted $wr_{\succ_{\mathcal{A}}}(a)$, is the rank of the worst stable outcome for applicant a over all possible priority rankings $\succ_{\mathcal{S}}$.

$$wr_{\succ_{\mathcal{A}}}(a) \equiv \max_{\succ'_{\mathcal{S}}} \left\{ \text{Rank}_{\succ_a} \left(\mu_{DA}(a \mid \succ_{\mathcal{A}}, \succ'_{\mathcal{S}}) \right) \right\}$$

Definition (Worst-Case Feasible Set). Suppose the profile of reported applicant preferences is $\succ_{\mathcal{A}}$. Applicant a's *worst-case feasible set*, denoted $W_{\succ_{\mathcal{A}}}(a)$, is the set of match outcomes applicant a can achieve under DA when they have the worst possible priority across all spots. That is,

$$W_{\succ_{\mathcal{A}}}(a) = \left\{ \mu_{DA} \left(a \,\middle|\, (\succ_{\mathcal{A}}, \succ'_{\mathcal{S}}) \right) : \succ'_{\mathcal{S}} \text{ s.t. } \text{Rank}_{\succ'_s}(a) = n \text{ for all } s \in \mathcal{S} \right\}$$

[7] The spot-proposing DA is analogously described by simply swapping the roles of the applicants and spots in the description above.

Observe that $wr_{\succ_{\mathcal{A}}}(a) \in W_{\succ_{\mathcal{A}}}(a)$. If spot s is the worst-case outcome for applicant α given $\succ_{\mathcal{A}}$, then it is an element of $W_{\succ_{\mathcal{A}}}(a)$, when applicant α has the lowest priority across all spots.

Yet, $W_{\succ_{\mathcal{A}}}(a)$ is in general *not* the singleton set $\{wr_{\succ_{\mathcal{A}}}(a)\}$. Also, the worst-case outcome may be obtained even when applicant α has strong priority at some spots.[8] It is important to note that a strong driver of bad outcomes for applicant α is the relative ranking spots have over other applicants. Keeping fixed applicant a's priority rank at each spot, the ranks of other applicants matter and generate different outcomes.

Definition (Top-k Rank Guarantee). Applicant a has a *top-k rank guarantee* when the least-preferred element in their worst-case feasible set has a rank of k^{th} or better, according to their preferences $\succ_a$. That is,

$$\max_{s \in W_{\succ_{\mathcal{A}}}(a)} \{\text{Rank}_{\succ_a}(s)\} \leq k$$

Finally, we can formally define the object of key interest in this paper, the distribution of rank guarantees for an individual applicant a, given the distribution of applicants' preference rankings.

Definition (Distribution of Rank Guarantees). Let the distribution of preferences be $\psi_{\mathcal{A}}$. The distribution of applicant a's worst-case ranks is characterized by the conditional cumulative density function:

$$F_{\psi_{\mathcal{A}}}(j; \succ_a) = \psi_{\mathcal{A}}\left(\left\{\succ_{\mathcal{A}} : \ wr_{\succ_A}(a) \leq j\right\} \middle| \succ_a\right)$$

3 Meaningful Rank Guarantees are Rare

In this section, I focus on characterizing the distribution of rank guarantees for an individual applicant in a balanced market with IID uniform preference rankings among applicants. I consider an arbitrary applicant α, with a preference ranking $s_1 \succ s_2 \succ \ldots s_n$. Simplifying notation, $\mathbb{P}$(top-k guarantee) denotes the probability that applicant α has a

[8] An example is provided in the Appendix.

guarantee of at least their k^{th}-ranked outcome in this setting, for fixed k.[9] Analogously, $\mathbb{P}$(top-$k(n)$ guarantee) refers to the probability that applicant α can guarantee an outcome ranked above some $k(n)$, such as the rank that corresponds to the 90^{th} percentile, in an IID uniform market of size n.

This section presents the main result: meaningful worst-case guarantees are rare for individual applicants in balanced markets. Specifically, the proportion of applicants' joint preferences that would yield a rank guarantee better than rank k^{th} for applicant α, exponentially decreases to 0 as k decreases. More importantly, the problem is exacerbated significantly as the market size n increases: the proportion of profiles that yield *any* percentile guarantee eventually diminishes to 0 as n grows.

Theorem 1. *For any rank level $k(n)$ such that for sufficiently large n, $k(n)/n \leq r \in (0,1)$,*

$$\lim_{n\to\infty} \mathbb{P}(\text{top-}k(n) \text{ guarantee}) = 0.$$

This result follows directly from the following proposition, characterizing an upper bound for the probability of rank guarantees.

Proposition 1 (Upper bound)**.** *For any fixed rank level k, given market size n,*

$$\mathbb{P}\left(\textit{top-k guarantee}\right) \leq \left\{1 - \gamma(n) \cdot \left(1 - \frac{1}{n}\right)^{n-1}\right\}^{n-k}$$

where $\lim_{n\to\infty} \gamma(n) = 1$.

To get a sense for why Proposition 1 holds, let us start by considering the case where $k = n-1$, i.e., the weakest non-trivial rank guarantee. We have a top-$(n-1)$ guarantee if and only if we can exclude s_n, applicant α's worst-ranked spot, from $W_{\succ_{\mathcal{A}}}(\alpha)$, their worst-case feasible set. Hence, $\mathbb{P}\left(\text{top-}(n-1) \text{ guarantee}\right) = 1 - Prob(\{s_n \in W_{\succ_{\mathcal{A}}}(\alpha)\})$.

We can characterize this probability by identifying a necessary and sufficient condition on applicants' preference rankings $\succ_{\mathcal{A}}$ should satisfy such that $s_n \in W_{\succ_{\mathcal{A}}}(\alpha)$ is true. We can then calculate the frequency of joint preference rankings for which such conditions hold.

[9] i.e., For all $k \in \{1, \cdots n\}$, $\mathbb{P}$(top-k guarantee) corresponds to $F_{\psi_{\mathcal{A}}}(k; \succ_\alpha)$ when $\psi_{\mathcal{A}} = \prod_{a\in\mathcal{A}} \psi_a$, where $\psi_a(\succ_a) = \frac{1}{|\mathcal{P}_S|}$ for all $\succ_a \in \mathcal{P}_S$, for all $a \in \mathcal{A}$.

Lemma 1. *Spot s is in applicant α's worst-case feasible set, that is, $s \in W_{\succ_{\mathcal{A}}}(a)$, if and only if given $\succ_{\mathcal{A}}$, there exists a perfect matching μ such that $\mu(a') \succ_{a'} s$ for all $a' \in \mathcal{A}/\{\alpha\}$.*

Lemma 1 implies that for $s_n \in W_{\succ_{\mathcal{A}}}(\alpha)$ to hold, the existence of a perfect matching where every applicant $a' \in \mathcal{A}/\{\alpha\}$ does better than s_n is both necessary and sufficient. The intuition for necessity is as follows. When is spot s_n feasible for applicant α when applicant α has the worst possible priority? For s_n to be feasible, applicant α should sometimes face no competition for s_n: otherwise, applicant α must be rejected from s_n. So, it must be possible that no other applicant applies to spot s_n under DA, at least for some $\succ_{\mathcal{S}}$. Since we have a balanced market where all spots and applicants are mutually acceptable, all applicants are for any preference realization. If no other applicant applies to s_n, then all other applicants must do better than s_n, which corresponds to the condition in Lemma 1.

To see why the condition in Lemma 1 is sufficient for $s_n \in W_{\succ_{\mathcal{A}}}(a)$, suppose there indeed exists a perfect matching $\tilde{\mu}$ meeting the condition. Consider any $\succ_{\mathcal{S}}$ where each applicant $a' \in \mathcal{A}/\{a\}$ has top-priority at their assigned spot under $\tilde{\mu}$. The outcome $\mu_{DA}(a')$, under applicants' preferences $\succ_{\mathcal{A}}$ and such priorities $\succ_{\mathcal{S}}$ must satisfy $\mu_{DA}(a') \succeq_{a'} \tilde{\mu}(a') \succ_{a'} s_n$ for all $a' \in \mathcal{A}/\{a\}$. Applicant α must be the only applicant to s_n, having been rejected by all other spots.

By Lemma 1, the calculation of $\mathbb{P}(\text{top-}n-1 \text{ guarantee}) = 1 - Prob(\{s_n \in W_{\succ_{\mathcal{A}}}(\alpha)\})$ can be reduced to considering the proportion of applicant preferences such that a perfect matching exists as specified in the lemma. We can utilize known results on perfect matchings to analyze $Prob(\{s_n \in W_{\succ_{\mathcal{A}}}(\alpha)\})$. In particular, Hall's Marriage Theorem provides an interpretation of Lemma 1 and serves as a framework to further break down the expression $Prob(\{s_n \in W_{\succ_{\mathcal{A}}}(\alpha)\})$.

Theorem (Hall's Marriage Theorem (Hall, 1935))**.** *Let $G = ((X, Y), E)$ be a bipartite graph with bipartite sets X and Y and edge set E. For each subset $\tilde{X} \subseteq X$, let $N_G(\tilde{X})$ denote the neighborhood of $\tilde{X}$.*[10] *There exists a X-perfect matching in G*[11] *if and only if*

$$|N_G(\tilde{X})| \geq |\tilde{X}| \quad \textit{for every } \tilde{X} \subset X.$$

[10] i.e., $N_G(\tilde{X}) = \{j \in Y : (i, j) \in E \text{ for some } i \in \tilde{X}\}$.

[11] i.e., There exists a set of disjoint edges covering every element in X.

Roughly speaking, all elements of X can be simultaneously matched in the bipartite graph G if and only if the graph is "sufficiently connected" so that every subset of X has sufficient neighbors. To apply the theorem directly, it is useful to consider the *truncated submarket* $\mathcal{M}_{-(\alpha,s_n)} = (\mathcal{A}/\{\alpha\}, \mathcal{S}/\{s_n\}, \tilde{\succ}_{\mathcal{A}/\{\alpha\}})$. Here, we exclude applicant α, spot s_n, and truncate all other applicants' preferences below (and including) spot s_n. It follows directly that the condition in Lemma 1 is equivalent to the existence of an individually rational perfect matching in $\mathcal{M}_{-(\alpha,s_n)}$. $\mathcal{M}_{-(\alpha,s_n)}$ can be represented as a bipartite graph, where one side consists of applicants $\mathcal{A}/\{\alpha\}$, the other side consists of spots $\mathcal{S}/\{s_n\}$, and $(a', s') \in E$ if and only if $s' \succ_{a'} s_n$; that is, applicant a' prefers spot s' to spot s_n.

