What makes a matching market congested?*

Justin Hadad† En Hua Hu‡

August 23, 2026

Abstract

We study a decentralized matching market where each applicant sends a fixed number of applications and then firms make one offer. Our game emulates matching markets under time constraints: agents who are unmatched after the round remain unmatched. Congestion arises from two basic market failures: some firms do not receive applications (an issue of coverage), and some applicants receive multiple offers (an issue of collisions). We study how the market size, degree of preference alignment, and number of applications affect congestion. Aligned preferences worsen congestion when there is no possibility to coordinate, and additional applications do not always alleviate congestion; optimal quotas are typically small.

1 Introduction

In decentralized matching markets, mutually beneficial matches may fail to form. This inefficiency persists across a variety of settings: even while on the short side of a large market imbalance, PhD programs often under-enroll students, search committees fail to hire, and low-income housing remains unfilled.

Market designers have broadly advocated for the centralization of matching markets to circumvent this exact failure. In centralized markets, agents observe the complete set of possible matches and either submit preference orderings to a mechanism or propose to prospective matches in sequence. But in practice, several markets, including university applications and sectors in the labor market, are still decentralized: no institution arranges the matching or direct proposals. In settings of this sort, applicants instead submit applications, firms offer an applicant who has applied, and applicants accept at most one offer. We study a stylized game of this form, and study how *congestion* — when market activity congregates around certain agents and thus others remain unmatched — naturally emerges.

*For their comments, we would like to thank Aytek Erdil, Irene Lo, Margaret Meyer, Daniel Quigley, Ludvig Sinander, Augustus Smith, Al Roth, and Alex Teytelboym, as well as seminar participants at the University of Oxford, Stanford University, Zurich Center for Market Design, the Transatlantic Theory Workshop 2025, the 2026 Conference on Mechanism and Institution Design, and the Twenty-Seventh ACM Conference on Economics and Computation (EC '26).

†University of Oxford; justin.hadad@economics.ox.ac.uk

‡University of Oxford; huenhua123@gmail.com

Decentralized markets vary in their market size and the degree of preference correlation (across either and/or both sides of the market). We study how a principal with limited power can regulate markets that vary in these primitives, toward maximizing some joint preference of match rate and applicant welfare.[1] While the principal cannot run a centralized mechanism, she can set a *quota* q that specifies the maximum number of applications that applicants can send in the first round of the game. Regulation in the form of application limits is practical in light of institutional details; centralization often requires authorizations and infrastructure that might not be feasible in the near term.[2] Further, application limits are a common policy, and a surprising variety of application limits exists in practice: the UCAS undergraduate application cycle allows only five applications, meanwhile the Common App allows 20; at the same time, in the Shizuoka Prefecture of Japan, prospective high-school students can send only one application to a public school. The question, then, is two-fold: why these policies vary, and whether they are indeed optimal for a given objective.

We begin by studying the game with one round of applications: applicants send q applications and then firms make one offer to an applicant who has applied to it; unmatched applicants and firms cannot reapply or re-offer. We find that increasing the quota has two opposing effects: while it increases the number of firms that receive at least one application (increasing *coverage*), it also increases the likelihood that offers concentrate on the same applicants (increasing *collision*). The match rate is exactly the number of firms that are *covered*, minus the number of offers that *collide*. Thus the principal sets the quota as a function of the *coverage-collision tradeoff*, and how this tradeoff manifests is itself a function of the market primitives.

Our first case of interest is the game where agents on each side of the market have independent preferences for agents on the other side. Under these independent preferences, there is a natural symmetric equilibrium in which applicants apply to their q highest-utility firms, firms offer the highest-utility applicant who has applied, and applicants accept their highest-utility offer.[3] We study how to set the quota under these strategies, and we find that the quota which maximizes the match rate is small. Indeed, when there is sufficient imbalance in the market — at least two applicants per firm — the optimal quota is one. This minimal quota is optimal because it maximizes the match rate, and because the utility conditional on a match is highest when $q = 1$, it also maximizes applicant welfare. Further, under any weak imbalance and with an additional large-market assumption, we find that no optimal quota can be larger than three. The large-market assumption simplifies the probability of receiving multiple offers; when the numbers of applicants and firms are finite, offers are not independent due to finite-sampling effects.[4]

[1] Match rate is a natural objective in settings that suffer from congestion, while maximizing applicant welfare (commonly student welfare) is standard in the matching literature.

[2] The seminal article promoting centralized kidney exchange, Roth et al. (2004), appeared in the Quarterly Journal of Economics in 2004. Congress clarified only in 2007 that kidney paired donation was not prohibited under the National Organ Transplant Act, and the pilot program became operational nationally only in 2010.

[3] When values are independent everywhere, agents cannot infer the probability of an application/offer turning into a match via their own values.

[4] The proof of our main result which bounds these dependencies in the finite market may be of independent interest; we discuss our proof strategy in Section 5.

That the optimal quota is small in this setting may bear some surprise, but can be understood through the coverage-collision tradeoff. Observe that for a fixed quota q, as the number of applicants increases relative to the number of firms, firms are more likely to get at least one application; when coverage is guaranteed, minimizing collisions then becomes the central concern. Small quotas minimize the number of collisions and thus are strong in this setting.

We show that the dominance of small quotas when preferences are independent is robust to uniformly increasing firm capacities, as well. We extend our game to a many-to-one matching setting where each firm has the same fixed number c of seats. In that setting, so long as $c \geq 2$ and there are weakly more applicants than total seats, the optimal quota is one.

In addition, our results shed light on which side — the long or the short side — should be granted proposal rights in a random matching market when only one round of a deferred-acceptance (or immediate-acceptance) game is feasible. Observe that when $q = 1$, applicants apply to only one firm, and firms offer (and ultimately match with) their favorite of the applicants who have applied. Meanwhile, when the quota is unlimited, applicants have the weakly dominant strategy to apply everywhere;[5] then each firm chooses one applicant to offer from the entire set, and applicants accept their best offer. Thus the $q = 1$ regime is identical to one round of the applicant-proposing game, and the unlimited quota is identical to one round of the firm-proposing game. We show that, when there are independent preferences, the match rate is always higher when the long side proposes. Intuitively, the match rate suffers less when the long side runs the risk of concentrating their actions on the same agents.

The preceding results rely on both applicants applying to individual firms with equal probability in expectation, and firms offering individual applicants with equal probability in expectation. We then adapt our model to allow for correlated preferences, where ex-ante uniform strategies no longer dominate and agents hedge in equilibrium.

We examine three kinds of correlation: firm-side, applicant-side, and cross-side (where applicants' values for firms are positively correlated with firms' values for applicants). We introduce each correlation while assuming unaffected preferences are independent, and we find that, when agents hedge symmetrically, the match rate generally decreases for a given quota. Ceteris paribus, if firms correlate their offers, collision increases without changing coverage, thus the match rate decreases for any quota satisfying $q > 1$.[6] Meanwhile, if applicants correlate their applications, coverage decreases and the effect on collision is ambiguous; we find a sufficient condition for the match rate to decrease for any quota $q < F$, where F is the number of firms. In the case of cross-side correlation, "star" applicants are generated endogenously: an applicant who receives an offer from one firm is then conditionally more likely to receive an offer from another. As some applicants are more likely to receive multiple offers, match rate decreases.[7]

[5] We assume all values are weakly positive.

[6] Observe that at $q = 1$, collisions are impossible, so the match rate is the same with and without firm correlation.

[7] For some intuition, observe that coverage stays the same when we introduce cross-side correlation (because values are still drawn i.i.d.), but the collisions conditional on the applications increase. This follows from an order-statistics argument, and a similar observation was made in the tie-breaking literature; see Ashlagi and Nikzad (2020).

When the decrease in match rate is due to an increase in collisions, small quotas stay powerful. Further, because the $q = 1$ regime is unaffected by collisions, whenever the minimal quota maximizes the match rate in the independent-preference setting, it also does when we introduce firm-side and cross-side correlation. However, if applicants are expected to concentrate their applications such that coverage falls, then the quota which maximizes the match rate cannot decrease; in that regime, the optimal quota can be large.

Taken together, our results offer justification for small quotas that we see in practice, but also present settings where large quotas can be optimal (i.e., if coverage lacks because applicants are the source of congestion). Still, we do not see our results as prescriptive for implementing quotas in lieu of centralizing the market. Indeed, sequential applications or offers can greatly mitigate congestion. To highlight the benefit of centralization, we adapt our model to a game with sequential rounds of proposals and acceptances. In particular, we are interested in how allowing unmatched firms to re-offer compares to implementing any static quota. We find a stark result, again using the large-market assumption in a setting with independent preferences: setting $q = F$ and letting firms send one additional offer if they are unmatched generates a higher match rate than any $q \geq 2$.

Toward practical insight, we show how a designer with limited power might regulate a time-constrained matching market. If relaxing these time constraints were feasible, our results suggest that such relaxation is productive for mitigating congestion. In turn, if centralization were infeasible, then the designer should set the quota while taking into account the source of congestion in the market.

1.1 Related Literature

Congestion has been a classic concern in matching markets and has proven influential on market and platform design. For instance, congestion has pushed markets to unravel or become centralized; in the seminal study on congestion in matching, it was shown that slow offer responses in the clinical psychology match led employers to make earlier, exploding offers (Roth and Xing, 1994). A body of literature has emerged that considers different institutional constraints and policies that mitigate the effect of congestion.

A prominent line of research treats congestion as arising from information frictions and search costs. Under uncertainty over which applications lead to offers, applicants may over-apply, creating a negative externality of wasted screening. When search becomes too cheap and decentralized, firm screening costs can be high, and the equilibrium features too many applications. A stark example of this is in school choice: a common application system lowers application costs and expands the set of schools each student applies to, but because the admissions system remains uncoordinated, congestion can arise and be costly to a social welfare function (Avery et al., 2025). Preference signaling has been proposed as a remedy: allowing candidates to send a limited signal of interest can improve matching outcomes by implicitly guiding search (Coles et al., 2013; Horton et al., 2024). While we do not study the use of preference signaling, we do study the benefit of limiting applications, and qualitatively support the idea that the number of allowed signals should likely

be small. Further, in our model, preference alignment can be interpreted as a form of screening; for instance, applicants apply to firms that in turn value the applicants. Hence we show that, in contrast, better screening can reduce the match rate and worsen congestion.

In platform design, restricting agents' actions can reduce wasteful screening by firms (Kanoria and Saban, 2021). Relatedly, Arnosti et al. (2021) distinguishes settings characterized by many applications from those characterized by intensive firm screening. A complementary approach treats congestion as an externality and studies Pigouvian taxation where applicants internalize the screening burden (He and Magnac, 2018). In contrast, we focus on markets absent endogenous screening costs. We also show that a lower application limit can improve the market, but through a different and more fundamental channel via the coverage-collision tradeoff.

In school choice, there is a growing literature specifically devoted to rectifying failures of contemporary application systems; for instance, Doğan and Erdil (2025) highlight how a lack of coordination makes schools miss enrollment targets. Related literature studies the portfolio problem where applicants apply strategically. In decentralized college admissions, Che and Koh (2016) shows that uncertainty over the other offers that applicants will receive leads colleges to avoid direct competition for top students. Hedging becomes an equilibrium response to correlated admissions risk (Ali and Shorrer, 2025). We abstract from the details of strategic choice and consider how applicants and firms who mix symmetrically affect the match rate.

2 Model

We study a decentralized, one-round matching market with applicants and firms. There is a finite set of applicants $\mathcal{A} = \{1, \ldots, A\}$ and a finite set of firms $\mathcal{F} = \{1, \ldots, F\}$, with $A \geq F$. We denote an economy by the market-size pair (A, F), and sometimes collectively call applicants and firms *agents*. The timing is as follows.

1. Applicants and firms privately observe their realized utilities.
2. Each applicant submits applications to q firms,[8] where $q \in \{1, \ldots, F\}$ is the *quota*.
3. Each firm makes an offer to an applicant who has applied to it (if any).
4. Each applicant accepts at most one offer, and the game ends.

The quota q is an integer and is uniform across applicants. The preceding game results in a matching, for which we use the usual definition. A matching is a function $\mu : \mathcal{A} \cup \mathcal{F} \to \mathcal{A} \cup \mathcal{F} \cup \{\varnothing\}$ such that, for every applicant $a \in \mathcal{A}$, $\mu(a) \in \mathcal{F} \cup \{\varnothing\}$; for every firm $f \in \mathcal{F}$, $\mu(f) \in \mathcal{A} \cup \{\varnothing\}$; and $\mu(a) = f$ if and only if $\mu(f) = a$. If $\mu(a) = \varnothing$, applicant a is unmatched; if $\mu(f) = \varnothing$, firm f is unmatched.

For each quota q, the game induces a distribution over the set of matchings satisfying the preceding definitions; we denote the random matching from this distribution μ_q.

[8] We will assume that utilities are weakly positive so applying to q firms is optimal.

For each applicant-firm pair (a, f), applicant a receives utility $u_{af} \geq 0$ from matching with firm f, and firm f receives utility $u_{fa} \geq 0$ from matching with applicant a. In the *baseline* model that we describe in this section, the random variables $\{u_{af}\}_{a\in\mathcal{A},f\in\mathcal{F}}$ are drawn independently from identical continuous distributions, as are the random variables $\{u_{fa}\}_{f\in\mathcal{F},a\in\mathcal{A}}$; we likewise assume independence across the sets. The utility of remaining unmatched is zero, $u_{a\varnothing} = 0$ and $u_{\varnothing f} = 0$.