Hence, $s_n \in W_{\succ_{\mathcal{A}}}(\alpha)$ if and only if there exists a perfect matching in the bipartite graph corresponding to $\mathcal{M}_{-(\alpha,s_n)}$. Hall's Marriage Theorem directly implies that this perfect matching exists if and only if every subset of applicants in $\tilde{A} \subseteq \mathcal{A}/\{\alpha\}$ are collectively connected to at least $|\tilde{A}|$ spots from $\mathcal{S}/\{s_n\}$. Given our mapping of $\mathcal{M}_{-(\alpha,s_n)}$ to a bipartite graph, the interpretation is that it is necessary and sufficient for every subgroup of applicants to find at least the same number of spots superior to s_n in the aggregate.

Conversely, there does *not* exist a perfect matching, so that $s_n \notin W_{\succ_{\mathcal{A}}}(\alpha)$, if and only if *some* subgroup of applicants $A \subseteq \mathcal{A}/\{\alpha\}$ *unanimously prefer* spot s_n to a sufficiently large group of at least $n - |A|$ spots. Hall's theorem yields a useful heuristic: we can exclude spot s_n as a worst-case outcome if and only if spot s_n is "sufficiently preferred" by other applicants, in a rather specific way.

Our discussion implies that we can write:

$$Prob(\{s_n \in W_{\succ_{\mathcal{A}}}(\alpha)\}) = Prob\left(\cap_{j=1}^{n-1} H(j)\right),$$

where $H(j)$ denotes the event that every subset of j applicants in $\mathcal{A}/\{\alpha\}$ find at least j spots superior to s_n in the aggregate, in the truncated submarket $\mathcal{M}_{-(\alpha,s_n)}$. With simple manipulation, we can write:

$$\begin{aligned}\mathbb{P}(\text{top-}n-1 \text{ guarantee}) &= 1 - Prob(\{s_n \in W_{\succ_{\mathcal{A}}}(\alpha)\}) \\ &= 1 - Prob(H(1)) \cdot Prob(\cap_{j=2}^{n-1} H(j) | H(1))\end{aligned}$$

With IID uniform ordinal rankings, it is straightforward to verify that $Prob(H(1)) = (1 - 1/n)^{n-1}$. $H(1)$ holds when no applicant in $\mathcal{A}/\{\alpha\}$ ranks s_n as their first choice. For each

applicant, this is true with probability $(1-1/n)$, and there are $n-1$ independent applicants. It remains for us to characterize $Prob(\cap_{j=2}^{n-1} H(j)|H(1))$, the probability of an individually rational perfect matching in $\mathcal{M}_{-(\alpha,s_n)}$ when all applicants $a' \in \mathcal{A}/\{\alpha\}$ rank at least one spot s' above s_n.

Lemma 2. *The probability that there exists an individually rational perfect matching in $\mathcal{M}_{-(\alpha,s)}$, conditional on all applicants $a' \in \mathcal{A}/\{\alpha\}$ ranking at least one spot s' above s, approaches 1 as the size of the market gets large. i.e.,*

$$\gamma(n) \equiv Prob(\cap_{j=2}^{n-1} H(j)|H(1)) \to 1$$

In this setting, if every applicant finds at least one spot acceptable in the truncated submarket, we are exceedingly likely to be able to match all applicants.

Intuitively, when ordinal preference rankings are IID uniform, there is significant variation in preferences. In the truncated submarket, the number of spots applicants find acceptable is uniformly distributed. So, the lengths of applicants' truncated preference lists are just as likely to be long as they are to be short. Conditional on the length of their truncated preference list, we again have a uniform distribution over the combination of spots applicants find acceptable. This means that when preference lists are short, they are more likely than not to be comprised of different combinations of spots, with limited overlap. This suggests that finding an individually rational perfect matching in the truncated submarket should be feasible. I include a proof for Lemma 2 in the Appendix.

Lemma 2 provides a direct proof of Proposition 1 for the case of $k = n-1$. Observe that given Lemma 2, we have shown Proposition 1 for the case where $k = n-1$. Generalizing to different values of k relies on observing that the probability that we can exclude some spot s from the applicant α's worst-case feasible set decreases given the condition that we also exclude another spot s'.[12] As a result, the joint probability of excluding spots $s_{k+1}, s_{k+2}, \ldots, s_n$ is bounded above by the product of the marginals $\{Prob(\{s_j \in W_{\succ_{\mathcal{A}}}(\alpha)\})\}_{j=k+1,k+2,\ldots,n}$.

[12] I include a short proof in the Appendix.

3.1 An Approximately Necessary and Sufficient Condition

This subsection highlights the observation that broadly speaking, applicant α has a top-k guarantee approximately only when every one of applicant α's $n-k$ *worst* outcomes are the *best* outcome for some other applicant. The peculiarity of this "requirement" provides a concrete flavor to our prior result of the rareness of meaningful rank guarantees.

Corollary 1. *If every spot $s \in \{s_{k+1}, s_{k+2}, \ldots, s_n\}$ which applicant α ranks $k+1^{th}$ or worse is the first choice of some applicant $a' \in \mathcal{A}/\{\alpha\}$, then applicant α has a top-k rank guarantee.*

From Lemma 1, we know that if applicant α has a top-k guarantee, then for *every* $s \in \{s_{k+1}, s_{k+2}, \ldots, s_n\}$, we should not be able to find a perfect matching where every applicant $a' \in \mathcal{A}/\{\alpha\}$ obtains an outcome they prefer to s. Mechanically, applicant α has a top-k guarantee if and only if spots $s_{k+1}, s_{k+2}, \ldots, s_n$ can be excluded from his worst-case feasible set $W_{\succ_{\mathcal{A}}}(\alpha)$. Corollary 1 very straightforwardly follows. Whenever some other applicant a' ranks some spot s as their first choice, we can never find a perfect matching where *every* $a' \in \mathcal{A}/\{\alpha\}$ does *better* than spot s.

We can demonstrate via simulation that Corollary 1 is asymptotically "almost necessary" for applicant α to have a top-k guarantee. Here, I derive a lower bound that exactly captures the probability of the sufficient condition in Corollary 1, and show that it is remarkably tight. In simulations, the lower bound derived here gets very close to the upper bound presented in Proposition 1 even for a market size as small as $n = 9$.

Stirling's numbers of the Second Kind play a role in the characterization of the lower bound. Formally, for any m, the r^{th} Stirling's number is given by:

$$S(m,r) = \sum_{i=0}^{r} \frac{(-1)^{r-i} i^m}{(r-i)!i!}$$

Intuitively, $S(m,r)$ counts the number of ways to partition a set of m objects into r non-empty subsets. They appear in various combinatorial problems and have been applied anywhere from calculating the moments of the Poisson distribution to counting the total number of possible rhyming schemes in poems.

To get a sense of how Stirling's numbers will be used, consider some arbitrarily fixed set $\tilde{S} \subseteq \mathcal{S}$ of $n-k$ spots, and some arbitrarily fixed set $\tilde{A}$ of j applicants, where $\tilde{A} \subseteq \mathcal{A}/\{\alpha\}, |\tilde{A}| = j$.

Notice that the number $S(j, n-k) \cdot (n-k)!$ counts how many ways we can have every applicant $a \in \tilde{A}$ be assigned to some spot in $s \in \tilde{S}$, where each of the $n-k$ spots in $\tilde{S}$ is assigned at least once.[13] So, analogously, $S(n-k, j) \cdot (n-k)!$ corresponds to the number of ways we can have every applicant $a \in \tilde{A}$ rank a spot in $\tilde{S}$ as their first choice, where every spot in $\tilde{S}$ is the first choice of some applicant in $\tilde{A}$.

Proposition 2 (Lower Bound)**.** *For any market size n, for any rank level k, let*

$$LB(n,k) \equiv \frac{\sum_{j=n-k}^{n} \binom{n-1}{j} \cdot S(j,\ n-k) \cdot (n-k)! \cdot k^{n-j-1}}{n^{n-1}}.$$

$LB(n,k)$ provides a lower bound for the probability that an individual applicant has a top-k rank guarantee in a balanced market of size n.

Applying Stirling's numbers, observe that the numerator of $LB(n,k)$ counts the number of other applicants' first choices where every spot applicant α ranks $k+1^{th}$ or below is ranked first by another applicant.[14] The denominator simply counts the total number of permutations of other applicants' first choices. Therefore, $LB(n,k)$ provides us with the probability, given IID uniform preference lists, that $\succ_{\mathcal{A}}$ satisfies a sufficient condition for applicant α to have a top-k guarantee.

[13] The factorial term arises since we may arbitrarily re-label spots.

[14] See Appendix for a detailed explanation of expressions in the numerator.

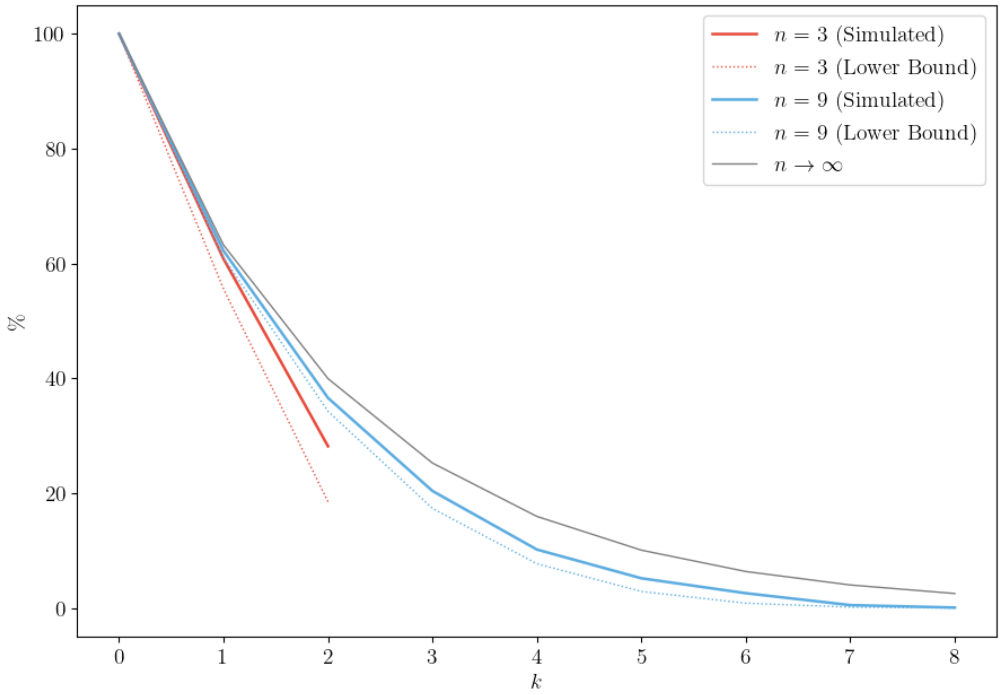


Figure 1: Simulated and analytical probabilities of top $n-k$ guarantees.