The expected match rate under quota q is

$$m_q = \mathbb{E}\left[\frac{1}{F}\sum_{f\in\mathcal{F}} \mathbf{1}\left\{\mu_q(f) \neq \varnothing\right\}\right], \tag{1}$$

where the expectation is taken over the realized utilities and equilibrium strategies. Applicant welfare under quota q is the expected average utility of applicants:

$$U_q = \mathbb{E}\left[\frac{1}{A}\sum_{a\in\mathcal{A}} u_{a\mu_q(a)}\right]. \tag{2}$$

When utilities are independent, for every distribution of values there is a symmetric equilibrium where each applicant applies to the q firms from which she receives the highest utility, and each firm offers the highest-utility applicant who has applied to it. To see that application and offer strategies defined in this way constitute an equilibrium, observe that an applicant's utility for a firm contains no information about her probability of receiving an offer from that firm, and a firm's utility for an applicant is independent of the utilities that other firms have for that applicant. This equilibrium assumes the inability to coordinate; for example, we rule out the possibility of agents to collude or condition their strategies on their identity.

Coverage and collision. We can decompose the match rate from (1) in the following way. The number of matches is the number of firms that receive at least one application, which we call *coverage*, less the number of rejected offers, which we call *collision*:

$$\mathbb{E}\left[\#\text{ matches}\right] = \mathbb{E}\left[\#\text{ firms receiving at least one application}\right] - \mathbb{E}\left[\#\text{ wasted offers}\right]. \tag{3}$$

This formula displays the tradeoff in how the designer sets the quota: while a higher quota increases the probability that each firm receives at least one application and thus can make an offer, it also increases the probability that firms collide when making such offers. Our results concern how the quota should be set according to this *coverage-collision tradeoff*.

Toward setting the quota, we allow the designer to have any joint preference over match rate and applicant welfare as in (1) and (2) respectively. We say a quota is *optimal* if there is no such quota which could be preferred to it for any joint preference.

Definition 1. *A quota q^* is optimal if it jointly maximizes both expected match rate and applicant*

welfare:

$$q^* \in \arg\max_{q\in\{1,\dots,F\}} m_q \quad \cap \quad \arg\max_{q\in\{1,\dots,F\}} U_q.$$

When an optimal quota exists, the designer does not face a tradeoff between the number of matches formed and applicant welfare. We additionally identify a notion by which quotas can be compared along the same two dimensions.

Definition 2. *A quota q dominates quota q' if and only if $m_q \geq m_{q'}$ and $U_q \geq U_{q'}$.*

Taking Definitions 1 and 2 together, the optimal quota dominates all other quotas.

Large market. Our baseline model has A and F finite, and we model the number of applications and offers via the binomial distribution. In this finite market, offer events are statistically dependent; the probability of getting an offer from one firm is negatively dependent on the probability of getting an offer from another firm. Conditional on an applicant not receiving an offer from one firm, then there is slightly less expected competition at the other firms to which she applied.

At several points we use a large-market approximation to simplify the offer probabilities. Specifically, we let $\lambda := \frac{A}{F}$ and study a sequence of markets in which $A, F \to \infty$ with λ fixed. Two simplifications arise in that limit: the Poisson distribution approximates the number of competing applications at a firm, and offer probabilities are independent.

Assumption 1. *In the large market, the Poisson distribution approximates the binomial distribution and offer probabilities are independent.*

While the large-market assumption greatly simplifies analysis, our results in the finite market that bound the effects of dependence may be of interest. We highlight our proof strategy in Section 5.

3 The Baseline Coverage-Collision Tradeoff

The application quota regulates a basic tradeoff: raising the quota increases coverage, but also increases the likelihood that offers collide. We analyze how the quota should be set in the baseline model espoused in Section 2.

Our first result compares the two extremal regimes: a quota of one and the unlimited quota. When $q = 1$, each applicant applies only to her most-preferred firm; and since each applicant submits only one application, any offer she receives is accepted. Similarly, when $q = F$ (equivalently, the game without any quota), every applicant applies to every firm; each firm then offers her most-preferred applicant from the full set of applicants, and applicants accept their best offer. The comparison between $q = 1$ and $q = F$ is thus insightful toward the following question: What side should propose in one round of a deferred- or immediate-acceptance game when preferences are independent?

Proposition 1. *In the baseline one-to-one market with $A \geq F$, $q = 1$ dominates $q = F$.*

Note in this proposition and several other results, we restate the key assumption that $A \geq F$; this market imbalance drives the results. Proposition 1 implies that, in one round of the canonical random matching model, the principal maximizes match rate by letting the long side propose. We provide the first steps of the proof to build intuition. Under $q = 1$, each applicant applies to her most-preferred firm, and since applicants submit only one application, any applicant who receives an offer accepts it. Thus the number of realized matches is exactly the number of firms that receive at least one application. And since each applicant applies uniformly to one of the F firms, the match rate under $q = 1$ is exactly

$$m_1 = 1 - \left(1 - \frac{1}{F}\right)^A.$$

Under $q = F$, each applicant applies to every firm, so every firm makes an offer to its most-preferred applicant from the full set of A applicants. Hence the number of matches is exactly the number of applicants who receive at least one offer. Since each firm's top applicant is uniformly distributed over the A applicants,

$$m_F = \frac{A}{F}\left(1 - \left(1 - \frac{1}{A}\right)^F\right).$$

We formally show in Appendix A that $m_1 \geq m_F$ for any $A \geq F$. To show dominance as in Definition 2, it remains to compare applicant welfare. Observe that under $q = 1$, every matched applicant is matched with her most-preferred firm, and under $q = F$, a matched applicant accepts her most-preferred offer among the firms that offered her, which need not be her most-preferred firm. Because the probability of matching is weakly higher under $q = 1$, *and* applicant utility is weakly higher conditional on a match, the claim follows.

Proposition 1 is the first indication that the short side concentrating their actions is especially costly. When firms choose one applicant from the complete set, they are likely to send competing offers to the same applicants. By contrast, when applicants are limited to one application, the market sacrifices some coverage but makes the probability of collision zero. Proposition 1 compares these two regimes, and finds that the zero-collision outcome generates a higher match rate than the full-coverage outcome.

The preceding result compares only the two extremal regimes. Our next result shows that, in the baseline model with sufficient imbalance, the optimal application limit is minimal.

Theorem 1. *In the baseline one-to-one market with $A \geq 2F$, $q^* = 1$.*

Once there are at least two applicants per firm, coverage ceases to be the central concern: there are already enough applicants for firms to be covered with high probability, and additional applications mainly create competing offers to the same applicants. We find that $q = 1$ maximizes the match rate when $A \geq 2F$, and because the utility conditional on a match is maximized at $q = 1$, dominance follows.

Observe that an optimal quota might not exist if the quota which maximizes the match rate satisfies $q > 1$. Hence whether one quota dominates another is sensitive to distributional assumptions. In our next result, we use Assumption 1 and show that the match rate is maximized at some $q \leq 3$, thus no optimal quota satisfies $q > 3$. Further, because utility conditional on a match is higher at smaller quotas, any $q > 3$ is dominated by $q \leq 3$.

Theorem 2. *Under Assumption 1 in the baseline one-to-one market with $A \geq F$, any $q > 3$ is dominated by $q \leq 3$. Further, the quota which maximizes the match rate decreases as the market imbalance increases.*

Under Assumption 1, if an applicant sends q applications, the number of competing applicants at any given firm is drawn from a Poisson distribution with mean λq. Hence the probability that a given application results in an offer under quota q is

$$s_q = \frac{1 - e^{-\lambda q}}{\lambda q}. \tag{4}$$

For exposition, let $u_{(1)} \geq u_{(2)} \geq \cdots \geq u_{(F)}$ denote a focal applicant's expected order statistics over firms, so she prefers firm 1 to firm 2, and so on. By assumption, her probability of getting an offer from any one of the F firms is equal and defined in (4). The applicant then accepts the offer from firm 1 if she receives it, and accepts the offer from firm 2 if she receives it and does not receive an offer from firm 1, and so on. We can then write the expected utility for this focal applicant under quota q in the large market as

$$U_q^L = \sum_{r=0}^{q-1} (1 - s_q)^r \, s_q \, u_{(r+1)}. \tag{5}$$

Differentiation reveals the quota that maximizes the match rate can be no larger than three, and is weakly decreasing in $\lambda = A/F$.[9] In Figure 1, we plot (5) for $u_{af}, u_{fa} \in \text{Uni}[0, 1]$ with different market imbalances.
In Figure 1, one can additionally see that $q^* = 1$ when $\lambda \geq 2$, giving evidence of Theorem 1.

We now extend our baseline model in two directions, both using Assumption 1. First, we show that our results are robust in a limited sense to relaxing the capacity constraints of firms. We allow each firm to have an equal integer capacity $c \geq 2$, and show that if $A \geq Fc$, then $q^* = 1$. Then we relax time constraints and compare a sequence of applications and offers to setting the static quota; we show that dynamic applications and offers generally outperform the optimal static quota, unless $q^* = 1$.

[9] The match rate can be gleaned from (5) by setting $u_{(f)} = 1$ for all $f \in \mathcal{F}$. Note that welfare computations penalize larger quotas more than match rate computations.

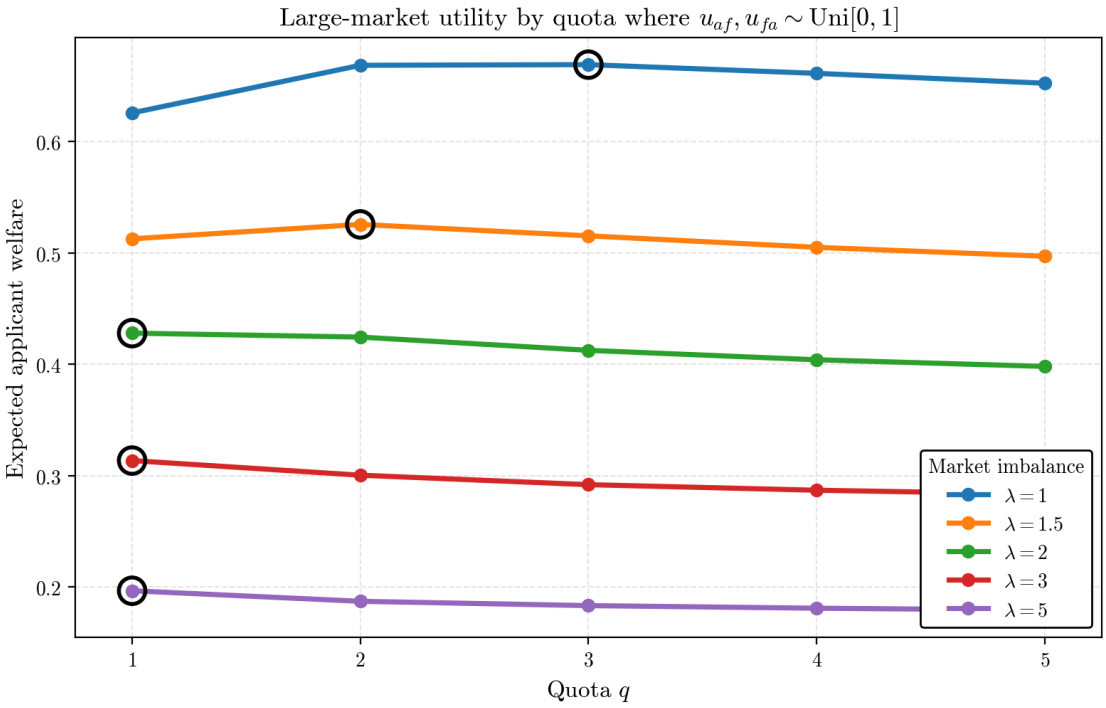


Figure 1: Expected applicant welfare when $u_{a,f}, u_{f,a} \in \text{Uni}[0,1]$ for different market imbalances λ. The quota which maximizes expected welfare is circled.

3.1 Many-to-One Markets

We slightly adapt the model. Let each firm have the same capacity $c \geq 2$ and suppose Assumption 1 holds; let $\lambda = A/F$ and suppose $\lambda \geq c$. We elide the definition of a matching in this section but the usual many-to-one matching definition applies; because we use Assumption 1, we can define the match rate more directly than as a function of the random matching.

We study the natural extension of the symmetric equilibrium from the baseline model. Each applicant applies to her q highest-utility firms,[10] each firm offers its $\min\{c, n\}$ highest-utility applicants among the n applicants who have applied, and each applicant accepts her highest-utility offer.

Let $\pi_{q,c}$ denote the probability that a focal applicant matches under quota q and uniform capacity c, and let $U_{q,c}$ denote expected average applicant welfare. Since each match fills one seat, the match rate satisfies

$$\widehat{m}_{q,c} = \frac{\lambda}{c}\pi_{q,c}. \tag{6}$$

Optimality and dominance are defined using $\widehat{m}_{q,c}$ and $U_{q,c}$ in extension of Definitions 1 and 2.[11] Fix a focal applicant and one of her applications. Under Assumption 1, the number of competitors at this focal firm is $K_q \sim \text{Poisson}(\lambda q)$. Conditional on $K_q = k$, the focal applicant receives an offer if

[10] Note that assuming equal capacities makes firms ex-ante symmetric, so applicants do not hedge toward firms with greater capacity.

[11] We elide a definition of $U_{q,c}$ in this section as our result only needs to compare match rates.

she ranks among the firm's c highest-utility applicants. Because firm preferences are independent, the probability that an individual application becomes an offer is therefore

$$s_{q,c} = \mathbb{E}\left[\min\left\{1, \frac{c}{K_q + 1}\right\}\right]. \tag{7}$$

Offer events across an applicant's applications are independent in the large market, so her probability of matching is

$$\pi_{q,c} = 1 - (1 - s_{q,c})^q. \tag{8}$$

Proposition 2. *Fix an integer $c \geq 2$. Under Assumption 1 and if $A \geq Fc$, then $q^* = 1$.*

The condition $A \geq Fc$ (equivalently, $\lambda \geq c$) when $c \geq 2$ is analogous to our condition in Theorem 1. In fact, a single application is more effective when each firm has several seats and the number of applicants scales accordingly. The applicant need only rank among the firm's top c applicants rather than rank first; if she does not get an offer at one seat, she may get an offer at another seat at the same firm.

3.2 Adding Rounds

We now relax the maximal time constraint and allow unmatched agents to act again. We consider two dynamic regimes, applicant-proposing and firm-proposing, and we assume that agents observe whether they themselves are matched, but not the full set of agents who remain unmatched. Thus the symmetric equilibrium from the baseline model extends to the dynamic game: agents apply to/offer their next most-preferred agent on the other side of the market who has not yet rejected them.

We study this extension in the large market using Assumption 1, which allows us to write the evolution of the match rate across rounds in a tractable way. We still assume preferences are independent within and across sides of the market.

Fix an integer $T \geq 1$. The **T-round applicant-proposing deferred-acceptance** game is as follows:[12]

1. In round $t = 1, \ldots, T$, every applicant who is not currently holding an offer applies to her most-preferred firm among those to which she has not previously applied.

2. Each firm considers all applicants who applied to it in the current round together with any applicant whose offer it was already holding from a previous round. The firm tentatively holds its most-preferred applicant among this set and rejects all others.

3. Rejected applicants proceed to the next round. Applicants whose offers are held do not apply again unless later rejected.

4. After round T, all tentative offers are finalized.

[12] Toward a simple exposition, we assume the equilibrium behavior when defining the two games.

Applicants and firms not matched after round T remain unmatched. The **T-round firm-proposing deferred-acceptance game** is defined symmetrically; in the initial stage, the firm chooses one applicant from the set of all applicants.

1. In round $t = 1, \ldots, T$, every firm not currently holding an acceptance makes an offer to its most-preferred applicant among those to whom it has not previously made an offer.

2. Each applicant considers all offers received in the current round together with any offer she was already holding from a previous round. The applicant tentatively holds her most-preferred offer among this set and rejects all others.

3. Rejected firms proceed to the next round. Firms whose offers are held do not offer again unless later rejected.

4. After round T, all tentative acceptances are finalized.

As in the applicant-proposing game, applicants and firms not matched after round T remain unmatched.

Let $\rho \in \{A, F\}$, where $\rho = A$ denotes the applicant-proposing game and $\rho = F$ denotes the firm-proposing game. For each $T \geq 1$, let μ_T^ρ denote the random matching induced by the T-round ρ-proposing game. Utilities are realized only from the final matching. Applicant welfare under the T-round ρ-proposing game is

$$U_T^\rho = \mathbb{E}\left[\frac{1}{A}\sum_{a \in \mathcal{A}} u_{a,\mu_T^\rho(a)}\right]. \tag{9}$$

The expected match rate under the T-round ρ-proposing process is

$$h_T^\rho = \mathbb{E}\left[\frac{1}{F}\sum_{f \in \mathcal{F}} \mathbf{1}\left\{\mu_T^\rho(f) \neq \varnothing\right\}\right]. \tag{10}$$

We set $h_0^A = h_0^F = 0$. Using Assumption 1 we can write (10) in a simple recursive form for each regime:

$$h_{t+1}^A = h_t^A + \left(1 - h_t^A\right)\left(1 - e^{-\left(\lambda - h_t^A\right)}\right), \tag{11}$$

$$h_{t+1}^F = h_t^F + \left(\lambda - h_t^F\right)\left(1 - e^{-\left(1 - h_t^F\right)/\lambda}\right). \tag{12}$$

We compare dynamic processes to quotas in the one round game by comparing the random matchings they induce. We adapt the definition of *domination* from Definition 2 to allow comparisons between the T-round ρ-proposing process and the quota q in the baseline game.

Definition 3. *The T-round ρ-proposing process dominates (is dominated by) a quota q in the baseline game if $h_T^\rho \geq (\leq) m_q$ and $U_T^\rho \geq (\leq) U_q$.*

As in Definition 2, the T-round ρ-proposing process dominates a quota q if it is preferred for any joint preference over match rate and applicant welfare. We begin our comparison of the dynamic regime and the quota game by making a simple observation about the two games: q rounds of applicant-proposing dynamics is lower bounded in terms of match rate and applicant welfare by a quota of q.

Remark. *The q-round applicant-proposing game dominates a quota of q.*

For intuition, consider an adapted setting where it is possible for each applicant to be rejected in each of the first $q-1$ rounds. This is pessimal under the equilibrium strategy we study and also reproduces the static quota-q outcome exactly. In any case which improves upon this, an applicant matches with a firm she prefers to her q-th favorite.[13]

We now ask whether the long side or short side should make sequential proposals when the number of rounds is finite.

Proposition 3. *In the large market with $A \geq F$, T rounds of applicant-proposing dynamics dominates T rounds of firm-proposing dynamics.*

Proposition 3 shows that the same force in Proposition 1 persists across rounds: offers are more likely to overlap on agents when the short side proposes. In other words, there is no number of rounds which reverses the comparison from Proposition 1. The proof for Proposition 3 is simple and follows from comparing the recursive form of the match rate as in (11) and (12). In Figure 2, we plot the match rate for different values of λ.

Thus q rounds of applicant-proposing dynamics dominates both a quota of q and q rounds of firm-proposing dynamics. We now evaluate the performance of quotas vis-à-vis firms sending additional offers. The condition for which the regime with additional offers could be preferred to the static quota is slight: allowing just one re-offer generates a higher match rate than any quota $q \geq 2$.

Theorem 3. *Under Assumption 1 when $A \geq F$, two rounds of firm-proposing dynamics generates a higher match rate than any $q \geq 2$.*

Thus when the quota that maximizes the match rate is greater than one, a designer who cares only about match rate prefers to abandon the quota and instead let firms send an additional offer if their first is rejected. This does not imply domination in the sense of Definition 3; applicant welfare may be higher when applicants apply to their top-q firms.[14]

Theorem 3 underscores the benefit of dynamic or centralized mechanisms: the benefit of a single re-offer is large and can improve upon several optimal static policies. Recall from Theorem 2 that

[13] Additionally, it will be shown in the following results that $h_q^A \geq h_2^A \geq h_2^F \geq m_q$ for $q \geq 2$, and when $q = 1$, the outcomes are the same. The applicant-welfare component follows from the fact that offers are independent of applicant utilities.

[14] Further, our approach allows comparisons between an arbitrary number of rounds and any static quota. To generate a higher match rate than a quota of one, there must be two rounds when $\lambda < 3.2$, three rounds when $\lambda < 6.3$, and so on.

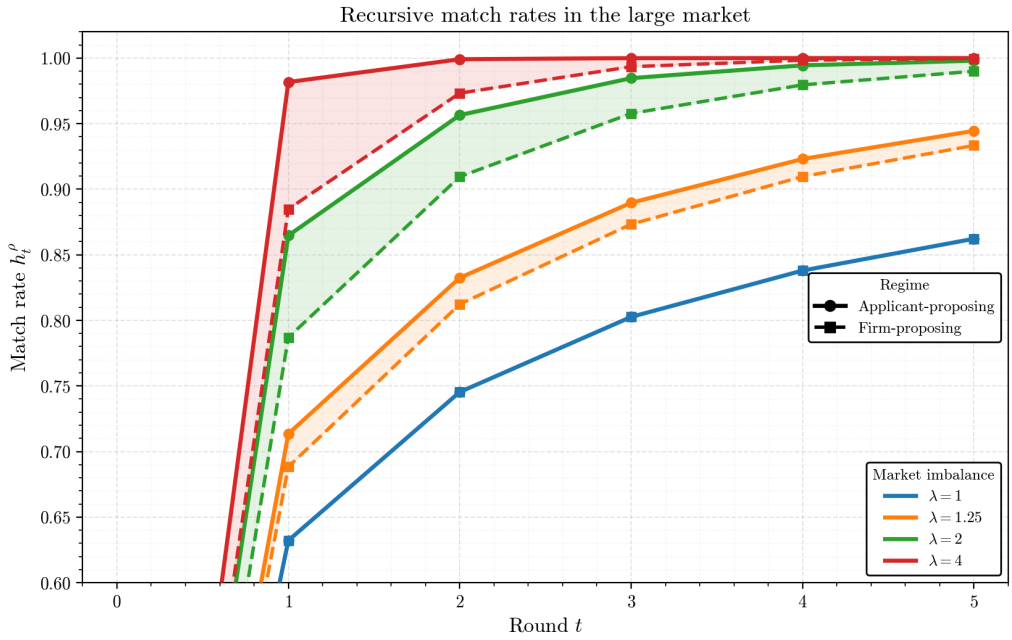


Figure 2: Large-market match rates under applicant-proposing and firm-proposing dynamics. Solid lines denote applicant-proposing and dashed lines denote firm-proposing; shaded regions show the gap between the two processes for each market imbalance $\lambda = A/F$.

the quota maximizing the match rate increases up to three as the market gets more balanced. This is precisely the environment where collisions are likely and allowing re-offers is most helpful.

4 Correlation

In Section 3, an applicant's utility for a firm contained no information about how many other applicants would apply to that firm, or whether the firm was likely to make her an offer. Similarly, a firm's utility for an applicant was uninformative about the applicant's probability of acceptance, and whether other firms would make offers to the same applicant. Thus, in the symmetric equilibrium studied previously, application and offer probabilities were ex-ante uniform, and agents could not condition their actions on likely congestion.

We study one key departure from this setting: agents mix symmetrically in expectation but the probabilities of applications and offers are non-uniform. This behavior is consequent to agents having correlated values.

We separately consider three forms of preference alignment.[15] In the *applicant-side* alignment case, applicants have correlated preferences over firms and thus may hedge; in the *firm-side* alignment case, firms have correlated preferences for applicants and thus may hedge. We then consider *cross-side* correlation, where individual applicants and firms have positively correlated values for

[15] When coverage and collision are both affected by different forms of correlation, identifying the quotas that maximize the match rate requires additional structure. This is difficult but promising for future work.

each other; the maximum cross-side correlation is that where the values are perfectly symmetric.[16]

In this section, some additional notation is necessary to formalize the three environments we study. For a summary, see Table 1.

For every applicant $a \in \mathcal{A}$ and firm $f \in \mathcal{F}$, let v_{af} denote a pair-specific common match component, let ε_{af} denote applicant a's idiosyncratic value for firm f, and let ε_{fa} denote firm f's idiosyncratic value for applicant a. For every applicant a, let θ_a denote applicant a's common attractiveness to firms, and for every firm f, let θ_f denote firm f's common attractiveness to applicants. The values $\{v_{af}\}_{a\in\mathcal{A},f\in\mathcal{F}}$, $\{\varepsilon_{af}\}_{a\in\mathcal{A},f\in\mathcal{F}}$, $\{\varepsilon_{fa}\}_{f\in\mathcal{F},a\in\mathcal{A}}$, $\{\theta_a\}_{a\in\mathcal{A}}$, and $\{\theta_f\}_{f\in\mathcal{F}}$ are mutually independent; within each set, random variables are i.i.d. and drawn from atomless distributions with support in $\mathbb{R}_+$. We assume that the idiosyncratic values ε_{af} and ε_{fa} have log-concave densities.[17] The parameters $\alpha, \beta^a, \beta^f \in [0,1]$ govern the weight placed on common components: cross-side alignment varies α, applicant-side alignment varies β^a, and firm-side alignment varies β^f.

Table 1: Correlated-preference environments.

Alignment type	Applicant value u_{af}	Firm value u_{fa}
Applicant-side	$\beta^a\theta_f + (1-\beta^a)\,\varepsilon_{af}$	ε_{fa}
Firm-side	ε_{af}	$\beta^f\theta_a + \left(1-\beta^f\right)\varepsilon_{fa}$
Cross-side	$\alpha v_{af} + (1-\alpha)\,\varepsilon_{af}$	$\alpha v_{af} + (1-\alpha)\,\varepsilon_{fa}$

In the environment with cross-side correlation, we study equilibrium as in the baseline model: applicants apply to their q highest-utility firms, and firms offer their highest-utility applicant among those who applied. In the applicant-side and firm-side environments, there may be multiple equilibria; we assume that agents hedge symmetrically toward higher common-value agents on the other side of the market. We formalize this hedging now.

The strategy of applicant a maps her utility vector u_a to a lottery over subsets of size q of $\mathcal{F}$; we write $\sigma_a(f \mid u_a)$ for the marginal probability that applicant a applies to firm f. For any realized application profile, let $P_f \subseteq \mathcal{A}$ denote the set of applicants who applied to firm f. A firm strategy maps (P_f, u_f) to a lottery over elements of P_f whenever $P_f \neq \varnothing$; we write $\sigma_f(a \mid P_f, u_f)$ for the probability that firm f offers applicant $a \in P_f$.

Fix a focal applicant a and let the firms to which she has applied be $\{f_1, \ldots, f_q\}$. For each $\ell \in \{1, \ldots, q\}$, let $N_\ell = |P_{f_\ell} \setminus \{a\}|$ denote the number of other applicants who applied to the ℓ-th firm to which a has applied.

For each correlated-preference environment, we restrict attention to *symmetric monotone equilibria.*

Definition 4. *A symmetric equilibrium is monotone if the following hold.*

[16] There is a growing literature on matching with aligned preferences (Fernandez et al., 2021; Ferdowsian et al., 2025); in these games, preferences can be represented by a single matrix.

[17] We only use log-concavity in the environment with cross-side correlation, and use it to ensure that a firm's value for an applicant is stochastically increasing in the applicant's value for that firm.