While the tightness of the lower bound may initially appear surprising, the reason it is tight stems from the same forces underpinning Lemma 2. Lemma 2 implies that if every applicant $a' \in \mathcal{A}/\{\alpha\}$ prefers at least one other spot s' to whichever spot s we want to exclude from $W_{\succ_A}(\alpha)$, we are almost certainly unable to exclude spot s. This should indicate that the first-order factor of whether or not $s \notin W_{\succ_A}(\alpha)$ is the existence of at least one applicant who finds all spots $s' \in \mathcal{S}/\{s\}$ inferior to spot s.

4 Rank Guarantees with Correlated Preferences

When applicants have correlated preferences, meaningful rank guarantees are even less prevalent. At a high level, we have seen that whenever applicant α has a good rank guarantee, the market must contain a sufficient number of other applicants who like the outcomes he dislikes. Intuitively, other applicants are less likely to have such preferences when preference rankings are more correlated.

In this section, I show that the distribution of rank guarantees when there is more

correlation first-order stochastically dominates the distribution when we have less correlation.[15] Using simulations, I show that this effect can be substantial even when introducing arguably small preference correlations.

Definition (Partial Order of Market Correlation). Compare $\psi_{\mathcal{A}}, \psi'_{\mathcal{A}} \in \Delta(\prod_{a\in\mathcal{A}} \mathcal{P}_{\mathcal{S}})$, with full support, where applicants in $\mathcal{A}$ are exchangeable[16]. A market governed by $\psi_{\mathcal{A}}$ is more correlated than one governed by $\psi'_{\mathcal{A}}$, i.e., $\psi_{\mathcal{A}} \succeq_{corr} \psi'_{\mathcal{A}}$, if and only if:

$$\mathbb{P}_{\psi_{\mathcal{A}}}(s \succ_a s' |\ s \succ_{a'} s', \{\succ_{\tilde{a}}\}_{\tilde{a}\in\mathcal{A}/\{a,a'\}}) \geq \mathbb{P}_{\psi'_{\mathcal{A}}}(s \succ_a s' |\ s \succ_{a'} s', \{\succ_{\tilde{a}}\}_{\tilde{a}\in\mathcal{A}/\{a,a'\}})$$

for all $s, s' \in \mathcal{S}$ and for all $a, a' \in \mathcal{A}$ for all $\{\succ_{\tilde{a}}\}_{\tilde{a}\in\mathcal{A}/\{a,a'\}}$ where such conditional probabilities are defined.

We say that a market where $\succ_{\mathcal{A}} \overset{d}{\sim} \psi_{\mathcal{A}}$ is more correlated than one where $\succ_{\mathcal{A}} \overset{d}{\sim} \psi'_{\mathcal{A}}$ if the following is true: for arbitrary pairs of spots and applicants, we are more confident that $s \succ_a s'$, when we are told that $s \succ_{a'} s'$ under $\psi_{\mathcal{A}}$, compared to under $\psi'_{\mathcal{A}}$.

This definition of correlation captures a natural notion of a "more correlated" joint distribution over preferences with the same marginal distributions over individual preference rankings. As a simple example, consider a market with two applicants a_1, a_2 and two spots s_1, s_2. Fixing the marginal (i.e., fix $\mathbb{P}(s_1 \succ_{a_1} s_2)$), increasing correlation directly corresponds to increasing the covariance between the random variables $X_1 = \mathbb{I}(\{\succ_{\mathcal{A}}: s_1 \succ_{a_1} s_2\})$ and $X_2 = \mathbb{I}(\{\succ_{\mathcal{A}}: s_1 \succ_{a_2} s_2\})$.

The order $\succeq_{corr}$ also captures correlation in a setting where ordinal preference rankings are generated by conditionally independent cardinal utilities, a common notion of preference correlation (e.g., in Agarwal and Somaini, 2018; Kapor et al., 2024). This is the model $u_{ij} = \rho v_j + \varepsilon_{ij}$, $v_j \sim F_j$, $\varepsilon_{ij} \sim G$, where F_j, G are CDFs. One may verify that a higher $|\rho|$ corresponds to more correlation among applicant preference rankings according to $\succeq_{corr}$.[17]

The following proposition shows that more correlation worsens the distribution of rank guarantees:

[15] A first-order stochastically dominant distribution is *worse*: Rank-$(k+1)$ is worse than Rank-k.

[16] The joint distributions $\psi_{\mathcal{A}},, \psi'_{\mathcal{A}}$ are invariant to any relabelling of applicants.

[17] Loosely: Each realization of $\vec{v} = (v_1, \ldots, v_n)$ is associated with some distribution of individual preference rankings ψ (non-degenerate due to ε_{ij}), and we have $\psi_{\mathcal{A}} = \prod_{\in\mathcal{A}} \psi$. Observing $\succ_\alpha$ of applicant α is informative about realized v, which in turn provides information about $\succ_{a'}$ of applicant a'. The higher is $|\rho|$, the more informative v is about $\succ_{a'}$.

Proposition 3. *If $\psi_{\mathcal{A}} \succeq_{corr} \psi'_{\mathcal{A}}$, then the distribution of rank guarantees generated by $\psi_{\mathcal{A}}$ first-order stochastically dominates by the distribution generated by $\psi'_{\mathcal{A}}$. Formally,*

$$\psi_{\mathcal{A}} \succeq_{corr} \psi'_{\mathcal{A}} \implies F_{\psi_{\mathcal{A}}}(\cdot; \succ_{a'}) \succ_{FOSD} F_{\psi'_{\mathcal{A}}}(\cdot; \succ_{a'}) \quad \text{for all } a' \in \mathcal{A}.$$

Broadly, correlated preferences generate poorer rank guarantees for applicant α due to reduced demand for his worst-ranked outcomes. Section 3 tells us that applicant α obtains a top-k guarantee when the spots $s_k, s_{k+1}, \ldots, s_n$ can be ruled out as feasible outcomes. We can rule out $s_k, s_{k+1}, \ldots, s_n$ only if these spots, which applicant α ranks poorly, are ranked relatively well by other applicants. Such a combination of preferences must be less likely when preferences are more correlated.

Since applicants are exchangeable, each joint distribution can be fully characterized by a marginal distribution over individual preference rankings ψ_α, and the conditional probability over unordered collections of other applicants' preferences given each realized individual ranking $\succ_\alpha$.[18] As seen in Section 3, applicant α's worst-case outcome, depends only on what spots can be excluded from his worst-case feasible set. The existence of particular perfect matchings in the market, which in turn only depends on the composition of preference lists, fully determines this. Hence, the distribution of worst-case outcomes given $\succ_\alpha$ is fully determined by the conditional probabilities over collections of other applicants' preference lists. Hence, given some $\succ_\alpha$, the profile of preferences $\{\succ_a\}_{a \in \mathcal{A}/\{\alpha\}}$ yields the same worst-case outcome for applicant α as $\{\succ'_a\}_{a \in \mathcal{A}/\{\alpha\}}$ if both contain the same collection of preference lists (ignoring applicant labels).

We want to show that when there is more correlation, greater weight is placed on collections of preference rankings which correspond to poorer worst-case ranks. The idea of the proof then utilizes two main observations. First, we can incrementally increase the level of correlation, starting from $\psi'_{\mathcal{A}}$ toward the level of $\psi_{\mathcal{A}}$, via a sequence of probabilistic swaps. In particular, at each step, for some fixed collection of preference rankings, swap an adjacent pair of spots in one preference ranking to reduce the number of order reversals relative to another preference ranking.[19] This corresponds to reducing the number of order reversals between applicant α and some applicant a', keeping the preference rankings of

[18] Given some collection of preference rankings, $\succ_{\mathcal{A}}$ corresponding to each permutation assigning preference rankings to different applicants occur with equal probability.

[19] In doing so, we make the same swap across all $\succ_{\mathcal{A}}$ corresponding to such a collection.

all other applicants constant. Second, a collection of preference lists with one fewer rank reversal relative to $\succ_\alpha$, all else constant, necessarily yields a weakly inferior worst-case outcome for applicant α.

The following figure illustrates simulations of rank guarantees under various market sizes, given that applicant preferences are generated according to $u_{ij} = \rho v_j + \varepsilon_{ij}$, $v_j \sim \text{Uniform}\{1,\dots,n\}$, $\varepsilon_{ij} \sim \text{Uniform}\{1,\dots,n\}$.

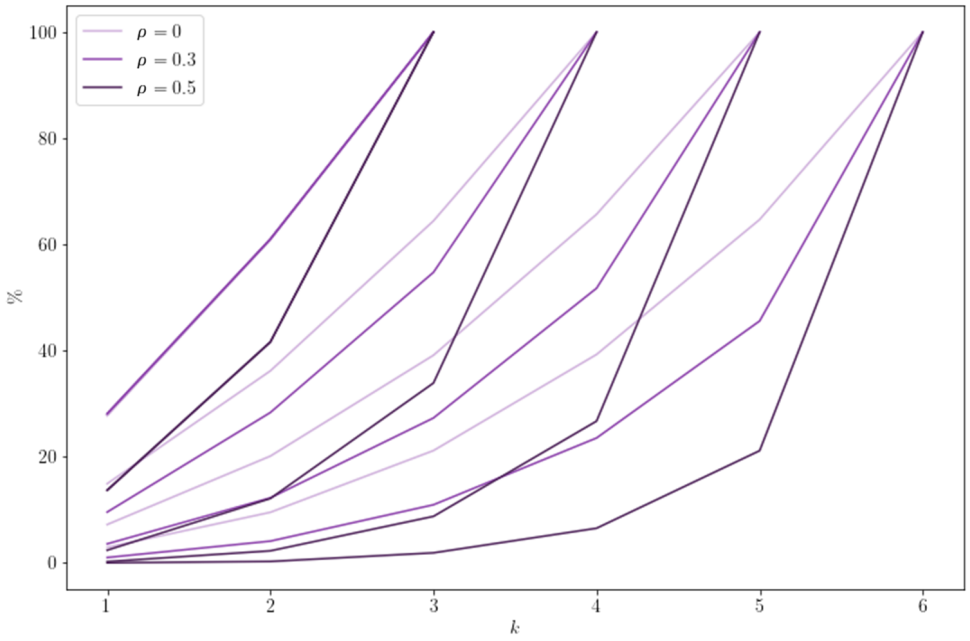


Figure 2: Cumulative Distribution Functions of top-k rank guarantees for different degrees of correlation, for markets of sizes $n = 3, 4, 5$ and 6.

We can see from the figure that even for relatively small levels of correlation, there is a noticeable worsening of the distribution in worst-case outcomes. Holding the size of the market constant, the CDF of rank guarantees under more correlation first-order stochastically dominates that under less correlation. Furthermore, this effect is exacerbated for larger markets. Holding the level of correlation constant, CDF of rank guarantees in larger markets first-order stochastically dominates that in smaller markets; this deterioration in rank guarantees appears to be more dramatic under more correlation.