(i) *For every applicant a and firms f, f', if $u_{af} \geq u_{af'}$, then $\sigma_a(f \mid u_a) \geq \sigma_a(f' \mid u_a)$. Likewise, for every firm f and applicants $a, a' \in P_f$, if $u_{fa} \geq u_{fa'}$, then $\sigma_f(a \mid P_f, u_f) \geq \sigma_f(a' \mid P_f, u_f)$.*

(ii) *Fix an applicant a and firm f, and let $\widetilde{u}_a$ differ from u_a only in the coordinate corresponding to f. If $\widetilde{u}_{af} \geq u_{af}$, then $\sigma_a(f \mid \widetilde{u}_a) \geq \sigma_a(f \mid u_a)$. Similarly, fix a firm f, an application pool P_f, and an applicant $a \in P_f$. Let $\widetilde{u}_f$ differ from u_f only in the coordinate corresponding to a. If $\widetilde{u}_{fa} \geq u_{fa}$, then $\sigma_f(a \mid P_f, \widetilde{u}_f) \geq \sigma_f(a \mid P_f, u_f)$. And for any $a' \in P_f \setminus \{a\}$, if $\widetilde{u}_f$ differs from u_f only in the coordinate corresponding to a' and $\widetilde{u}_{fa'} \geq u_{fa'}$, then $\sigma_f(a \mid P_f, \widetilde{u}_f) \leq \sigma_f(a \mid P_f, u_f)$.*

The monotone symmetric equilibrium is the natural response to common values on one side of the market. Our definition has to be quite involved to ensure that a higher common component θ_f attracts weakly more applications, and a higher θ_a attracts weakly more offers. For the following result, we assume agents on the correlated-value side play strategies according to the monotone symmetric equilibria, and the other side applies to/offers their highest-utility agents on the other side of the market. Again, in the cross-side-correlated environment, we assume the same strategies as in the baseline game.

Recall m_q denotes the match rate in the baseline game with independent preferences. Let m_q^{α}, $m_q^{\beta^f}$, and $m_q^{\beta^a}$ denote the expected match rates under cross-side, firm-side, and applicant-side alignment, respectively, evaluated at quota q.

Theorem 4. *Fix a quota q. Under the equilibrium selections described above:*

1. $m_q \geq m_q^{\alpha}$.
2. $m_q \geq m_q^{\beta^f}$.
3. $m_q \geq m_q^{\beta^a}$ *if the competitor counts faced by a focal applicant satisfy*

$$\mathbb{E}\left[\prod_{\ell=1}^{q} \frac{N_\ell}{N_\ell + 1}\right] \geq \prod_{\ell=1}^{q} \mathbb{E}\left[\frac{N_\ell}{N_\ell + 1}\right].$$

The remainder of this section is devoted to discussing Theorem 4.

Firm-side and cross-side correlation. We first focus on firm-side correlation. In this case, firms playing strategies as in the symmetric monotone equilibria is sufficient to guarantee a weakly decreased match rate for any size quota. Application strategies are the same as in the baseline game, but firms are now more likely to collide. Hence any quota will induce a weakly smaller match rate, and the match rates are equal at $q = 1$.

When there is cross-side correlation, we assume applicants play the same strategies as in the baseline game. The decrease in match rate is again due to an increase in collisions: conditional on already receiving an offer, an applicant is more likely to receive another. This follows from an order-statistics argument; a simple example illustrates, which we provide in Figure 3.



(a) First-firm competition

(b) Induced comparison at the second firm

Figure 3: Intuition for the increase in collisions due to cross-side correlation. Conditional on f_1 being each applicant's most-preferred firm, a higher value at f_1 induces a stochastically higher value at f_2. The applicant preferred by f_1 is therefore more likely also to be preferred by f_2.

In Figure 3, two applicants compete at two firms; preferences are perfectly symmetric and the value distribution is Uni$[0, 1]$. If both applicants apply to the same firm under $q = 1$ (which happens with positive probability in expectation), then when $q = 2$, the utility for the losing applicant is drawn from a truncated distribution that is a subset of the winning applicant's distribution.

Recall that when $q = 1$, collisions are impossible. Because the decrease in match rate in both the firm-side and cross-side environments is due to an increase in collisions, setting $q = 1$ generates the same match rate with and without these kinds of correlation, hence we have the following corollary.

Corollary 1. *Under cross-side and firm-side correlation, if $A \geq F$, then $q = 1$ dominates $q = F$; if $A \geq 2F$, then $q^* = 1$.*

We now devote attention to a sketch of the proof that $m_1 \geq m_F$ under perfectly cross-side-correlated preferences. When cross-side correlation is maximal, preferences for applicants and firms can be represented by a single matrix. It will be helpful to think about the ranks of values; specifically, for $\eta \leq A \cdot F$, the η-th highest rank is in some entry of the matrix, and it is equally likely that a rank occurs in any given entry. We will give a sketch that there is order dominance in the ranks that enter the matching when $q = 1$ versus $q = F$.

Let there be A rows to represent applicants and F columns to represent the firms. Since $A \geq F$, we refer to rows as the long side and columns as the short side. Under $q = 1$ — equivalently, when applicants propose — the lowest rank (highest value) in a row is always sent as an application. Each firm then offers the lowest-rank applicant among the applications it receives, and that becomes a match. Meanwhile, when $q = F$ — equivalently, when firms propose — the columns and rows are reversed. Thus the comparison between applicant-proposing and firm-proposing regimes is a comparison between rows first choosing their lowest rank versus columns first choosing their lowest rank. Figure 4 has one particular instance of values which illustrates this.

Observe that no matter if rows or columns select first, the lowest rank will always enter the matching: the corresponding application is sent, then the corresponding offer is made and accepted. Now consider the η-th rank. When rows select first, a rank is certainly blocked (meaning not able to be in the final matching) if a lower rank η' is in its same row, but is not certainly

	f_1	f_2	f_3
a_1	7	$\boxed{2}$	11
a_2	$\boxed{1}$	8	10
a_3	9	12	$\boxed{3}$
a_4	6	~~$\boxed{4}$~~	5

(a) Applicant-proposing: each row selects first

	f_1	f_2	f_3
a_1	7	$\boxed{2}$	11
a_2	$\boxed{1}$	8	10
a_3	9	12	$\boxed{3}$
a_4	6	4	5

(b) Firm-proposing: each column selects first

Figure 4: A common-value rank matrix under perfect cross-side correlation. Lower numbers are better ranks. In the applicant-proposing game, each applicant selects her best-ranked firm first, meaning rows choose first; firm f_2 receives two applications so rejects a_4, hence rank 4 is crossed out. In the firm-proposing game, each firm selects its best-ranked applicant first, meaning columns choose first.

Case 1: η blocked by $\eta' < \eta$

$$\begin{bmatrix} \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots \\ \eta & \boxed{\eta'} & \cdots \\ \vdots & \vdots & \vdots \end{bmatrix}$$

Case 2: η not blocked by $\eta' < \eta$ because $\eta'' < \eta'$

$$\begin{bmatrix} \cdots & \cdots & \cdots \\ \eta' & \boxed{\eta''} & \cdots \\ \boxed{\eta} & \cdots & \cdots \\ \vdots & \vdots & \vdots \end{bmatrix}$$

Figure 5: Blocking logic in the common-value rank matrix. In the left panel, entry η is blocked because $\eta' < \eta$ lies in the same row and is therefore chosen. In the right panel, $\eta' < \eta$ does not block η because it is itself blocked by $\eta'' < \eta$ in its row. Thus lower ranks in the same row certainly block, while lower ranks in the same column are only potential blockers.

blocked if a lower rank η' is in its same column; indeed, there could be an even lower rank $\eta'' < \eta' < \eta$ in the same row as η'. Thus when there are fewer columns than rows, the probability of a value being blocked by another value in the column is smaller. The designer wants there to be as few certain blockers as possible, and so the designer should choose the longer side to select first — which corresponds to setting $q = 1$ rather than $q = F$. Figure 5 illustrates.

Applicant-side correlation. We now discuss the case of applicant-side correlation and its corresponding condition in Theorem 4. Under applicant-side alignment, the direct effect is on coverage: applications become concentrated on the same firms and firms receive zero applications with higher probability. To see the force at the applicant level, fix a focal applicant who applies to q firms, and recall N_ℓ is the number of competing applicants at the ℓ-th firm in her portfolio. Conditional on the competitor counts, the probability that she is rejected by firm ℓ is $\frac{N_\ell}{N_\ell+1}$. Hence the probability that she is rejected by all q firms is $\prod_{\ell=1}^{q} \frac{N_\ell}{N_\ell+1}$. This rejection probability is supermodular in the competitor counts: the marginal harm of facing more competition at one firm is larger when the applicant already faces high competition at the other firms. Indeed, if competition is low at some

firm to which she applied, another firm being crowded is not very costly for the applicant because she is still likely to receive an offer.

In the baseline model with independent preferences, the numbers of competitors across firms to which an applicant has applied are negatively associated. This negative dependence is to the benefit of the applicant: congestion at one firm is partially offset by lower congestion elsewhere. Our sufficient condition for applicant-side correlation to decrease the match rate for any quota $q < F$ essentially reverses the logic from the independent case. The condition

$$\mathbb{E}\left[\prod_{\ell=1}^{q} \frac{N_\ell}{N_\ell + 1}\right] \geq \prod_{\ell=1}^{q} \mathbb{E}\left[\frac{N_\ell}{N_\ell + 1}\right]$$

formalizes the positive association: we require the probability of being rejected by all firms to be weakly higher than if the competitor counts were independent with the same marginal distributions.

Further, when applicants correlate, the quota which maximizes the match rate can be large. Indeed, suppose firms have independent preferences and the applicant-side alignment and equilibrium are such that, for every quota q, all applicants apply to the same firms. In that case, large quotas are useful to obtain coverage; indeed, the match rate would then increase in q.

Broadly speaking, Theorem 4 gives practical guidance for how congestion should be regulated by quotas. If congestion is generated on the firm side — by way of either firm-side or cross-side correlation — a designer has incentive to decrease the quota and limit collision. Meanwhile, when congestion has its source in coverage issues, the best policy may be to increase the quota, and resolve issues of coverage at the expense of increasing the probability of offers going to the same applicants.

5 Proof Foundations of Theorem 1

Recall we used Assumption 1 to treat an applicant's chances of receiving offers from different firms as independent. In the finite market, this independence does not hold: the same competing applicants may or may not apply to the same firms as the focal applicant. The proof of Theorem 1 shows that this dependence can be controlled sharply enough to recover simple quota comparisons.[18]

Our proofs require translating the coverage-collision tradeoff from (3) into an inclusion-exclusion formula. Fix a quota q and a focal applicant a. Conditional on the set of firms to which a has applied, let O_r denote the event that the r-th firm in her portfolio makes her an offer, and let π_q denote the probability that a receives at least one offer. Since an applicant is matched if and only if she receives at least one offer, the expected number of matches is equal to $A\pi_q$, and comparing

[18] A finite-market version of Theorem 2 is harder because it must address the intermediate region $F \leq A < 2F$, and there, $q = 1$ need not dominate quota $q = 2$. One approach which bears promise is to show that $q = 2$ achieves a higher match rate than every $q \geq 4$. But to address the region $q \geq 4$, we need to bound the probability of receiving offers from sets of four (and more) firms. The bookkeeping there gets more involved but the underlying intuition is the same as we discuss here.

match rates across quotas is equivalent to comparing the probabilities that a focal applicant receives at least one offer. Let $I_{r,q} = \Pr(O_1 \cap \cdots \cap O_r)$ be the probability that a particular fixed set of r firms all make offers to the focal applicant.

The first step is to derive a closed form for being offered by one exact firm. Conditional on the focal applicant applying to the firm, the number of other applicants who applied there is binomial with parameters $A-1$ and q/F. The firm then offers the focal applicant with probability equal to the reciprocal of the number of applications it received.

Lemma 1. *For any quota q, the probability that a given firm offers the focal applicant is $I_{1,q} = \frac{F}{Aq}\left[1-\left(1-\frac{q}{F}\right)^A\right]$.*

If offer events were independent, the probability of receiving two offers would be $I_{1,q}^2$; but in the finite market, it is weakly below. We bound that difference.

Lemma 2. *For every $q \geq 2$, the two-offer probability satisfies $I_{1,q}^2 - \delta_q \leq I_{2,q} \leq I_{1,q}^2$, where $\delta_q = \frac{F^2}{q^3(A-1)A^2}$.*

Again consider a focal applicant and observe that a competing applicant chooses q distinct firms as well. Should this competing applicant apply to a firm to which the focal applicant applied, she has fewer remaining applications available to compete at some other firm to which the focal applicant also applied. Hence there is negative dependence in the competition faced by the focal applicant. Indeed, the upper bound $I_{2,q} \leq I_{1,q}^2$ says that two offers are no more likely than they would be under independence.

When $q = 2$, the additional coverage created by the second application must be compared to the probability that the same applicant receives both offers. Lemma 2 effectively helps us predict enough collision to show that once $A \geq 2F$, $q = 2$ does not generate a higher match rate than $q = 1$ in expectation.

To make sure that higher-order inclusion-exclusion terms do not undo the preceding bound on collisions, we have to show that a third offer to the focal applicant is not arbitrarily likely once two offers have already occurred.

Lemma 3. *For every $q \geq 3$, the probability of receiving three offers satisfies*

$$I_{3,q} \leq I_{1,q} I_{2,q}.$$

By Lemma 3, the probability of a third collision is bounded by the probability of a second collision times the marginal offer probability. This bound is useful because the inclusion-exclusion formula alternates signs: for $q = 3$, we have $\pi_3 = 3I_{1,3} - 3I_{2,3} + I_{3,3}$. The bound on receiving two offers makes $3I_{2,3}$ large, but the positive term $I_{3,3}$ then offsets it. Lemma 3 limits this offset by showing that $I_{3,3} \leq I_{1,3}I_{2,3}$. Since $I_{1,3}$ is small when $A \geq 2F$, the third-order term cannot erase the collision cost that exists pairwise.

The proof for all $q \geq 4$ in Theorem 1 uses the same logic. Then what remains is a comparison in the imbalance A/F; when $A/F \geq 2$, we show the gain in coverage from raising q is already

small, so the collision term is large enough to dominate. This concludes Theorem 1: when there are at least two applicants per firm, $q = 1$ maximizes the match rate. Our method is not specific to the quota comparison, and could be useful for finding bounds on intersections of events generated by random subset choice without replacement. Appendix B.1 gives more technical detail for the preceding bounds, and Appendix B.3 proves them.

6 Conclusion

We capture how congestion naturally emerges in decentralized matching markets via a simple application game: applicants apply to q firms, firms make an offer to an applicant who has applied, and applicants accept one offer. The principal sets the quota to maximize some joint preference over match rate and applicant welfare, and we find in settings with independent preferences that, so long as there is weak over-demand of capacity, the optimal quota (if it exists) is small. This rationalizes small quotas that we see in practice. When firms correlate their actions, the optimal quota (if it exists) weakly decreases; when applicants correlate their actions, the principal might set the quota to be large. We adapt our model for both the many-to-one market and for an environment without the maximal time constraint, and we consider part of our contribution to be the parsimonious model which enables us to study different aspects of the application game.

The stark optimality of small quotas relies on the maximality of the time constraint. Hence, while a quota of one has been implemented in practice, we do not insist that this result be interpreted as prescriptive. Indeed, our results in the absence of time constraints emphasize that more applications can be desirable if they can be sent sequentially. In general, and as is commonplace in studies of matching markets, our results support centralization.

We focus on both match formation (the extensive margin) and match quality (the intensive margin). Our results highlight a tradeoff between those two dimensions: policies can raise surplus conditional on a match while simultaneously reducing the number of matches formed. Making the principal's preferences explicit could lead to more informed policy recommendations.

Future work is in order. One promising angle includes information disclosure via strong signals. Several decentralized matching markets rely on the use of signals to communicate preferences, which in turn affect where applications and offers are sent. Future work can also study optimal strategies under a fixed quota when each side has correlated values; finding a general statement for the optimal quota in that setting would be a worthy result and would be useful for designing decentralized markets in practice. Our large-market model in particular could be tractable for answering these kinds of questions.

References

Ali, S. N. and Shorrer, R. I. (2025). Hedging when applying: Simultaneous search with correlation. *American Economic Review*, 115(2):571–598.

Arnosti, N., Johari, R., and Kanoria, Y. (2021). Managing congestion in matching markets. *Manufacturing & Service Operations Management*, 23(3):620–636.

Ashlagi, I. and Nikzad, A. (2020). What matters in school choice tie-breaking? how competition guides design. *Journal of Economic Theory*, 190:105120.

Avery, C., Kocks, G., and Pathak, P. A. (2025). The Algorithm Advantage: Ranked Application Systems Outperform Decentralized and Common Applications in Boston and Beyond.

Che, Y.-K. and Koh, Y. (2016). Decentralized college admissions. *Journal of Political Economy*, 124(5):1295–1338.

Coles, P., Kushnir, A., and Niederle, M. (2013). Preference Signaling in Matching Markets. *American Economic Journal: Microeconomics*, 5(2):99–134.

Doğan, B. and Erdil, A. (2025). Widening access in university admissions. *Journal of Economic Theory*, 224:105973.

Ferdowsian, A., Niederle, M., and Yariv, L. (2025). Strategic Decentralized Matching: The Effects of Information Frictions.

Fernandez, M. A., Rudov, K., and Yariv, L. (2021). Centralized Matching with Incomplete Information.

He, Y. and Magnac, T. (2018). A Pigouvian Approach to Congestion in Matching Markets. *IZA Discussion Papers*. Number: 11967 Publisher: Institute of Labor Economics (IZA).

Horton, Vasserman, and Watt (2024). Reducing congestion in labor markets: A case study in simple market design. *Working Paper*.

Kanoria, Y. and Saban, D. (2021). Facilitating the search for partners on matching platforms. *Management Science*, 67(10):5990–6029.

Roth, A. and Xing, X. (1994). Jumping the gun: Imperfections and institutions related to the timing of market transactions. *American Economic Review*, 84(4):992–1044.

Roth, A. E., Sönmez, T., and Ünver, M. U. (2004). Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488.

A Proposition 1

Proof. We continue the proof from the body. Recall the utility conditional on a match is higher under $q = 1$ than $q = F$, so it remains to compare the match rate. We want to show

$$F\left[1 - \left(1 - \frac{1}{F}\right)^{A}\right] \geq A\left[1 - \left(1 - \frac{1}{A}\right)^{F}\right]. \tag{13}$$

Let us define

$$C_A := F - A\left[1 - \left(1 - \frac{1}{A}\right)^F\right] = \sum_{i=1}^{F-1}\left[1 - \left(1 - \frac{1}{A}\right)^i\right] \tag{14}$$

where the second equality in (14) follows from the formula for finite geometric series: $A\left[1 - \left(1 - \frac{1}{A}\right)^F\right] = \sum_{i=0}^{F-1}\left(1 - \frac{1}{A}\right)^i$, and we elide $i = 0$ in the summation above because it evaluates to zero. Note that the target (13) then reduces to showing that

$$C_A \geq F\left(1 - \frac{1}{F}\right)^A. \tag{15}$$

First, observe that (15) holds at equality when $A = F$. We now will show that C_A cannot decrease enough by adding applicants so that the inequality fails to hold.

Define a function $f_i(x) = x\left(1 - (1 - \frac{1}{x})^i\right)$ for $x \geq 1$. Note this function is increasing.[19] Thus $f_i(A+1) \geq f_i(A)$ and so

$$(A+1)\left[1 - \left(1 - \frac{1}{A+1}\right)^i\right] \geq A\left[1 - \left(1 - \frac{1}{A}\right)^i\right].$$

Take sums to recreate the form of (15), and noting that $A \geq F$, we have $\frac{A}{A+1} \geq \frac{F-1}{F}$, and so $C_{A+1} \geq \frac{F-1}{F}C_A$. Then

$$\begin{aligned} C_{A+1} &\geq \frac{F-1}{F}C_A \\ &\geq \frac{F-1}{F}F\left(1 - \frac{1}{F}\right)^A \\ &= F\left(1 - \frac{1}{F}\right)^{A+1} \end{aligned}$$

and so the claim holds by induction. □

B Proof of Theorem 1

B.1 Finite-Market Bounds

We first collect the finite-market bounds used in the proof. We give more technical foundations than in Section 5 and defer the complete technical proofs to Subsection B.3.

Fix a quota q and a focal applicant i. Conditional on the set of firms to which i applies, relabel those firms $1, ..., q$, and let O_j be the event that the j-th firm makes an offer to i. Let $\pi_q = \Pr(\bigcup_{j=1}^q O_j)$ be the probability that the focal applicant receives at least one offer; since an applicant is matched if and only if she receives at least one offer, we have by symmetry that $m_q = (A/F)\pi_q$.

[19] Letting $y = 1 - \frac{1}{x}$, the first-order condition is $f_i'(x) = (1-y)\left(1 + y + \cdots + y^{i-1} - iy^{i-1}\right)$ which is positive everywhere because $0 \leq y \leq 1$.

Restating the definition from Section 5, for each $r \leq q$, let $I_{r,q} = \Pr(O_1 \cap \cdots \cap O_r)$ denote the probability that any particular fixed set of r firms in the focal applicant's portfolio all make offers to her. Here, we lighten notation; write $a_q = I_{1,q}$, $b_q = I_{2,q}$, $c_q = I_{3,q}$.

Lemma 4. *For every quota q, the probability that a fixed firm in the focal applicant's portfolio makes her an offer is*

$$a_q = \frac{F}{Aq}\left[1 - \left(1 - \frac{q}{F}\right)^A\right].$$

We next introduce auxiliary quantities used to bound the finite-market dependence among offer events. Let $n = A - 1$ be the number of competitors faced by the focal applicant, and at times we will write $p_q = q/F$ to be the probability that any particular competitor applies to a given firm in her portfolio.

For $x \in [0, 1]$, interpret x as the focal applicant's percentile in a fixed firm's ranking, so a randomly selected competitor from the firm's pool of applicants is ranked below her with probability x. A given competitor fails to get an offer over the focal applicant either because the competitor does not apply (which happens with probability $1 - p_q$) or because the competitor applies but is ranked below her (which happens with probability $p_q x$). We define $\nu_q(x) = 1 - p_q + p_q x$. Conditional on percentile x, the probability that none of the n competitors defeats the focal applicant is $\nu_q(x)^n$, and so we also have the equality $a_q = \int_0^1 \nu_q(x)^n \, dx$.

The bounds isolate a focal competitor and treat the remaining $n-1$ competitors symmetrically. Define $\omega_q = \int_0^1 \nu_q(x)^{n-1} \, dx$ as the marginal offer probability in the corresponding calculation with one competitor removed. Also define $\kappa_q = \int_0^1 (1-x)\nu_q(x)^{n-1} \, dx$; conditional on the focal competitor applying, κ_q is the probability that this competitor outranks the focal applicant while all remaining competitors fail to defeat her. So κ_q measures the probability that a designated competitor is pivotal.

Lemma 5. *For every quota q, the marginal offer probability with one competitor removed, and the probability that this competitor outranks the focal applicant while all remaining competitors fail to defeat her, satisfy*

$$\omega_q = \frac{1 - (1 - p_q)^{A-1}}{(A-1)p_q} \quad \textit{and} \quad \kappa_q = (\omega_q - a_q)/p_q.$$

Further,

$$\omega_q \leq \frac{A}{A-1} a_q \leq \frac{F}{q(A-1)} \quad \textit{and} \quad \kappa_q \leq \frac{a_q}{p_q(A-1)} \leq \frac{F^2}{q^2 A(A-1)}.$$

The bounds in Lemma 5 follow directly from the relationship between the calculations with and without the focal competitor.

We now compare the probability of receiving offers from two fixed firms with the corresponding probability under independence. If the two offer events were independent, their joint probability would be a_q^2. However, in the finite market, each competitor selects q distinct firms without replacement, and a competitor who "spends" one application at the first focal firm is then slightly less likely to apply to the second focal firm.

We write a term which captures the (pairwise) departure from independent application decisions. First, let

$$\Delta_{2,q} = (q/F)^2 - \frac{q(q-1)}{F(F-1)} = \frac{q(F-q)}{F^2(F-1)}.$$

The first term is the probability that the competitor would apply to both firms if the two application decisions were independent; the second is the actual probability when she selects a portfolio of q distinct firms. For this departure to affect whether the focal applicant receives both offers, the competitor must also be pivotal at both firms. The probability of being pivotal at one firm, conditional on applying, is κ_q, so the corresponding term is κ_q^2; since there are $n = A - 1$ potential competitors, the total first-order correction is $n\Delta_{2,q}\kappa_q^2$. This gives intuition for the following bound.

Lemma 6. *For every $q \geq 2$, the probability of receiving two offers satisfies*

$$a_q^2 - n\Delta_{2,q}\kappa_q^2 \leq b_q \leq a_q^2.$$

Consequently,

$$a_q^2 - \delta_q \leq b_q \leq a_q^2 \quad \textit{where} \quad \delta_q = \frac{F^2}{q^3(A-1)A^2}$$

.

The wedge δ_q will be central, and has the interpretation that receiving an offer from one fixed firm is weakly bad news about receiving an offer from another fixed firm — which, broadly, bodes positively toward natural mitigation of congestion. Indeed, an offer from the first focal firm indicates that the focal applicant effectively faced relatively little competition there; because competitors allocate their applications without replacement, this implies relatively more competition at the other firm. The lower bound controls the magnitude of this dependence.[20]

The two sides of the inequality derived in Lemma 6 will serve different purposes. When an intersection enters our inclusion-exclusion formula with a negative sign, its *lower* bound provides an *upper* bound on the match probability under a larger quota; when $q = 2$ is used as the comparison benchmark, the *upper* bound $b_2 \leq a_2^2$ provides a *lower* bound on the match probability.[21] Our last bound makes sure the probability of three offers cannot be more than the preceding pair bound times the marginal offer probability.

Lemma 7. *For every $q \geq 3$, the probability of receiving three offers satisfies*

$$c_q \leq a_q b_q.$$

[20] In accordance with Assumption 1, it is easy to see that for a fixed quota and a fixed market ratio A/F, the wedge converges to zero as the market grows.

[21] We believe this approach will serve well for a finite-market proof of Theorem 2. There, we need bounds for intersections of offers of higher order. The next result bounds these probabilities; the "intersection" probability cannot increase by more than the marginal probability of an individual offer.

B.2 Comparison of Match Rates

Our proof approach is to compare the match rate under $q = 1$ with the match rates under quotas $q = 2$, $q = 3$, and $q \geq 4$. It will be helpful to collect all of the formulas:

$$\pi_1 = a_1 = \frac{F}{A}\left[1 - \left(1 - \frac{1}{F}\right)^A\right], \qquad m_1 = 1 - \left(1 - \frac{1}{F}\right)^A,$$

$$\pi_2 = 2a_2 - b_2, \qquad m_2 = \frac{A}{F}(2a_2 - b_2) = 1 - \left(1 - \frac{2}{F}\right)^A - \frac{A}{F}b_2,$$

$$\pi_3 = 3a_3 - 3b_3 + c_3, \qquad m_3 = \frac{A}{F}(3a_3 - 3b_3 + c_3) = 1 - \left(1 - \frac{3}{F}\right)^A - \frac{A}{F}(3b_3 - c_3).$$

For $q \geq 4$, the first three terms of the inclusion-exclusion formula give

$$\pi_q \leq qa_q - \binom{q}{2}b_q + \binom{q}{3}c_q,$$
$$m_q \leq \frac{A}{F}\left(qa_q - \binom{q}{2}b_q + \binom{q}{3}c_q\right).$$

Lemma 8. *For every $F \geq 2$ and $A \geq 2F$, $m_1 \geq m_2$.*

Proof. First suppose $F = 2$. Under $q = 2$, every applicant applies to both firms, and the claim follows from Proposition 1.