5 When do Meaningful Guarantees Exist?

This section discusses two approaches to understanding when meaningful rank guarantees exist. First, I define and discuss a notion of "anti-correlation" in applicant preferences. Sufficiently "anti-correlated" preferences generate strong rank guarantees for all applicants. Second, I explore how different ways of restricting the domain of spots' priorities may improve applicants' rank guarantees.

5.1 Anti-Correlated Preferences

The ideas presented in Section 3 and Section 4 suggest that meaningful guarantees are achievable only when applicants' preference rankings are somewhat complementary.[20] Here, I formalize a sense in which applicant preferences are complementary, by defining a partial order of "anti-correlation" in preferences.

Definition (Block-k Anti-Correlation). Applicant preferences $\succ_{\mathcal{A}}$ are block-k correlated if there exists a partition of spots, $\mathcal{B}_{\mathcal{S}} = \{S_1, S_2, \ldots, S_T\}$, and a partition of applicants, $\mathcal{B}_{\mathcal{A}} = \{A_1, A_2, \ldots, A_T\}$, such that for all $t \in \{1, \ldots, T\}$:

- $|A_t| = |S_t| \leq k$.
- If $a \in A_t$, then $b \in B_t,\ b' \in B_j \neq B_t \implies b \succ_a b'$.
- There does not exist an alternative partition where $|A_j| < k$ for all j.

Block-1 anti-correlation is stronger than block-2 anti-correlation, and so on.

Lemma 3. *If $\succ_{\mathcal{A}}$ is block-k anti-correlated, then all applicants are guaranteed at least their top-k outcome.*

Intuitively, block-k anti-correlation breaks down the matching market into smaller "sub-markets" so that applicants can truncate their preference rankings without loss. Each applicant only faces at most k applicants for their k most preferred spots regardless of priority rankings.[21] Worst-case outcomes are alleviated since long rejection chains are circumscribed.

[20] For example, Section 3 demonstrated that an applicant α obtains a top-k rank guarantee approximately only when all his *worst* ranked spots $s_{k+1}, s_{k+2}, \ldots, s_n$ are the *best* outcome for another applicant.

[21] A formal proof is available in the Appendix.

Block anti-correlation is, in fact, necessary for top-1 and top-2 market-wide guarantees. Necessity for top-1 guarantees is extremely intuitive: every applicant needs to have a different first choice. The more interesting result is that for a top-2 guarantee to exist for every applicant, $\succ_{\mathcal{A}}$ must be block-2 anti-correlated.

The question of when a top-2 guarantee exists connects to obvious strategy-proofness (Li, 2017) in DA. For a mechanism to be obviously strategy-proof, each applicant's worst-case outcome by truthfully revealing their preferences must be as good as any feasible outcome when they are not truthful. In this case, applicants should be guaranteed at least their second favorite outcome for it to be "obvious" for applicants to reveal their first choice truthfully. That is, we require a top-2 guarantee for all applicants.

Lemma 4. *Every applicant in $\mathcal{A}$ has a top-2 guarantee if and only if $\succ_{\mathcal{A}}$ is block-2 anti-correlated.*

Lemma 4 implies that we have a top-2 guarantee only when the entire market can be divided into extremely small sub-markets of at most two applicants and two spots. Sufficiency follows immediately from Lemma 3. In the Appendix, I show how a market-wide top-2 rank guarantee means we must be able to construct a partition of market participants corresponding to block-2 anti-correlation.

A possible interpretation of Lemma 4 is that it quantifies the importance of dynamic implementation. Ashlagi and Gonczarowski (2018) show that when priorities are acyclic, we have a dynamic implementation of DA that is obviously strategy-proof. In contrast, if dynamic implementation is impossible, Lemma 4 implies that obvious strategy-proofness requires us to restrict applicant preferences to an extremely particular set. Convincing unsophisticated participants that truthtelling is optimal would require them to believe that applicants' preference rankings are block-2 anti-correlated, which is implausible in most applications.

While I leave a formal quantification for future work, it should not be surprising if block-k anti-correlation is approximately necessary for top-k market-wide guarantees more generally. I outline a loose intuition here. Consider some $\succ_{\mathcal{A}}$ that is block-$(k+1)$ anti-correlated. What is the probability that $\succ_{\mathcal{A}}$ yields a top-k market-wide guarantee? By definition of block-$(k+1)$ anti-correlation, we know there must exist a subset of applicants A_j and corresponding spots S_j which comprise a $(k+1)$-sized balanced market. If a top-k market-wide guarantee holds overall, a top-k market-wide guarantee must hold this

submarket. To the extent that a top-k market-wide guarantee in such a submarket is unlikely,[22] block-k anti-correlation is approximately necessary for a top-k market-wide guarantee.

5.2 Restricting Priorities

Mechanically, limiting the domain of possible outcomes must weakly improve the worst-case outcome. Hence, a natural question is the extent to which reducing the domain of priorities $\succ_{\mathcal{S}}$ can improve the distribution of worst-case ranks. The question of priority restrictions also has a direct practical appeal, given priorities are generally perceived to be more tractable to designers than applicant preferences. In this section, I discuss two practically salient priority restrictions: aligned priorities and the assignment of "safety schools".

5.2.1 Aligned Priorities

Aligned priorities are common in real-world applications: for instance, prospective students and residents may be ordered by quality or test scores, or prospective public housing residents may enter a single lottery across available properties. Limiting possible priority realizations to aligned priorities does not impact the distribution of worst-case ranks.

Lemma 5. *Fix any $\succ_{\mathcal{A}}$. An individual applicant's worst-case rank does not change when we restrict spots to have aligned priorities over applicants. i.e.,*

$$wr_{\succ_{\mathcal{A}}}(a) \equiv \max_{\succ'_{\mathcal{S}} \in \prod_{s\in\mathcal{S}} \mathcal{P}_{\mathcal{A}}} \left\{ Rank_{\succ_a}\left(\mu_{DA}(a|\succ_{\mathcal{A}}, \succ'_{\mathcal{S}})\right)\right\} = \max_{\succ'_{\mathcal{S}} \in \mathcal{D}_{aligned}} \left\{ Rank_{\succ_a}\left(\mu_{DA}(a|\succ_{\mathcal{A}}, \succ'_{\mathcal{S}})\right)\right\}$$

where $\mathcal{D}_{aligned} := \{\succ_{\mathcal{S}} \in \prod_{s\in\mathcal{S}} \mathcal{P}_{\mathcal{A}} \ : \succ_s = \succ_{s'} \quad \forall s, s' \in \mathcal{S}\}$.

The insight underlying Lemma 5 is that if there exists a perfect matching $\tilde{\mu}$ where every $a \in \mathcal{A}/\{\alpha\}$ is matched to some $\tilde{\mu}(a) \succ_a s$, then we can find an aligned priority which implements an even better outcome under DA for every applicant $a \in \mathcal{A}/\{\alpha\}$. Then,

[22] A top-k market-wide guarantee in such a submarket should be unlikely. Section 3 suggests the probability of even an individual top-$(k-1)$ guarantee in an IID uniform setting is limited. In this case, we require a market-wide guarantee across all individuals and additionally rule out meaningful block anti-correlation within the $k+1$ submarket.

Lemma 1 implies that spot s is still a worst-case feasible outcome. I include the constructive proof of Lemma 5 in the Appendix.

5.2.2 Safety Schools

An intuitive way to improve rank guarantees is to assign each applicant a spot where they always have top priority.[23] This can be interpreted as the assignment of "safety schools": when an applicant applies to such a spot, they will be matched to the spot. In school choice applications, we may see this when students are guaranteed admission to their nearest neighborhood school.

In a setting where applicant preferences are IID uniform, such a priority restriction improves outcomes in a first-order stochastic sense.

Corollary 2. *Let $\psi_{\mathcal{A}}$ correspond to the IID uniform distribution, and suppose $\mathcal{R}_{\tilde{\mu}} = \{\succ_{\mathcal{S}} \in \prod_{s\in\mathcal{S}} \mathcal{P}_{\mathcal{A}} : Rank_{\succ_{\underline{\mu}(a)}}(a) = 1 \quad \forall a \in \mathcal{A}\}$ for some matching $\tilde{\mu}$. Then,*

$$\psi_{\mathcal{A}}(\max_{\succ_{\mathcal{S}}\in\mathcal{R}}\{Rank(\mu(\alpha;\succ_{\mathcal{A}}))\} \leq k) \geq \frac{k}{n} > \psi_{\mathcal{A}}(\max_{\succ_{\mathcal{S}}\in\prod_{s\in\mathcal{S}}\mathcal{P}_{\mathcal{A}}}\{Rank(\mu(\alpha;\succ_{\mathcal{A}}))\} \leq k).$$

for all $k \in \{1, 2 \ldots, n\}$.

Each applicant's worst case outcome must be at least as good as the spot at which they were assigned top priority. In an IID uniform setting, each applicant's rank for their assigned spot is uniformly distributed. So, the distribution of each applicant's worst-case rank must be first-order stochastically dominated by the uniform distribution. Following Proposition 1, the uniform distribution is in turn first-order stochastically dominated by the distribution of worst-case ranks without priority restrictions.

At the same time, an important caveat to Corollary 2 is that improvements via such priority restrictions are sensitive to preference correlation. In the extreme where applicant preferences were perfectly aligned, the assignment of safety schools reduces to fixing an aligned priority order over applicants. There is no improvement to the worst-case outcome of the applicant assigned top priority the least preferred spot, relative to no priority restrictions. When preference rankings are significantly correlated, the loss in symmetry

[23] This means restricting the domain of possible preferences to $\{\succ_{\mathcal{S}}: \text{Rank}_{\succ_{\underline{\mu}(a)}}(a) = 1 \quad \forall a \in \mathcal{A}\}$ for some matching $\tilde{\mu}$.

among applicants may be of concern despite improvement to worst-case outcomes for a subset of applicants.

5.2.3 Implications for Tie-Breaking

Should we run a universal lottery, or assign specialized priorities? Universal lotteries, which randomly select an aligned priority order of applicants that applies across all spots, have reported efficiency advantages (e.g., Abdulkadiroğlu et al., 2009, 2017). Yet, Lemma 5, together with results in Section 3, tells us that individual distributions of worst-case ranks under aligned priorities are rather dismal. Meanwhile, Corollary 2 implies that distributions of worst-case ranks improve if applicants have asymmetric tie-breaking probability at different spots especially if applicants and spots are relatively symmetric.

The observations made in this section imply that when we care about both efficiency and applicants' perceptions of their worst-case outcomes, there may be a trade-off. Such a trade-off may be amplified when applicants have idiosyncratic preferences, and less salient of a concern if applicants have highly correlated preferences.