Now suppose $F \geq 3$. Define two terms: $u = (1 - 1/F)^A$ and $v = (1 - 2/F)^A$. The exact inclusion-exclusion formula gives $\pi_2 = 2a_2 - b_2$, where by Lemma 4 we have $2a_2 = (F/A)(1 - v)$. We can then write the difference in match rates as

$$m_1 - m_2 = (1 - u) - \frac{A}{F}(2a_2 - b_2)$$
$$= v - u + \frac{A}{F}b_2.$$

It is then enough to show that $b_2 \geq (F/A)(u - v)$. We can use the lower bound from Lemma 6 to derive a sufficient condition:

$$a_2^2 - \delta_2 \geq \frac{F}{A}(u - v). \tag{16}$$

The rest of the proof amounts to showing that this inequality holds. We present one such way of showing this.

First, write a_2^2 in terms of v, multiply through (16) by A^2/F^2, and move all the terms to one

side; define

$$\Phi_F(A) = \frac{1}{4}(1-v)^2 - \frac{A}{F}(u-v) - \frac{1}{8(A-1)}. \tag{17}$$

Thus showing the target inequality (16) is equivalent to showing that $\Phi_F(A) \geq 0$.

Let $\ell_1 = -\log(1-1/F)$ and $\ell_2 = -\log(1-2/F)$, so then $u = e^{-A\ell_1}$ and $v = e^{-A\ell_2}$. Taking the first-order condition of (17) and then grouping terms gives

$$\begin{aligned}\Phi'_F(A) = \frac{1}{2}(1-v)v\ell_2 + \frac{1}{8(A-1)^2} \\ + \frac{1}{F}\left[u(A\ell_1 - 1) - v(A\ell_2 - 1)\right].\end{aligned}$$

The first line is nonnegative; note $u, v \in [0,1]$. For the term in brackets in the second line, define the function $g(x) = e^{-Ax}(Ax-1)$, and note it is decreasing on $[\ell_1, \ell_2]$.[22] Then because $\ell_2 > \ell_1$, it follows that $g(\ell_1) \geq g(\ell_2)$, or equivalently $u(A\ell_1 - 1) \geq v(A\ell_2 - 1)$. Hence $\Phi'_F(A) \geq 0$ for every $A \geq 2F$.

Confirming that $\Phi_F(2F) \geq 0$ amounts to algebra. Because Φ_F is weakly increasing on $[2F, \infty)$, we obtain $\Phi_F(A) \geq 0$ for every $A \geq 2F$, which proves the claim. □

Lemma 9. *For every $F \geq 3$ and $A \geq 2F$, $m_1 \geq m_3$.*

Proof. The inclusion-exclusion formula for $q = 3$ gives the offer probability $\pi_3 = 3a_3 - 3b_3 + c_3$. Using again the equality from Lemma 4, we can write the difference in match rates in the following way:

$$m_1 - m_3 = \left(1 - \frac{3}{F}\right)^A - \left(1 - \frac{1}{F}\right)^A + \frac{A}{F}(3b_3 - c_3)$$

Because the first term is positive, the dominance of $q = 1$ follows if we can show the second term is smaller in magnitude than the third term. Because $(1-1/F)^A \leq e^{-A/F}$, it is sufficient to prove that

$$3b_3 - c_3 \geq \frac{F}{A}e^{-A/F}. \tag{18}$$

We can use the finite-market bounds from Lemmas 6 and 7 to write a lower bound for the left-hand side of (18):

$$3b_3 - c_3 \geq (3 - a_3)\left(a_3^2 - \frac{F^2}{27(A-1)A^2}\right). \tag{19}$$

Now, $A \geq 2F$ implies that $a_3 \leq \frac{F}{3A} \leq 1/6$, so that coefficient term satisfies $3 - a_3 \geq 17/6$. Also

[22] Since $A \geq 2F$ and $\ell_1 \geq 1/F$, we have $A\ell_1 \geq 2$, and $g'(x) = Ae^{-Ax}(2 - Ax) \leq 0$ in the domain.

observe that $A - 1 \geq 5$. We can bound the other terms then as follows:

$$a_3^2 \geq \frac{F^2}{9A^2}\left(1 - e^{-3A/F}\right)^2, \qquad \frac{F^2}{27(A-1)A^2} \leq \frac{F^2}{135A^2},$$

A computation will complete the proof. Substituting these bounds directly into (19) gives

$$\begin{aligned} 3b_3 - c_3 &\geq C\frac{F^2}{A^2}, \\ C &= \frac{17}{6}\left[\frac{(1-e^{-6})^2}{9} - \frac{1}{135}\right] \approx 0.2923. \end{aligned}$$

Note this C is just a constant (and is evaluated as such). We now compare it with the right-hand side of (18); it remains to show that $C \geq \frac{A}{F}e^{-A/F}$. The constant $2e^{-2}$ is less than C but upper bounds $\frac{A}{F}e^{-A/F}$, so the result follows. □

Lemma 10. *For every $F \geq 4$, every $A \geq 2F$, and every $q \in \{4, \ldots, F\}$, $m_1 \geq m_q$.*

Proof. We will compare m_1 to m_q by writing a lower bound for m_1 and an upper bound of m_q, and making a computation as we did in the proof of Lemma 9. We can use the inclusion-exclusion formula to upper bound the probability of receiving an offer under $q \geq 4$:

$$\pi_q \leq qa_q - \binom{q}{2}b_q + \binom{q}{3}c_q.$$

By Lemma 7, we have $c_q \leq a_q b_q$, and note that $\binom{q}{3} = \binom{q}{2}\frac{q-2}{3}$. Hence we can rewrite the upper bound as

$$\pi_q \leq qa_q - \binom{q}{2}\left(1 - \frac{q-2}{3}a_q\right)b_q. \tag{20}$$

Since $a_q \leq \frac{F}{Aq} \leq \frac{1}{2q}$, we can bound that new term on the right-hand side of (20), as $1 - \frac{q-2}{3}a_q \geq \frac{5}{6}$. Then using Lemma 6, we can rearrange into the following:

$$\pi_q \leq qa_q - \frac{5}{6}\binom{q}{2}a_q^2 + \frac{5}{6}\binom{q}{2}\frac{F^2}{q^3(A-1)A^2}. \tag{21}$$

We bound the three terms separately. For the first term, we have $qa_q \leq F/A$. For the second term, $(1-x)^A \leq e^{-Ax}$ and $q \geq 4$ imply that

$$\begin{aligned} a_q &= \frac{F}{Aq}\left[1 - \left(1 - \frac{q}{F}\right)^A\right] \\ &\geq \frac{F}{Aq}\left(1 - e^{-Aq/F}\right) \\ &\geq \frac{F}{Aq}\left(1 - e^{-4A/F}\right). \end{aligned}$$

And then plugging in that new bound into the second term and using that $q \geq 4$:

$$\begin{aligned}\binom{q}{2} a_q^2 &\geq \frac{q-1}{2q}\frac{F^2}{A^2}\left(1-e^{-4A/F}\right)^2 \\ &\geq \frac{3}{8}\frac{F^2}{A^2}\left(1-e^{-4A/F}\right)^2. \end{aligned} \tag{22}$$

For the third term, $\frac{q-1}{2q^2} \leq 3/32$ for $q \geq 4$, and $A \geq 8$ implies $1/(A-1) \leq 1/7$. We merely evaluate to get the bound. Hence

$$\frac{5}{6}\binom{q}{2}\frac{F^2}{q^3(A-1)A^2} \leq \frac{5}{448}\frac{F^2}{A^2}. \tag{23}$$

Substituting (22) and (23) (as well as $qa_q \leq F/A$) into (21) yields

$$\pi_q \leq \frac{F}{A} - \frac{5}{16}\frac{F^2}{A^2}\left(1-e^{-4A/F}\right)^2 + \frac{5}{448}\frac{F^2}{A^2}. \tag{24}$$

Now the objective is to show that $m_1 - m_q$ is positive using this upper bound we have just found. First, multiply (24) through by A/F to get m_q:

$$m_q = \frac{A}{F}\pi_q \leq 1 - \frac{F}{A}\left[\frac{5}{16}\left(1-e^{-4A/F}\right)^2 - \frac{5}{448}\right].$$

Now toward the comparison, observe that a suitable lower bound for m_1 is simply

$$m_1 = 1 - \left(1-\frac{1}{F}\right)^A \geq 1 - e^{-A/F},$$

and so, after writing $m_1 - m_q$ and rearranging terms, it is sufficient to show that

$$\frac{A}{F}e^{-A/F} \leq \frac{5}{16}\left(1-e^{-4A/F}\right)^2 - \frac{5}{448}. \tag{25}$$

The left-hand side of (25) is decreasing in A/F for $A/F \geq 2$, while the right-hand side is increasing. So it is enough to evaluate both sides at $A/F = 2$. At this boundary,

$$2e^{-2} \approx 0.271 < 0.301 \approx \frac{5}{16}(1-e^{-8})^2 - \frac{5}{448}.$$

Thus (25) holds whenever $A/F \geq 2$, and the claim follows. □

B.3 Proofs of the Finite-Market Bounds

We now prove Lemmas 4–7. First, we derive a preparatory result:

Lemma 11. *If $Z \sim \text{Bin}(N, p)$ with $N \in \mathbb{Z}_{\geq 0}$ and $p \in (0, 1]$, then*

$$\mathbb{E}\left[\frac{1}{1+Z}\right] = \frac{1-(1-p)^{N+1}}{(N+1)p}. \tag{26}$$

Proof. The key steps are using $\frac{1}{z+1}\binom{N}{z} = \frac{1}{N+1}\binom{N+1}{z+1}$ and the binomial theorem:

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{1+Z}\right] &= \sum_{z=0}^{N} \frac{1}{z+1}\binom{N}{z} p^z (1-p)^{N-z} \\
&= \frac{1}{(N+1)p} \sum_{w=1}^{N+1} \binom{N+1}{w} p^w (1-p)^{N+1-w} \\
&= \frac{1}{(N+1)p} \left[\sum_{w=0}^{N+1} \binom{N+1}{w} p^w (1-p)^{N+1-w} - (1-p)^{N+1}\right] \\
&= \frac{1}{(N+1)p} \left[\big(p + (1-p)\big)^{N+1} - (1-p)^{N+1}\right] \\
&= \frac{1}{(N+1)p} \left[1 - (1-p)^{N+1}\right] \\
&= \frac{1-(1-p)^{N+1}}{(N+1)p}.
\end{aligned} \tag{27}$$

□

Proof of Lemma 4. Conditional on the focal applicant applying to a fixed firm, each of the other $A-1$ applicants applies there with probability $p_q = q/F$ (independently across applicants). Thus the number of competitors is $Z \sim \text{Bin}(A-1, p_q)$, and conditional on $Z = z$, the focal applicant is highest ranked in the application pool with probability $1/(z+1)$. By Lemma 11,

$$a_q = \mathbb{E}\left[\frac{1}{1+Z}\right] = \frac{1-(1-p_q)^A}{Ap_q} = \frac{F}{Aq}\left[1 - \left(1 - \frac{q}{F}\right)^A\right].$$

□

Proof of Lemma 5. Integrating $\nu_q(x)^{n-1}$ directly and substituting $n = A-1$ gives

$$\omega_q = \int_0^1 \nu_q(x)^{n-1}\, dx = \frac{1-(1-p_q)^n}{np_q} = \frac{1-(1-p_q)^{A-1}}{(A-1)p_q}.$$

Then because $\nu_q(x) = 1 - p_q(1-x)$, we can also write $p_q(1-x) = 1 - \nu_q(x)$, and so

$$p_q\kappa_q = \int_0^1 \left[\nu_q(x)^{n-1} - \nu_q(x)^n\right] dx = \omega_q - a_q,$$

which establishes the target $\kappa_q = (\omega_q - a_q)/p_q$.

That $\omega_q \leq \frac{A}{A-1} a_q$ trivially follows from $1-(1-p_q)^{A-1} \leq 1-(1-p_q)^A$. As a result, $\omega_q - a_q \leq$

$a_q/(A-1)$ and $\kappa_q \leq a_q/[p_q(A-1)]$. Then using that, $a_q \leq 1/(Ap_q) = F/(Aq)$, we can write the final bounds

$$\omega_q \leq \frac{1}{p_q(A-1)} = \frac{F}{q(A-1)},$$
$$\kappa_q \leq \frac{1}{Ap_q^2(A-1)} = \frac{F^2}{q^2A(A-1)}.$$

□

Proof of Lemma 6. Fix two distinct firms to which the focal applicant has applied, and for each competitor $\ell \in \{1, \ldots, n\}$, let X_ℓ and Y_ℓ indicate whether that competitor applies to the first and second firm, respectively. Set $Z_1 = \sum_{\ell=1}^n X_\ell$ and $Z_2 = \sum_{\ell=1}^n Y_\ell$.

Conditional on (Z_1, Z_2), rankings at the two firms are independent and the focal applicant is top ranked with probability $1/(1+Z_j)$ at firm j. Therefore the probability of getting both offers is

$$b_q = \mathbb{E}\left[\frac{1}{(1+Z_1)(1+Z_2)}\right].$$

We will define two functions which consider a single competitor. First, define $G_q(x,y) = \mathbb{E}[x^{X_\ell}y^{Y_\ell}]$ which will naturally emerge from analyzing b_q; it is the joint probability-generating function for the two applications. Then let $H_q(x,y) = \nu_q(x)\nu_q(y)$ be the same but with the independence assumption; under independence, the joint probability of applying to both would be p_q^2. But the actual probability is smaller: $\frac{q(q-1)}{F(F-1)}$. Their difference is

$$\begin{aligned}\Delta_{2,q} &= p_q^2 - \frac{q(q-1)}{F(F-1)} \\ &= \left(\frac{q}{F}\right)^2 - \frac{q(q-1)}{F(F-1)} \\ &= \frac{q(F-q)}{F^2(F-1)}\end{aligned}$$

So $\Pr(X_\ell = 1, Y_\ell = 1) = \frac{q(q-1)}{F(F-1)} = p_q^2 - \Delta_{2,q}$. Now, since $\Pr(X_\ell = 1) = \Pr(Y_\ell = 1) = p_q$, we also can write

$$\begin{aligned}\Pr(X_\ell = 1, Y_\ell = 0) &= p_q - \Pr(X_\ell = 1, Y_\ell = 1) \\ &= p_q(1-p_q) + \Delta_{2,q},\end{aligned}$$

and symmetrically,

$$\Pr(X_\ell = 0, Y_\ell = 1) = p_q(1-p_q) + \Delta_{2,q}.$$

The probability of no applications is then simply $(1-p_q)^2 - \Delta_{2,q}$. Using all of these terms, we can

write out $G_q(x, y)$ fully:

$$\begin{aligned}
G_q(x,y) &= \mathbb{E}\left[x^{X_\ell} y^{Y_\ell}\right] \\
&= \Pr(X_\ell = 0, Y_\ell = 0) + \Pr(X_\ell = 1, Y_\ell = 0)x \\
&\quad + \Pr(X_\ell = 0, Y_\ell = 1)y + \Pr(X_\ell = 1, Y_\ell = 1)xy \\
&= \left[(1-p_q)^2 - \Delta_{2,q}\right] \\
&\quad + \left[p_q(1-p_q) + \Delta_{2,q}\right] x \\
&\quad + \left[p_q(1-p_q) + \Delta_{2,q}\right] y \\
&\quad + \left[p_q^2 - \Delta_{2,q}\right] xy.
\end{aligned}$$

Meanwhile, we can write $H_q(x, y)$ in the following way:

$$\begin{aligned}
H_q(x,y) &= \nu_q(x)\nu_q(y) \\
&= (1 - p_q + p_q x)(1 - p_q + p_q y) \\
&= (1-p_q)^2 + p_q(1-p_q)x + p_q(1-p_q)y + p_q^2 xy.
\end{aligned}$$

Note the $\Delta_{2,q}$ terms emerge in the expected places: the events of neither or both applications are slightly more likely under the independence assumption than without. Then the difference between the above equations is

$$\begin{aligned}
H_q(x,y) - G_q(x,y) &= \Delta_{2,q} - \Delta_{2,q}x - \Delta_{2,q}y + \Delta_{2,q}xy \\
&= \Delta_{2,q}(1 - x - y + xy) \\
&= \Delta_{2,q}(1-x)(1-y).
\end{aligned}$$

And so we have the key equality:

$$G_q(x,y) = H_q(x,y) - \Delta_{2,q}(1-x)(1-y).$$

Indeed, this preceding equality captures the cost of sampling without replacement. A competitor who spends one of her q applications at the first firm is then slightly less likely to apply to the second.

Because competitors apply independently, we can consider all n competitors simultaneously:

$$\begin{aligned}
\mathbb{E}\left[x^{Z_1} y^{Z_2}\right] &= \mathbb{E}\left[x^{\Sigma_\ell X_\ell} y^{\Sigma_\ell Y_\ell}\right] \\
&= \mathbb{E}\left[\prod_{\ell=1}^{n} x^{X_\ell} y^{Y_\ell}\right] \\
&= \prod_{\ell=1}^{n} \mathbb{E}\left[x^{X_\ell} y^{Y_\ell}\right] \\
&= G_q(x,y)^n.
\end{aligned} \tag{28}$$

We now work toward a bound for b_q. First, we want to write it as an integral; use $\frac{1}{1+z} = \int_0^1 t^z\,dt$ twice to get

$$\begin{aligned} b_q &= \mathbb{E}\left[\int_0^1\int_0^1 x^{Z_1}y^{Z_2}\,dx\,dy\right] \\ &= \int_0^1\int_0^1 \mathbb{E}\left[x^{Z_1}y^{Z_2}\right]\,dx\,dy \\ &= \int_0^1\int_0^1 G_q(x,y)^n\,dx\,dy. \end{aligned}$$

Swapping expectation and integral in the second step is permitted because the integrand $x^{Z_1}y^{Z_2}$ is nonnegative on $[0,1]^2$.

Now, because $\Delta_{2,q} \geq 0$, we know $G_q \leq H_q$ pointwise. It follows that

$$b_q \leq \int_0^1\int_0^1 H_q(x,y)^n\,dx\,dy = a_q^2. \tag{29}$$

For the lower bound, if $q = F$, then $\Delta_{2,q} = 0$ and equality holds in (29). So suppose $q < F$. Then $H_q > 0$ on $[0,1]^2$, and $0 \leq \Delta_{2,q}\frac{(1-x)(1-y)}{H_q(x,y)} \leq 1$. Thus we have

$$\begin{aligned} G_q(x,y)^n &= H_q(x,y)^n\left[1 - \frac{\Delta_{2,q}(1-x)(1-y)}{H_q(x,y)}\right]^n \\ &\geq H_q(x,y)^n\left[1 - n\frac{\Delta_{2,q}(1-x)(1-y)}{H_q(x,y)}\right] \\ &= H_q(x,y)^n - n\Delta_{2,q}(1-x)(1-y)H_q(x,y)^{n-1} \end{aligned}$$

where the bound in the second step followed from using Bernoulli's inequality: $(1-u)^n \geq 1-nu$ for $u \in [0,1]$ and integer n. Integrating the preceding inequality gives

$$\begin{aligned} b_q &\geq \int_0^1\int_0^1 H_q(x,y)^n\,dx\,dy \\ &\qquad - n\Delta_{2,q}\int_0^1\int_0^1 (1-x)(1-y)H_q(x,y)^{n-1}\,dx\,dy \\ &= a_q^2 - n\Delta_{2,q}\left[\int_0^1 (1-x)\nu_q(x)^{n-1}\,dx\right]^2 \\ &= a_q^2 - n\Delta_{2,q}\kappa_q^2. \end{aligned}$$

And Lemma 5 implies

$$\begin{aligned} n\Delta_{2,q}\kappa_q^2 &\leq (A-1)\frac{q(F-q)}{F^2(F-1)}\left(\frac{F^2}{q^2A(A-1)}\right)^2 \\ &= \frac{F^2(F-q)}{q^3(F-1)(A-1)A^2} \leq \frac{F^2}{q^3(A-1)A^2} = \delta_q. \end{aligned}$$

□

Proof of Lemma 7. This proof closely follows that of the preceding claim. Fix three distinct firms j_1, j_2, j_3 to which the focal applicant applied. For each other applicant $\ell \neq i$ define the application indicators X_ℓ, Y_ℓ, Z_ℓ, and write the corresponding application counts

$$S_1 := \sum_{\ell \neq i} X_\ell, \qquad S_2 := \sum_{\ell \neq i} Y_\ell, \qquad S_3 := \sum_{\ell \neq i} Z_\ell.$$

Conditional on (S_1, S_2, S_3), firms randomize independently, so

$$\begin{aligned} c_q &= \mathbb{E}\left[\frac{1}{(1+S_1)(1+S_2)(1+S_3)}\right] \\ &= \mathbb{E}\left[\int_0^1\int_0^1\int_0^1 x^{S_1} y^{S_2} z^{S_3}\, dx\, dy\, dz\right] \\ &= \int_0^1\int_0^1\int_0^1 \mathbb{E}\big[x^{S_1} y^{S_2} z^{S_3}\big]\, dx\, dy\, dz, \end{aligned}$$

(This follows Lemma 6 for three variables.) Now let $G(x, y, z) := \mathbb{E}[x^{X_\ell} y^{Y_\ell} z^{Z_\ell}]$ be the single-applicant probability-generating function, so $\mathbb{E}[x^{S_1} y^{S_2} z^{S_3}] = G(x, y, z)^n$. Then write for the pair probability $G_{xy}(x, y) := G(x, y, 1)$ and single probability $G_z(z) := G(1, 1, z)$; the subscripts denote which terms are "active." We will now show that for all $(x, y, z) \in [0, 1]^3$,

$$G(x, y, z) \;\leq\; G_{xy}(x, y)\, G_z(z). \tag{30}$$

With slight abuse of notation, let $p_t = \frac{\binom{F-3}{q-t}}{\binom{F}{q}}$ be the probability that a given other applicant chooses a prespecified set of t firms among $\{j_1, j_2, j_3\}$ and chooses the remaining $q - t$ firms from the other $F - 3$. Expanding $G(x, y, z)$ by conditioning on how many of the set $\{j_1, j_2, j_3\}$ are chosen yields

$$G(x, y, z) = p_0 + p_1(x + y + z) + p_2(xy + xz + yz) + p_3 xyz,$$

Then we can write $G_{xy}(x, y)$ and $G_z(z)$ as

$$\begin{aligned} G_{xy}(x, y) &= p_0 + p_1(x + y + 1) + p_2(xy + x + y) + p_3 xy, \\ G_z(z) &= p_0 + p_1(2 + z) + p_2(1 + 2z) + p_3 z. \end{aligned}$$

Closed forms follow from the binomial form of p_t. Writing out each p_t for $t \in \{0, 3\}$, we get

$$p_0 = \frac{(F-q)(F-q-1)(F-q-2)}{F(F-1)(F-2)}, \qquad p_1 = \frac{q(F-q)(F-q-1)}{F(F-1)(F-2)},$$

$$p_2 = \frac{q(q-1)(F-q)}{F(F-1)(F-2)}, \qquad p_3 = \frac{q(q-1)(q-2)}{F(F-1)(F-2)}.$$

Toward the bound, set

$$\Delta(x,y,z) := G_{xy}(x,y)G_z(z) - G(x,y,z).$$

Using $G(x,y,z) - G(x,y,1) = (z-1)\big(p_1 + p_2(x+y) + p_3 xy\big)$ and $G_z(z) - 1 = (z-1)\big(p_1 + 2p_2 + p_3\big)$, we can rewrite

$$\begin{aligned}\Delta(x,y,z) &= G_{xy}(x,y)\big(G_z(z) - 1\big) - \big(G(x,y,z) - G_{xy}(x,y)\big) \\ &= (z-1)\Big[\big(p_1 + 2p_2 + p_3\big)G_{xy}(x,y) - \big(p_1 + p_2(x+y) + p_3 xy\big)\Big].\end{aligned}$$

Substituting the expression for G_{xy} and the closed forms for p_0, p_1, p_2, p_3 and simplifying (collecting terms in $(1-x)$ and $(1-y)$) yields

$$\begin{aligned}\Delta(x,y,z) = \frac{q(F-q)}{F^2(F-1)(F-2)}\,(1-z) \\ \cdot \Big((F-2)\big((1-x)+(1-y)\big) - 2(q-1)(1-x)(1-y)\Big) \geq 0.\end{aligned}$$

Proving positivity amounts to proving the large term within brackets is weakly positive, which follows from $x, y \in [0,1]$. The coefficient $\frac{q(F-q)}{F^2(F-1)(F-2)}(1-z) \geq 0$. This proves (30). Then $G(x,y,z)^n \leq (G_{xy}(x,y)G_z(z))^n$ on $[0,1]^3$, so

$$\begin{aligned}c_q &= \int_0^1\int_0^1\int_0^1 G(x,y,z)^n\,dx\,dy\,dz \\ &\leq \int_0^1\int_0^1\int_0^1 (G_{xy}(x,y)G_z(z))^n\,dx\,dy\,dz \\ &= \Big(\int_0^1\int_0^1 G_{xy}(x,y)^n\,dx\,dy\Big)\Big(\int_0^1 G_z(z)^n\,dz\Big).\end{aligned}$$

The first factor is exactly b_q and the second factor is exactly a_q, hence we have shown the claim.

□

C Proof of Theorem 2

Proof. Recall our definition from the main text: the probability that any particular application leads to an offer is

$$s_q = \frac{1 - e^{-\lambda q}}{\lambda q}.$$

Then again letting π_q denote the probability that a representative applicant matches under quota q, we can write $\pi_q = 1 - (1 - s_q)^q$, and by symmetry, the match rate per firm is $m_q = \lambda \pi_q$. Differentiation reveals s_q is decreasing for $q \geq 1$ and π_q is decreasing for $q \geq 3$; we elide the algebra.

Toward comparing applicant welfare, consider a focal applicant and let R_q be the rank of the firm whose offer she accepts under quota q, and set $R_q = \infty$ if she receives no offer. We will show

stochastic dominance. Because offers are independent,

$$\Pr(R_q \leq r) = 1 - (1 - s_q)^{\min\{r,q\}} \qquad \text{for } r = 1, \ldots, F.$$

For $r = 1, 2$, that $s_3 \geq s_q$ implies

$$\Pr(R_3 \leq r) = 1 - (1 - s_3)^r \geq 1 - (1 - s_q)^r = \Pr(R_q \leq r).$$

For $r \geq 3$,

$$\Pr(R_3 \leq r) = \pi_3 \geq \pi_q \geq \Pr(R_q \leq r).$$

Thus R_3 first-order stochastically dominates R_q (where a lower rank is better). Hence $q = 3$ weakly increases both the match rate and applicant welfare relative to every $q > 3$. Therefore every quota $q > 3$ is dominated by a quota $q \leq 3$. □

D Proof of Theorem 4

Recall that a random vector X is associated if $\mathrm{Cov}(\phi(X), \psi(X)) \geq 0$ for every pair of bounded, coordinatewise nondecreasing functions ϕ and ψ. We state key properties of associated random variables in Lemma 12 which we will use in the proofs.