6 Discussion

This paper demonstrates that in many settings, applicants are unlikely to have meaningful rank guarantees under deferred acceptance. Worst-case guarantees are relevant whenever we are concerned about ex-post equality in applicants' outcomes, and when we are concerned about participation incentives. Where markets involve many participants, the proportion of markets where an individual applicant has a guarantee of their k^{th} favorite outcome is approximately the proportion of markets where all of his most undesirable outcomes are the best outcome for someone else in the market. The probability that an individual applicant has a guarantee of achieving an outcome he ranks above any fixed percentile is vanishing. Correlation in preferences exacerbates this result. Together, these insights suggest that efforts to reassure participants about potential outcomes may be warranted in many applications when there is a realistic option to opt-out.

Meaningful market-wide rank guarantees can be achieved when preferences are sufficiently anti-correlated. I further demonstrate that block-2 anti-correlation is necessary and sufficient for a market-wide top-2 rank guarantee, which is in turn necessary for DA to be OSP. A potentially fruitful direction of future research would be to formally quantify the

extent to which block-k anti-correlation is not just sufficient, but approximately necessary for the existence of meaningful rank guarantees, as alluded to in Section 5.1.

Section 5 discussed how priority restrictions may improve rank guarantees. In particular, assigning applicants top priority at different spots (assigning "safety schools"), can improve worst-case outcomes. An interesting question for future research is the extent to which this type of priority restriction yields market-wide improvements. For example, it may be infeasible to assign everyone a good safety school, as evident from the extreme case where applicant preferences are perfectly aligned. When is it possible to assign priorities so that every applicant achieves a top-k guarantee? Observe that this requires the existence of some Pareto-efficient matching where all applicants achieve an outcome of at least rank k. It may be worthwhile to characterize the distribution of markets where the profile of applicant preferences satisfies such a requirement.[24]

Another avenue for future work may be determining how much preference truncation may change the profile of worst-case outcomes. Suppose every applicant. except some applicant α, truncates their preference list: applicant α's worst-case outcome must improve, following the logic outlined in Lemma 1. Yet, if applicant α also truncates their preference list, they may be unmatched. Conditional on being matched, applicants' worst-case outcomes should improve. On the other hand, how frequently might applicants wind up unmatched? To the extent that applicants have limited outside options, we should care about such a trade-off.

Finally, a limitation of the worst-case approach is that it does not allow significant commentary on applicants' optimal reporting strategy. If applicants only care about their worst-case outcomes, then any strategy is optimal as worst-case outcomes only depend on other applicants' reports. At the same time, if applicants care about the reported rank of their worst-case outcome, they may adjust their strategy. A potential question arises: if indeed applicants care about the reported rank of their worst-case outcome, what would equilibrium reports look like, and how would worst-case outcomes change? The answer may allow us to generate predictions that can be compared to reported preferences in real applications.

[24]This would be an upper limit to our ability to achieve market-wide outcome guarantees via such preference restrictions since the issue of eliciting enough information to assign appropriate priorities remains.

A Appendix

Proposition 1

In conclude Proposition 1, we loosely argued that

$$P(s_{n-j} \notin W_{\succ_{\mathcal{A}}}(\alpha) | \cap_{l=1}^{j-1} \{s_{n-l} \notin W_{\succ_{\mathcal{A}}}(\alpha)\}) \leq P(s_{n-j} \notin W_{\succ_{\mathcal{A}}}(\alpha)) = P(s_n \notin W_{\succ_{\mathcal{A}}}(\alpha))$$

for all j. Intuitively, when one spot is not sufficiently preferred (and so excluded from the worst-case set), another spot is more likely to be sufficiently preferred (and so included in the worst-case set). I provide a formal proof here. The proof formalizes this natural notion of negative association.

Suppose we fix some set of applicant preferences $\{\succ_a\}_{a \in \mathcal{A}/\{\tilde{a}\}}$ and fix applicant $\tilde{a}$'s ranking of spots $\mathcal{S}/\{s_n\}$. i.e., we fix how applicant $\tilde{a}$ orders $s_1, s_2, \ldots, s_{n-1}$. Given we are in an IID uniform setting, every such arbitrary restriction on applicants' preferences holds with equal probability. We can partition the domain of joint applicants' preferences into equal sizes according to these restrictions. Within each element of the partition, i.e., conditional on such preferences, applicant $\tilde{a}$ ranks spot s_n at each position with equal probability.

Focus on one particular element of the partition, which I denote as E. For simplicity, I discuss without loss of generality a case where $s_1 \succ_{\tilde{a}} s_2 \succ_{\tilde{a}} \cdots \succ_{\tilde{a}} s_{n-1}$. How should we think about the case where s_n is ranked k as opposed to $k+1$? Crucially, moving s_n one position ahead is equivalent to adding the edge $(\tilde{a}, s_n)$ to the bipartite graph associated with the truncated submarket $\mathcal{M}_{-(\alpha, s_k)}$. Analogously, removing the edge $(\tilde{a}, s_k)$ from the bipartite graph associated with the truncated submarket $\mathcal{M}_{-(\alpha, s_n)}$. Straightforwardly, adding edges to $\mathcal{M}_{-(\alpha, s_k)}$ must make it easier to find a perfect matching, while removing edges from $\mathcal{M}_{-(\alpha, s_n)}$ makes it harder to find a perfect matching.

Hence, we have monotonicity with respect to applicant $\tilde{a}$'s rank of s_n, in the following sense: for all $s \neq s_n$, if $s \in W_{\succ_{\mathcal{A}}}(\alpha)$ when s_n is ranked k^{th} by applicant $\tilde{a}$, we should still have $s \in W_{\succ_{\mathcal{A}}}(\alpha)$ when s_n is ranked j^{th} by applicant $\tilde{a}$, where $j < k$. Similarly, if $s_n \notin W_{\succ_{\mathcal{A}}}(\alpha)$ when s_n is ranked k^{th} by applicant $\tilde{a}$, we should still have $s_n \notin W_{\succ_{\mathcal{A}}}(\alpha)$ when s_n is ranked j^{th} by applicant $\tilde{a}$, where $j < k$.[25] The number of $s \in \mathcal{S}/\{s_n\}$ for which $s \in W_{\succ_{\mathcal{A}}}(\alpha)$ must weakly decrease in applicant $\tilde{a}$'s rank of s_n; at the same time, the incidence of $s_n \in W_{\succ_{\mathcal{A}}}(\alpha)$ must weakly increase in applicant $\tilde{a}$'s rank of s_n. It directly follows that within this

[25] In the extreme, when s_n is ranked 1 by applicant $\tilde{a}$, there are no edges involving applicant $\tilde{a}$ in $\mathcal{M}_{-(\alpha, s_n)}$ and so there is no perfect matching, which implies $s_n \notin W_{\succ_{\mathcal{A}}}(\alpha)$ in that case.

partition element, $P(\{s \in W_{\succ_{\mathcal{A}}}(\alpha) \text{ for all } s \in \tilde{S}\} | s_n \notin W_{\succ_{\mathcal{A}}}(\alpha), E) \geq P(\{s \in W_{\succ_{\mathcal{A}}}(\alpha) \text{ for all } s \in \tilde{S}\} | E)$, for all $\tilde{S} \subseteq \mathcal{S}/\{s_n\}$.

Leveraging the symmetry of the IID uniform setting, straightforward probability manipulations imply that the following holds for arbitrary s':

$$P(\{s \in W_{\succ_{\mathcal{A}}}(\alpha) \text{ for all } s \in \tilde{S}\} | s' \notin W_{\succ_{\mathcal{A}}}(\alpha)) \geq P(\{s \in W_{\succ_{\mathcal{A}}}(\alpha) \text{ for all } s \in \tilde{S}\}), \quad \text{for all } \tilde{S} \subseteq \mathcal{S}/\{s'\}$$

Slight further manipulation yields:

$$P(s' \notin W_{\succ_{\mathcal{A}}}(\alpha) | \{s \notin W_{\succ_{\mathcal{A}}}(\alpha) \text{ for all } s \in \tilde{S}\}) \leq P(s' \notin W_{\succ_{\mathcal{A}}}(\alpha)) \quad \text{for all } \tilde{S} \subseteq \mathcal{S}/\{s'\}$$

Lemma 2

Simplifying notation, let us denote the event that $\cap_{j=1}^{n-1} H(j)$ as PM.

First, we may breakdown the $\gamma(n)$ into a weighted sum of conditional probabilities: $\gamma(n) = \sum_{j=1}^{n} \text{Prob}(PM|\text{min len} = j) \cdot \text{Prob}(\text{min len} = j)$ where $\text{Prob}(\text{min len} = j)$ denotes the probability that the longest preference list in the market is of length j.

Second, observe that the probability that a perfect matching exists increases when applicants have longer preference lists, ceteris paribus. This implies that $j > l \iff \text{Prob}(PM|\text{min len} = j) > \text{Prob}(PM|\text{min len} = l)$.

Next, consider $\text{Prob}(PM|\text{min len} = 1)$. Observe that

$$\gamma(n) \geq \text{Prob}(PM|\text{min len} = 1) = \left(1 - \frac{C(1)}{C(n)}\right)^{n-1} \cdot \gamma(n-1)$$

- Given min len = 1, we know that some applicant a finds only one spot s acceptable. In any perfect matching in this subset of markets, applicant a must be matched to spot s. There is a perfect matching if and only if we can find a perfect matching in the market where we additionally exclude (a, s).
- Consider the submarket excluding applicant a and spot s. The probability of a perfect matching in the submarket is the probability of a perfect matching in the $n-1 \times n-1$ market where the length of preference lists is uniformly distributed. Conditional on the event that all preference lists are at least length 1, this probability takes the

value of $\gamma(n-1)$. Meanwhile, the probability that some applicant has a preference list of at least length 1 is:

$$\frac{C(1)}{C(n)} = \frac{\binom{n-1}{1}1!(n-2)!}{\sum_{k=1}^{n-1}\binom{n-1}{k}k!(n-1-k)!};$$

Given applicant independence, the probability that all applicants find at least one spot other than s acceptable is: $m(n) = \left(1 - \frac{C(1)}{C(n)}\right)^{n-1}$.[26]

In fact, the inequality must be strict, since Prob(min len > 1) > 0 and we know that Prob(PM|min len = j) > Prob(PM|min len = 1) for all $j > 1$.

Therefore,

$$\gamma(n) = \varepsilon + \left(1 - \frac{C(1)}{C(n)}\right)^{n-1} \cdot \gamma(n-1) \tag{1}$$

$$\frac{\gamma(n)}{\gamma(n-1)} = \frac{\varepsilon}{\gamma(n)} + \left(1 - \frac{C(1)}{C(n)}\right)^{n-1} \tag{2}$$

$$\frac{\gamma(n)}{\gamma(n-1)} = \frac{\varepsilon}{\gamma(n)} + [1 - \nu(n)] \tag{3}$$

$$\geq \varepsilon + 1 - \nu(n) \tag{4}$$

For all $n \geq N_\varepsilon$, $|\nu(n)| = 1 - \left(1 - \frac{C(1)}{C(n)}\right) \leq \varepsilon$ and hence, for all $n \geq N_\varepsilon$, $\frac{\gamma(n)}{\gamma(n-1)} = 1 + \varepsilon - \nu(n) > 1$. i.e., for large enough n, we know that $\gamma(n)$ is strictly increasing.

Since $\gamma(n) \leq 1$, it must be that $\gamma(n) \to 1$.

Proposition 2 details

Let $\tilde{S} = \{s \in \mathcal{S} \ : \ \text{Rank}_{\succ_\alpha}(s) \geq k+1\}$ be the set of $n-k$ spots which applicant α ranks $k+1^{th}$ or below.

- Fixing a particular subset of applicants $\tilde{A} \subseteq \mathcal{A}/\{\alpha\}$, $|\tilde{A}| = j$, $S(j, n-k) \cdot (n-k)!$ counts the number of profiles of first choices among applicants in $\tilde{A}$ where the first choice of each $a' \in \tilde{A}$ is an element of $\tilde{S}$, and every spot in $\tilde{S}$ is the first choice of some $a \in \tilde{A}$.
- Fixing a particular subset of applicants $\tilde{A} \subseteq \mathcal{A}/\{\alpha\}$, $|\tilde{A}| = j$, $S(n-k, j) \cdot (n-k)! \cdot k^{n-j}$ counts the number of profiles of first choices among applicants in $\mathcal{A}/\{\alpha\}$ where the

[26] Note that $m(n)$ converges to 1, and takes a minimum value at $n = 3$. Intuitively, other candidates are extremely likely to find at least one more spot acceptable other than spot s.

first choice of each $a' \in \mathcal{A}/\{\alpha\}$ is an element of $\tilde{S}$ if and only if applicant $a' \in \tilde{A}$, and every spot in $\tilde{S}$ is the first choice of some $a \in \tilde{A}$.

Relative to $S(n-k,j)\cdot(n-k)!$, we now additionally account for the number of permutations of first choices among applicants in $\mathcal{A}/\{\tilde{A}\cup\alpha\}$.

- The first choice of every applicant in $\mathcal{A}/\{\tilde{A}\cup\alpha\}$ is among k spots which applicant α ranks k^{th} or higher.
- Fixing the first choices of applicants in $\tilde{A}$, given IID uniform preference lists, there are k^{n-j} possible first choice profiles of applicants in $\mathcal{A}/\{\tilde{A}\cup\alpha\}$.

- Hence, $\sum_{j=n-k}^{n-1}\binom{n-1}{j}S(n-k,j)\cdot(n-k)!\cdot k^{n-j}$ counts the number of profiles of first choices among members of $\mathcal{A}/\{\alpha\}$ where the sufficient condition in Lemma 1 is satisfied.

  Relative to $S(n-k,j)\cdot(n-k)!\cdot k^{n-j}$, we now additionally account for the fact that the group $\tilde{A}$, and the size of $\tilde{A}$, can be selected in many ways.

Proposition 3 details

Consider the set of all possible preference rankings. As there are finite number of possible preference rankings, let us assign each ranking an index i. i.e.,

$$\mathcal{P}_{\mathcal{S}} = \{\succ_{(1)},\cdots,\succ_{(|\mathcal{P}_{\mathcal{S}}|)}\}$$

A collection of rankings is defined by the size-n set of indexes of unique individual preference rankings, $\mathcal{C} = \{1,\ldots,|\mathcal{P}_{\mathcal{S}}|\}^n$. Let $\succ_{(\mathcal{C})}$ denote the collection of individual preference rankings where

$$\succ_{(\mathcal{C})} = \{\succ_{(i)} : i \in \mathcal{C}\}$$

Conditional on some collection of preference rankings, the probability that applicant a' has a preference ranking corresponding to $\succ_{(i)}$ is

$$P(\{\succ_{a'}=\succ_{(i)}\}|\succ_{(\mathcal{C})}) = \frac{\text{Number of } i \text{ in } \mathcal{C}}{n}$$

For notational simplicity, let $\{s_i \succ s_j\}$ denote $\{\succ \in \mathcal{P}_{\mathcal{S}} : s_i \succ s_j\}$,

Consider

$$P(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}) = \frac{P(\{s_i \succ_{a'} s_j\} \cap \{s_i \succ_\alpha s_j\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\})}{P(\{s_i \succ_\alpha s_j\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\})}$$

for some $\tilde{\succ} = \{\succ_{\sigma(1)}, \ldots, \succ_{\sigma(n-2)}\}$. Note:

$$\begin{aligned} P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \mathcal{P}_{\mathcal{S}}\}) &= P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) \\ &\quad + P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\}) \end{aligned}$$

and

$$\begin{aligned} P(\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}, s_i \succ_\alpha s_j) &= P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) \cdot \frac{2!(n-2)!}{n!} \\ &\quad + P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\}) \cdot \frac{1(n-2)!}{n!} \end{aligned}$$

At each step:

Suppose that $P(s_i \succ s_j) > P(s_j \succ s_i)$. Then,

- Switch with some probability ε, the collection $\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\} \times \{s_1 \succ s_2 \succ \ldots s_j \succ s_i \ldots s_{n-1} \succ s_n\}\}$ to the collection $\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\} \times \{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\}\}$ to $\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\} \times \{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\}\}$.
- And switch with some probability ε, the collection $\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_1 \succ s_2 \succ \ldots s_j \succ s_i \ldots s_{n-1} \succ s_n\} \times \{s_1 \succ s_2 \succ \ldots s_j \succ s_i \ldots s_{n-1} \succ s_n\}\}$ to $\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}} \text{ and } \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\} \times \{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\}\}$.

Hence:

- $P'(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}, \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\}) = P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}, \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\}) - \varepsilon$
- $P'(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}, \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) = P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}, \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) + 2\varepsilon$
- $P'(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}, \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\}) = P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}, \succ_{\mathcal{C}} / \tilde{\succ} \in \{s_j \succ s_i\} \times \{s_j \succ s_i\}\}) - \varepsilon$
- There are no changes for all $(l,k) \neq (i,j)$.

So,

$$P'(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}) = \frac{P'(\{s_i \succ_{a'} s_j\} \cap \{s_i \succ_\alpha s_j\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\})}{P'(\{s_i \succ_\alpha s_j\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\})}$$

$$= \frac{P(\{s_i \succ_{a'} s_j\} \cap \{s_i \succ_\alpha s_j\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\}) + \varepsilon \frac{2}{n(n-1)}}{P(\{s_i \succ_\alpha s_j\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\}) + \varepsilon \frac{2}{n(n-1)} - \varepsilon \frac{1}{n(n-1)}}$$
$$= P(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}) + \eta_\varepsilon, \quad \eta_\varepsilon > 0$$

and

$$P'(s_j \succ_{a'} s_i | s_j \succ_\alpha s_i, \tilde{\succ}) = \frac{P'(\{s_j \succ_{a'} s_i\} \cap \{s_j \succ_\alpha s_i\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\})}{P'(\{s_j \succ_\alpha s_i\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\})}$$
$$= \frac{P(\{s_j \succ_{a'} s_i\} \cap \{s_j \succ_\alpha s_i\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\}) - \varepsilon \frac{2}{n(n-1)}}{P(\{s_j \succ_\alpha s_i\} \cap \{\succ_{\mathcal{A}/\{\alpha,a'\}} \equiv \tilde{\succ}\}) - \varepsilon \frac{2}{n(n-1)} - \varepsilon \frac{1}{n(n-1)}}$$
$$= P(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}) + \eta'_\varepsilon, \quad \eta'_\varepsilon > 0$$

and $P'(s_l \succ_{a'} s_k | s_l \succ_\alpha s_k, \tilde{\succ}) = P(s_l \succ_{a'} s_k | s_l \succ_\alpha s_k, \tilde{\succ})$

Conditioning on other collections Let $\succ_*$ denote $\{s_1 \succ s_2 \succ \ldots s_i \succ s_j \ldots s_{n-1} \succ s_n\}$ and let $\succ_{**}$ denote $\{s_1 \succ s_2 \succ \ldots s_j \succ s_i \ldots s_{n-1} \succ s_n\}$.

Case 1a: Conditioning on $\tilde{\succ}' = \{\succ_{\sigma(1)}, \succ_{\sigma(2)}, \ldots, \succ_{\sigma(j-1)}, \succ_{\sigma(j+1)}, \cdots \succ_{\sigma(n-2)}\} \cup \{\succ_*\}$ and suppose that $s_i \succ_{\sigma(j)} s_j$. Observe that $p \equiv P(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) = P(\{\succ_\mathcal{C} : \tilde{\succ} \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ} \in \{\succ_*\} \times \{\succ_*\}\}) + P(\{\succ_\mathcal{C} : \tilde{\succ} \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ} \in \{\succ_*\} \times \{s_i \succ s_j\} / \{\succ_*\}\}) + P(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ} \in \{s_i \succ s_j\} / \{\succ_{\sigma(j)}\} \times \{s_i \succ s_j\}\})$ and the modification step only altered (increased) the first term. So, $P'(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) = P(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) + \varepsilon = p + \varepsilon$. Also, denoting $q \equiv P(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\})$, $P'(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\}) = P'(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{\succ_{\sigma(j)}\} \times \{s_j \succ s_i\}\}) + P'(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} / \{\succ_{\sigma(j)}\} \times \{s_j \succ s_i\}\}) = q - \varepsilon$ is decreased by the same amount, by an analogous logic, via changes to the first term. Applying similar logic from the first part of the proof, $P'(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}') = \frac{2(p+\varepsilon')}{2(p+\varepsilon')+(q-\varepsilon')} > \frac{2p}{2p+q} = P(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}')$.

Case 1b: Condition on $\tilde{\succ}' = \{\succ_{\sigma(1)}, \succ_{\sigma(2)}, \ldots, \succ_{\sigma(j-1)}, \succ_{\sigma(j+1)}, \cdots \succ_{\sigma(n-2)}\} \cup \{\succ_*\}$ and suppose that $s_j \succ_{\sigma(j)} s_i$. Let $p \equiv P(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\})$ and $q \equiv P(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\})$. $P'(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\})$ remains unchanged, and $P'(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_j \succ s_i\} \times \{s_i \succ s_j\}\}) = P'(\{\succ_\mathcal{C} : \tilde{\succ} \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ} \in \{\succ_*\} \times \{s_i \succ s_j\}\}) + P'(\{\succ_\mathcal{C} : \tilde{\succ}' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}' \in \{s_j \succ s_i\} / \{\succ_{\sigma(j)}\} \times \{s_i \succ s_j\}\}) = q - \varepsilon$. So, $P'(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}') = \frac{2p}{2p+(q-\varepsilon')} > \frac{2p}{2p+q} = P(s_i \succ_{a'} s_j | s_i \succ_\alpha s_j, \tilde{\succ}')$.

Case 2a: Condition on $\tilde{\succ}'' = \{\succ_{\sigma(1)}, \succ_{\sigma(2)}, \ldots, \succ_{\sigma(j-1)}, \succ_{\sigma(j+1)}, \cdots \succ_{\sigma(n-2)}\} \cup \{\succ_{**}\}$ and suppose that $s_j \succ_{\sigma(j)} s_i$. Observe that $P(\{\succ_\mathcal{C} : \tilde{\succ}'' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}'' \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\})$ remains constant, and $P(\{\succ_\mathcal{C} : \tilde{\succ}'' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}'' \in \{s_j \succ s_i\} \times \{s_i \succ s_j\}\}) = P(\{\succ_\mathcal{C} : \tilde{\succ} \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ} \in \{\succ_*\} \times \{s_i \succ s_j\}\}) + P(\{\succ_\mathcal{C} : \tilde{\succ}'' \in \succ_\mathcal{C} \text{ and } \succ_\mathcal{C} / \tilde{\succ}'' \in \{s_j \succ s_i\} / \{\succ_{\sigma(j)}\} \times \{s_i \succ s_j\}\})$ decreases by ε. Increased correlation

is immediate.

Case 2b: Condition on $\tilde{\succ}'' = \{\succ_{\sigma(1)}, \succ_{\sigma(2)}, \ldots, \succ_{\sigma(j-1)}, \succ_{\sigma(j+1)}, \cdots \succ_{\sigma(n-2)}\} \cup \{\succ_{**}\}$ and suppose that $s_i \succ_{\sigma(j)} s_j$. $P(\{\succ_{\mathcal{C}} : \tilde{\succ}'' \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ}'' \in \{s_i \succ s_j\} \times \{s_i \succ s_j\}\}) = P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ} \in \{\succ_{**}\} \times \{\succ_{*}\}\}) + P(\{\succ_{\mathcal{C}} : \tilde{\succ} \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ} \in \{\succ_{**}\} \times \{s_i \succ s_j\}/\{\succ_{*}\}\}) + P(\{\succ_{\mathcal{C}} : \tilde{\succ}'' \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ}'' \in \{s_i \succ s_j\}/\{\succ_{\sigma(j)}\} \times \{s_i \succ s_j\}\})$ decreases by ε, while the denominator decreases by the same amount since $P(\{\succ_{\mathcal{C}} : \tilde{\succ}'' \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ}'' \in \{s_i \succ s_j\} \times \{s_j \succ s_i\}\}) = P(\{\succ_{\mathcal{C}} : \tilde{\succ}'' \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ}'' \in \{\succ_{**}\} \times \{\succ_{**}\}\}) + P(\{\succ_{\mathcal{C}} : \tilde{\succ}'' \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ}'' \in \{\succ_{\sigma(j)}\} \times \{s_j \succ s_i\}/\{\succ_{**}\}\}) + P(\{\succ_{\mathcal{C}} : \tilde{\succ}'' \in \succ_{\mathcal{C}}$ and $\succ_{\mathcal{C}} / \tilde{\succ}'' \in \{s_i \succ s_j\}/\{\succ_{\sigma(j)}\} \times \{s_j \succ s_i\}/\{\succ_{**}\}\})$. Overall, $P'(s_i \succ_{a'} s_j | s_i \succ_{\alpha} s_j, \tilde{\succ}'') = \frac{2(p-\varepsilon')}{2(p-\varepsilon')+(q-\varepsilon')} > \frac{2p}{2p+q} = P(s_i \succ_{a'} s_j | s_i \succ_{\alpha} s_j, \tilde{\succ}'')$.

Lemma 3

Proof by Contradiction: Suppose block-k anti-correlation is insufficient. Then, for some block-k anti-correlated $\succ_{\mathcal{A}}$, there exists some priority order $\tilde{\succ}_{\mathcal{S}}$ where at least one applicant will be rejected from their k^{th} favorite spot. It follows that some applicant $a' \in A_l$ must be rejected from all $s \in S_l$, since $|S_l| \leq k$. Without loss of generality, let a' be the first such applicant. Let $\underline{s}$ be the spot applicant a' that prefers the least within S_k.[27] Suppose also that a' is rejected from $\underline{s}$ at Step t of the DA mechanism.

At step t, every spot $\tilde{s} \in S_l$ is tentatively matched to some applicant other than a'.[28] Let $\tilde{A}_l(t)$ denote the set of applicants tentatively matched to some $\tilde{s} \in S_k$ at step t. By construction, $|\tilde{A}_l(t)| = |S_l| = |A_l| > |A_l/\{a'\}|$. So, we know $\tilde{a} \in \tilde{A}_l(t)$ for some $\tilde{a} \in A_j$, $j \neq k$. However, if such an $\tilde{a}$ applied to some $s \in S_l$ by step t of the mechanism, $\tilde{a}$ must have been rejected by every spot in S_j, their own set of favorite spots, before step t. This contradicts our premise that a' is the first such spot.

Lemma 4

If all applicants are guaranteed their top-1 or top-2 outcomes, we can coarsely partition $\mathcal{A}$ into the subsets $\mathcal{A}_{(1)}$ and $\mathcal{A}_{(2)}$. $\mathcal{A}_{(1)}$ is a (potentially empty) set $\mathcal{A}_{(1)}$ of applicants who are guaranteed their first choice, while $\mathcal{A}_{(2)}$ is a (non-empty) set of applicants who are only guaranteed their second choice.

The first choice of every applicant in $\mathcal{A}_{(1)}$ must be ranked below the worst-case outcome of every $a \in \mathcal{A}/\{a'\}$. Let $\mathcal{S}_{(1)}$ denote the set of spots that are the first choice of applicants in $\mathcal{A}_{(1)}$. Directly, $s \in \mathcal{S}_{(1)}$ must be ranked 2^{nd} or below by all but one. It must also be true that every $s \in \mathcal{S}_{(1)}$ is ranked 3^{rd} or below by applicants in $\mathcal{A}_{(2)}$. Suppose not, then some $s \in \mathcal{S}_{(1)}$ is ranked 2^{nd} for some

[27] i.e., $\text{Rank}_{\succ_{a'}}(\underline{s}) = |S_k|$.

[28] We have assumed that every applicant is acceptable, so applicant a' could only have been rejected by spot s if a preferable applicant applied to spot s.

$a' \in \mathcal{A}_{(2)}$. Applicant a' must be guaranteed their first choice if they have a top-2 guarantee since it is impossible to guarantee spot s for applicant a'. Yet, this contradicts $a' \in \mathcal{A}_{(2)}$.

The top-2 outcomes of every $a' \in \mathcal{A}_{(2)}$ must be ranked after the worst-case outcome of every $a \in \mathcal{A}/\{a', \tilde{a}\}$, where applicant $\tilde{a} \neq a'$ has the same first and second choice as applicant a'. To see this, consider some $a' \in \mathcal{A}_{(2)}$ whose first two choices are s_1 and s_2 respectively. First, at least one other applicant must share s_2 as a second choice. Whenever applicant a' is matched to spot s_1, some other applicant $\tilde{a}$ must be matched to s_2.[29] Next, at least one pair of applicants who rank s_2 as a second choice must have a common first choice. Otherwise, it is possible that no applicant applies to s_2 within their first two applications, which is incongruent with a top-2 market-wide guarantee. Finally, since applicant a' has a top-2 guarantee, it must be that *exactly* one pair of applicants with s_2 as a second choice have the same first choice, and applicant a' must be one of these applicants. Furthermore, no third applicant must rank s_1 as their first choice, and no third applicant in $\mathcal{A}_{(2)}$ must rank s_2 as their second choice. If not, at least one applicant does not have a top-2 guarantee.

Lemma 5

Fixing a profile of applicant preferences $\succ_{\mathcal{A}}$ and given some matching $\tilde{\mu}$ implemented by $\succ_{\mathcal{S}}$ under DA, we can always construct an aligned priority such that every applicant weakly prefers their match outcome under this priority compared to $\tilde{\mu}$. So, if a perfect matching that matches the condition in *Lemma* 1 can be implemented by some $\succ_{\mathcal{S}}$, then there exists an aligned priority that also implements some perfect matching that matches the condition in *Lemma* 1.

I describe the process of constructing such an aligned priority here, where one sequentially assigns each applicant a priority rank.

Suppose we are given some matching $\tilde{\mu}$, and want to construct an aligned priority such that every applicant does at least as well as under $\tilde{\mu}$.

- Step 1: Consider $P(1) = \{a \in \mathcal{A} \ : \ \text{Rank}_{\succ_a}(\tilde{\mu}(a)) = 1\}$: The set of applicants who are matched to their first choice under $\tilde{\mu}$. Assign each $a \in P(1)$ one of best available priorities: $\{1, \ldots, |P(1)|\}$. Denote the set of spots that these applicants are matched to as $S(1) = \{s \in \mathcal{S} \ : \ \tilde{\mu}(s) \in P(1)\}$.
- Step 2: Consider $P(2) = \{a \in \mathcal{A} : \text{Rank}_{\succ_a}(\tilde{\mu}(a)) = 2\} \cap \{a \in \mathcal{A} \ : \ \{s \in S : \text{Rank}^{-1}_{\succ_a}(s) \leq 1\} \subseteq S(1)\}$. i.e., $P(2)$ consists of all applicants who (1) are matched to their second choice under $\tilde{\mu}$, and (2) whose first choice is matched under $\tilde{\mu}$ to another applicant assigned priority in Step 1. Assign each $a \in P(2)$ one of best available priorities: $\{|P(1)| + 1, \ldots, |P(1)| + |P(2)|\}$. Denote the set of spots these applicants are matched to as $S(2) = \{s \in \mathcal{S} \ : \ \tilde{\mu}(s) \in P(2)\}$.

[29] We have assumed a balanced market where all spots and applicants are mutually acceptable. Also, s_2 cannot be the first choice of another applicant if it can be guaranteed.

- Step k: Consider $P(k) = \{a \in \mathcal{A} : \text{Rank}_{\succ_a}(\tilde{\mu}(a)) \leq k\} \cap \{a \in \mathcal{A} : \{s \in S : \text{Rank}^{-1}_{\succ_a}(s) \leq \text{Rank}_{\succ_a}(\tilde{\mu}(a))\} \subseteq \cup_{j=1}^{k-1} S(j)\}$. i.e., $P(k)$ consists of all applicants who (1) are matched to their k^{th} or better choice under $\tilde{\mu}$, and (2) all spots they prefer to $\tilde{\mu}(a)$ are matched under $\tilde{\mu}$ to another applicant assigned priority in an earlier step.

  Assign each $a \in P(k)$ one of best available priorities: $\{1 + |\cup_{j=1}^{k-1} P(j)|, \ldots, |P(k)| + |\cup_{j=1}^{k-1} P(j)|\}$. Denote the set of spots these applicants are matched to as $S(k) = \{s \in \mathcal{S} \ : \ \tilde{\mu}(s) \in P(k)\}$.

  We can see via induction that each applicant $a \in P(k)$ must be matched to $\tilde{\mu}(a)$ under DA given such an assignment of priorities. Directly, all applicants $a \in P(1)$ are matched to their partner under $\tilde{\mu}$. Suppose the induction hypothesis is true, every $a \in \cup_{j=1}^{k-1} P(j)$ must be matched to $\tilde{\mu}(a')$. By construction, $\tilde{\mu}(a') \succ_{a'} \tilde{\mu}(a)$.

- Step $n+1$: Consider $\mathcal{A} / \cup_{j=1}^{n} P(j)$.
  - Case 1: $\mathcal{A} / \cup_{j=1}^{n} P(j) = \emptyset$. All applicants have been assigned a priority ranking. Under such a priority ranking, each applicant $a \in \mathcal{A}$ must be matched to $\tilde{\mu}$. We have found an aligned priority that works.
  - Case 2: $\mathcal{A} / \cup_{j=1}^{n} P(j) \neq \emptyset$. For each remaining applicant $a \in \mathcal{A} / \cup_{j=1}^{n} P(j)$, $\tilde{\mu}(a) \in \mathcal{S} / \cup_{j=1}^{n} S(j)$. Furthermore, for each $a \in \mathcal{A} / \cup_{j=1}^{n} P(j)$, there exists some $s \in \mathcal{S} / \cup_{j=1}^{n} S(j)$ such that $s \succ_a \tilde{\mu}(a)$. Run the following process:
    - Start with $P(n+1) = \emptyset$ and $S(n+1) = \emptyset$
    - Step A) Pick some $a \in \mathcal{A} / \cup_{j=1}^{n+1} P(j))$, and assign this applicant the priority rank $|\cup_{j=1}^{n} P(j)| + 1$. Add a to $P(n+1)$, and let $\bar{\mu}(a)$ denote applicant a's favorite spot in $\mathcal{S} / \cup_{j=1}^{n+1} S(j)$. Add $\bar{\mu}(a)$ to $S(n+1)$.
    - Step B) For every $a' \in \mathcal{A} / \cup_{j=1}^{n+1} P(j)$, check that there exists some $s \in \mathcal{S} / \cup_{j=1}^{n+1} S(j)$. If this condition is true, go back to Step A. Otherwise, the remaining applicants and spots form a balanced market $\tilde{A} \subset \mathcal{A}$ and $\tilde{S} \subset \mathcal{S}$.

  Repeat Steps 1 to $n+1$ for the market comprising only $\tilde{A} \subset \mathcal{A}$ and $\tilde{S} \subset \mathcal{S}$, to determine priority rankings among low priority applicants in $\mathcal{A} / \cup_{j=1}^{n} P(j)$.

Ambiguity over Priorities versus Ambiguity over Preferences

What does the distribution of worst-case outcomes look like if we are instead robust over applicant preferences, and are probabilistic over priorities?

Suppose $\succ_\alpha = s_1 \succ_\alpha s_2 \succ_\alpha \cdots \succ_\alpha s_n$. For any possible priority ranking $\succ_\mathcal{S}$, the set of $\succ_\mathcal{A}$ that produces the worst outcome for applicant α a set always includes the completely aligned profile where every other applicant has the same ranking as applicant α. Hence, the probability that

applicant α is guaranteed at least his k^{th} outcome is the probability that priorities generate a DA outcome where applicant α is matched to a spot ranked k^{th} or better given all applicants share the preference ranking $\succ_\alpha$.

If applicant preferences are aligned, then the DA outcome of any priority $\succ_{\mathcal{S}}$ is equivalent to the DA outcome of some aligned priority.[30] Without loss of generality, we can restrict our domain of preferences to aligned priorities.

If we suppose that applicants and spots are ex-ante identical, every applicant must have the same probability of being in each priority rank. In such a setting, applicant α must have a uniform probability over worst-case ranks. i.e.,

$$\text{Prob}_{\succ_{\mathcal{S}}}(\min_{\succ_{\mathcal{A}}}\{\text{Rank}(\mu(\alpha;\succ_{\mathcal{S}}))\} \leq k) = \frac{k}{n}$$

Implications Given our result in Proposition 1, we can make the following observation:

$$\text{Prob}_{\succ_{\mathcal{S}}}(\min_{\succ_{\mathcal{A}}}\{\text{Rank}(\mu(\alpha;\succ_{\mathcal{S}}))\} \leq k) \geq \text{Prob}_{\succ_{\mathcal{A}}}(\min_{\succ_{\mathcal{S}}}\{\text{Rank}(\mu(\alpha;\succ_{\mathcal{A}}))\} \leq k) \quad \forall k \in \{1,\ldots,n\}.$$

In words, in a setting with symmetric applicants and spots, the worst-case distribution when we are robust against priorities and probabilistic over spots first-order stochastically dominates the worst-case distribution if we were robust against applicants' priorities but probabilistic over priorities. We have a "better" distribution in the latter case.

Holding fixed any profile of applicant preferences $\succ_{\mathcal{A}}$, we can in most cases find a priority profile $\succ_{\mathcal{S}}$ which generates a bad outcome for applicant α. Conversely, holding fixed a profile of priorities $\succ_{\mathcal{S}}$, in a smaller number of cases we can find $\succ_{\mathcal{A}}$ which generates a bad outcome for α. In that sense, knowledge about one's priority among other applicants is much more "informative" for applicant α about how well or badly they might do in a centralized matching process employing DA. Without information about one's priorities, or some assurance that one will have a decent priority, an ambiguity-averse applicant will be much more pessimistic about their potential outcomes from participating. TLDR: telling people about other applicants' preferences is much less helpful for an ambiguity-averse applicant than assuring them about their priorities.

[30] Namely, assign the applicant ranked first by s_1 top priority; then, assign the applicant ranked first by s_2 (excluding the first applicant) second priority, and so on.

References

Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match. *American Economic Review*, 99(5):1954–1978, December 2009. ISSN 0002-8282. doi: 10.1257/aer.99.5.1954.

Atila Abdulkadiroğlu, Nikhil Agarwal, and Parag A. Pathak. The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match. *American Economic Review*, 107(12):3635–3689, December 2017. ISSN 0002-8282. doi: 10.1257/aer.20151425.

Nikhil Agarwal and Paulo Somaini. Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism. *Econometrica*, 86(2):391–444, 2018. ISSN 0012-9682. doi: 10.3982/ECTA13615.

Itai Ashlagi and Yannai A. Gonczarowski. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425, September 2018. ISSN 0022-0531. doi: 10.1016/j.jet.2018.07.001.

Itai Ashlagi, Yash Kanoria, and Jacob D. Leshno. Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, 125(1):69–98, February 2017. ISSN 0022-3808. doi: 10.1086/689869.

Itai Ashlagi, Afshin Nikzad, and Assaf Romm. Assigning more students to their top choices: A comparison of tie-breaking rules. *Games and Economic Behavior*, 115:167–187, May 2019. ISSN 0899-8256. doi: 10.1016/j.geb.2019.02.015.

Yeon-Koo Che and Olivier Tercieux. Efficiency and stability in large matching markets. *Journal of Political Economy*, 127(5):2301–2342, October 2019. ISSN 0022-3808. doi: 10.1086/701791.

Yan Chen and Tayfun Sönmez. School choice: An experimental study. *Journal of Economic Theory*, 127(1):202–231, March 2006. ISSN 00220531. doi: 10.1016/j.jet.2004.10.006.

Yi-Chun Chen and Gaoji Hu. A Theory of Stability in Matching with Incomplete Information. *American Economic Journal: Microeconomics*, 15(1):288–322, February 2023. ISSN 1945-7669, 1945-7685. doi: 10.1257/mic.20200411.

Marcelo Ariel Fernandez. Deferred Acceptance and Regret-Free Truth-Telling. *Working Paper*, 2024.

Marcelo Ariel Fernandez, Kirill Rudov, and Leeat Yariv. Centralized matching with incomplete information. *American Economic Review: Insights*, 4(1):18–33, March 2022. ISSN 2640-2068. doi: 10.1257/aeri.20210123.

D. Gale and L. S. Shapley. College Admissions and the Stability of Marriage. *The American*

Mathematical Monthly, 69(1):9–15, January 1962. ISSN 0002-9890, 1930-0972. doi: 10.1080/00029890.1962.11989827.

Itzhak Gilboa and Massimo Marinacci. Ambiguity and the Bayesian Paradigm. In Horacio Arló-Costa, Vincent F. Hendricks, and Johan van Benthem, editors, *Readings in Formal Epistemology: Sourcebook*, pages 385–439. Springer International Publishing, Cham, 2016. doi: 10.1007/978-3-319-20451-2_21.

Avinatan Hassidim, Assaf Romm, and Ran I. Shorrer. "strategic" behavior in a strategy-proof environment. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 763–764, New York, NY, USA, July 2016. Association for Computing Machinery. doi: 10.1145/2940716.2940751.

Nicole Immorlica and Mohammad Mahdian. Incentives in Large Random Two-Sided Markets. *ACM Transactions on Economics and Computation*, 3(3):14:1–14:25, June 2015. ISSN 2167-8375. doi: 10.1145/2656202.

Adam Kapor, Mohit Karnani, and Christopher Neilson. Aftermarket Frictions and the Cost of Off-Platform Options in Centralized Assignment Mechanisms. *Journal of Political Economy*, 132 (7):2346–2395, July 2024. ISSN 0022-3808. doi: 10.1086/729068.

Fuhito Kojima and Parag A. Pathak. Incentives and stability in large two-sided matching markets. *The American Economic Review*, 99(3):608–627, 2009. ISSN 0002-8282.

Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287, November 2017. ISSN 0002-8282. doi: 10.1257/aer.20160425.

Qingmin Liu, George J. Mailath, Andrew Postlewaite, and Larry Samuelson. Stable Matching With Incomplete Information. *Econometrica*, 82(2):541–587, 2014. ISSN 0012-9682. doi: 10.3982/ECTA11183.

Marek Pycia and Peter Troyan. Strategy-proof, Efficient, and Fair Allocation: Beyond Random Priority. *Working Paper*, February 2023.

Alex Rees-Jones. Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match. *Games and Economic Behavior*, 108:317–330, March 2018. ISSN 0899-8256. doi: 10.1016/j.geb.2017.04.011.

Clayton Thomas. Classification of Priorities Such That Deferred Acceptance is Obviously Strategyproof. (arXiv:2011.12367), February 2021. doi: 10.48550/arXiv.2011.12367.

Peter Troyan. Obviously Strategy-Proof Implementation of Top Trading Cycles. *International Economic Review*, 60(3):1249–1261, 2019. ISSN 0020-6598.