Lemma 12. *Subvectors and independent unions of associated vectors are associated, coordinatewise nondecreasing transformations preserve association, and if X is associated, so is $-X$. If $Z_1, \ldots, Z_q$ are positively associated random variables taking values in $[0, 1]$, then*

$$\mathbb{E}\left[\prod_{\ell=1}^{q} Z_\ell\right] \geq \prod_{\ell=1}^{q} \mathbb{E}[Z_\ell].$$

The key idea for the next result is that choosing the top q from a set creates a shared cutoff; conditional on that cutoff the selected variables are independent, but unconditionally the cutoff induces positive dependence.

Lemma 13. *Let $(X_f, Y_f)_{f=1}^F$ be i.i.d., and suppose Y_f is stochastically increasing in X_f. Fix a set $S \subseteq \mathcal{F}$ with $|S| = q$. Conditional on S being the set of indices of the q largest values in X, the vector $(Y_f)_{f \in S}$ is associated.*

Proof. If $q = F$ the result is immediate, so suppose $q < F$, and let $\tau = \max_{f \notin S} X_f$. Conditional on $\tau = t$ and on S being the set of firms that rank in the top q for a focal applicant, the pairs $(X_f, Y_f)_{f \in S}$ are independent draws from the distribution of (X, Y) conditional on $X > t$. Now as t increases, the distribution of X conditional on $X > t$ increases in the first-order stochastic order. Since Y is stochastically increasing in X, the distribution of Y conditional on $X > t$ is also stochastically increasing in t. We can then write $Y_f = Q(\tau, V_f)$ where $Q(\cdot)$ is a quantile function and the V_f variables are independent uniform random variables; Q is nondecreasing in both arguments.

The vector $(\tau, (V_f)_{f\in S})$ is associated, and $(Y_f)_{f\in S}$ is a coordinatewise nondecreasing transformation of that vector; by Lemma 12 we have the result. □

D.1 Cross-Side Alignment

Recall that we study the equilibrium where applicants apply to their q-best firms for every α, and firms offer their favorite applicant from those who have applied. So the distribution of applications is independent of α; we study how collision changes because of the preference alignment.

Fix an applicant a and write $X_f = u_{af}$ and $Y_f = u_{fa}$. The first step is to show that Y_f is stochastically increasing in X_f. Note the claim is immediate when $\alpha = 0$ or $\alpha = 1$; now suppose $\alpha \in (0, 1)$. Let f_ε denote the density of the applicant-side idiosyncratic shock. Conditional on match-specific value $v_{af} = v$, the density of X_f is

$$f_{X|v}(x \mid v) = \frac{1}{1-\alpha} f_\varepsilon \left(\frac{x - \alpha v}{1-\alpha} \right)$$

where the term $\frac{x-\alpha v}{1-\alpha}$ is X rearranged for ε. Recall we assumed log-concavity of f_ε; this implies the monotone-likelihood-ratio property in v. To see this directly where the derivative exists, let $\gamma = \log f_\varepsilon$. For $v_2 > v_1$,

$$\frac{\partial}{\partial x} \log \frac{f_{X|v}(x \mid v_2)}{f_{X|v}(x \mid v_1)} = \frac{1}{1-\alpha} \left[\gamma' \left(\frac{x - \alpha v_2}{1-\alpha} \right) - \gamma' \left(\frac{x - \alpha v_1}{1-\alpha} \right) \right] \geq 0.$$

The inequality follows because γ' is nonincreasing. By Bayes' rule we then have that the posterior distribution of v_{af} is stochastically increasing in X_f. Since $Y_f = \alpha v_{af} + (1-\alpha)\varepsilon_{fa}$ is stochastically increasing in v_{af} by the same construction, it follows that Y_f is also stochastically increasing in X_f.

We can think of the applications as a graph where an application corresponds to an edge between the respective applicant and firm nodes. We elide formalism and will consider a realized application graph $\mathcal{G}$. For each applicant a, let S_a denote the applications she has sent. Conditional on $\mathcal{G}$, the definition of S_a is that it contains the q highest (applicant-side) values of applicant a. Lemma 13 gives that $(u_{fa})_{f\in S_a}$ is associated. Because the primitive values composing applicants' utilities are independent, the applicant-specific utility vectors remain independent across applicants conditional on $\mathcal{G}$.

Fix the focal applicant a, and for every $f \in S_a$; recall that O_f indicates that firm f offers to a and that $N_f = |P_f \setminus \{a\}|$ is the number of competing applicants at firm f.

Since $\sigma_f(a \mid P_f, u_f)$ is weakly increasing in u_{fa} and weakly decreasing in every competitor value $u_{fa'}$, it follows that $(\sigma_f(a \mid P_f, u_f))_{f\in S_a}$ is associated. Hence so is $(1 - \sigma_f(a \mid P_f, u_f))_{f\in S_a}$.

Conditional on the realized utilities, firms randomize independently. Therefore

$$\begin{aligned}\Pr_{\alpha}(a \text{ receives no offer} \mid \mathcal{G}) &= \mathbb{E}_{\alpha}\left[\prod_{f \in S_a}\left(1-\sigma_f(a \mid P_f, u_f)\right) \,\middle|\, \mathcal{G}\right] \\ &\geq \prod_{f \in S_a} \mathbb{E}_{\alpha}\left[1-\sigma_f(a \mid P_f, u_f) \,\middle|\, \mathcal{G}\right] \\ &= \prod_{f \in S_a} \frac{N_f}{N_f+1} \\ &= \Pr_{0}(a \text{ receives no offer} \mid \mathcal{G}).\end{aligned}$$

The third line follows because the applicants in P_f are exchangeable conditional on $\mathcal{G}$ and firm f makes exactly one offer. The last line follows because firm-side values are independent across firms when $\alpha = 0$. Since the distribution of $\mathcal{G}$ is the same for every α, averaging over $\mathcal{G}$ gives $m_q^{\alpha} \leq m_q$. □

D.2 Firm-Side Alignment

This case is simpler than that for cross-side correlation. Applicant utilities do not depend on β^f, so applicants apply to their top q firms; the distribution of the application graph is the same as in the case of cross-side correlation and the baseline independent game. We now show firms are more likely to collide.

Fix a realized application graph $\mathcal{G}$. The vector $(u_{fa'})_{f \in S_{a'}}$ is associated because its coordinates are coordinatewise nondecreasing functions of the independent variables $\theta_{a'}$ and $(\varepsilon_{fa'})_{f \in S_{a'}}$. These vectors are independent across applicants.

By the same argument as above, the focal applicant's offer indicators $(O_f)_{f \in S_a}$ are associated conditional on $\mathcal{G}$. Their marginal probabilities are again $1/(N_f+1)$. Note when $\beta^f = 0$, the offer indicators are independent conditional on $\mathcal{G}$. Thus

$$\Pr_{\beta^f}(a \text{ receives no offer} \mid \mathcal{G}) \geq \prod_{f \in S_a} \frac{N_f}{N_f+1} = \Pr_{0}(a \text{ receives no offer} \mid \mathcal{G}).$$

Averaging over $\mathcal{G}$ yields $\Pr_{\beta f}(a \text{ matches}) \leq \Pr_0(a \text{ matches})$, and hence $m_q^{\beta^f} \leq m_q^0$. □

D.3 Applicant-Side Alignment

Fix a focal applicant a and order the firms to which she has applied by $f_1, ..., f_q$, and again let $N_1, \ldots, N_q$ denote the corresponding competitor counts, each with the same marginal. Conditional on $(N_1, \ldots, N_q)$, firm preferences are independent across firms; also, symmetry implies that an applicant facing N_ℓ competitors receives the firm's offer with probability $1/(N_\ell+1)$. Hence her

probability of receiving no offer is

$$\Pr_{\beta^a}(a \text{ receives no offer}) = \mathbb{E}\left[\prod_{\ell=1}^{q} \frac{N_\ell}{N_\ell + 1}\right].$$

Let $r_{\beta^a} := \mathbb{E}[N_\ell/(N_\ell + 1)]$ denote the marginal rejection probability of a uniformly selected application. By the sufficient condition in the theorem,

$$\Pr_{\beta^a}(a \text{ receives no offer}) \geq r_{\beta^a}^q.$$

But in the baseline game, negative association (see, e.g., Appendix B) implies $\Pr_0(a \text{ receives no offer}) \leq r_0^q$. Combining the preceding inequalities gives

$$\Pr_{\beta^a}(a \text{ receives no offer}) \geq r_{\beta^a}^q \geq r_0^q \geq \Pr_0(a \text{ receives no offer}).$$

Thus $\Pr_{\beta^a}(a \text{ matches}) \leq \Pr_0(a \text{ matches})$ and by symmetry the claim follows. □

E Proof of Proposition 2

We will show that the match rate is maximized at $q = 1$ when $\lambda = A/F \geq c$. As in our other results that use Assumption 1, we elide some of the routine algebra. It will be helpful to define $\gamma := c/\lambda \in (0, 1]$ to be the ratio of capacity and market imbalance. We are going to compare the match rate $\widehat{m}_{q,c}$ to $s_{1,c}$, which is equal to the match rate when $q = 1$. Recall from the body that under Assumption 1, the probability that a focal applicant matches is $\pi_{q,c} = 1 - (1 - s_{q,c})^q$.

Our strategy is to upper bound the match rate under $q \geq 2$ and lower bound the match rate under $q = 1$.

For $q \geq 2$, note the expected total number of offers is at most Fc. There are Aq applications and by symmetry each receives an offer with probability $s_{q,c}$. Therefore $Aq \cdot s_{q,c} \leq Fc$, which implies

$$s_{q,c} \leq \frac{c}{\lambda q} = \frac{\gamma}{q},$$

$$\pi_{q,c} \leq 1 - \left(1 - \frac{\gamma}{q}\right)^q = g_\gamma(q). \tag{31}$$

The function $g_\gamma(q)$ in (31) will be used to write an upper bound of $\widehat{m}_{q,c}$ and is strictly decreasing for $q > \gamma$; in particular, it is strictly decreasing over the integers $q \geq 2$.

Now we lower bound the match rate under $q = 1$. Let $K \sim \text{Poisson}(\lambda)$ be the number of competitors at the focal applicant's firm under $q = 1$. Conditional on $K = k$, the focal applicant

receives an offer with probability $\min\{1, c/(k+1)\}$. Hence

$$\begin{aligned} s_{1,c} &= \Pr(K \le c-1) + \sum_{k=c}^{\infty} \frac{c}{k+1} \Pr(K = k) \\ &= \Pr(K \le c-1) + \frac{c}{\lambda} \Pr(K \ge c+1) \\ &= \gamma + (1-\gamma) \Pr(K \le c-1) - \gamma \Pr(K = c). \end{aligned} \tag{32}$$

The second equality uses $\Pr(K = k)/(k+1) = \Pr(K = k+1)/\lambda$; further, $\Pr(K = c-1) = (c/\lambda)\Pr(K = c)$. Since $\Pr(K \le c-1) \ge \Pr(K = c-1)$, using (32) in fact gives

$$s_{1,c} \ge \gamma - \gamma^2 \Pr(K = c). \tag{33}$$

Now for fixed c, the Poisson mass $e^{-\lambda}\lambda^c/c!$ is decreasing over $\lambda \ge c$. Therefore

$$\Pr(K = c) \le e^{-c}\frac{c^c}{c!} =: t_c,$$

and the sequence $\{t_c\}$ we just defined is strictly decreasing because $\frac{t_{c+1}}{t_c} = e^{-1}\left(1 + \frac{1}{c}\right)^c < 1$.

Now we break our comparisons down into cases, as we did in the one-to-one game.

First, compare the match rate from $q = 1$ to every quota $q \ge 3$ for all $c \ge 2$. When $c \ge 2$, we have $t_c \le t_2 = 2e^{-2}$, and use the bound $2e^{-2} < 8/27$. Combining this with equation (33), we get $s_{1,c} > \gamma - \frac{8}{27}\gamma^2$. But then $s_{1,c} > g_\gamma(3)$ because

$$g_\gamma(3) = 1 - \left(1 - \frac{\gamma}{3}\right)^3 = \gamma - \frac{\gamma^2}{3} + \frac{\gamma^3}{27}$$

and

$$\gamma - \frac{8}{27}\gamma^2 - g_\gamma(3) = \frac{\gamma^2(1-\gamma)}{27} \ge 0.$$

And since $\pi_{1,c} = s_{1,c}$ and $g_\gamma(q) \le g_\gamma(3)$ for every integer $q \ge 3$, by (31) we have

$$\pi_{1,c} > g_\gamma(3) \ge g_\gamma(q) \ge \pi_{q,c} \qquad \text{for every } q \ge 3.$$

Now make the comparison to $q = 2$ when $c \ge 3$. The mechanics are the same as above: if $c \ge 3$, then $t_c \le t_3 = 9/(2e^3) < 1/4$; by (33) we have $s_{1,c} > \gamma - \frac{\gamma^2}{4}$, and $g_\gamma(2) = \gamma - \gamma^2/4$. Thus

$$\pi_{1,c} = s_{1,c} > g_\gamma(2) \ge \pi_{2,c}.$$

Lastly we have $c = 2$ and $q = 2$, in which case $\lambda \ge 2$; here, the bounds we derived above are

too loose for our purposes. If $L \sim \text{Poisson}(x)$, then

$$\begin{aligned}\mathbb{E}\left[\min\left\{1, \frac{2}{L+1}\right\}\right] &= \Pr(L \leq 1) + \frac{2}{x}\Pr(L \geq 3) \\ &= \frac{2}{x} - \left(1 + \frac{2}{x}\right)e^{-x}. \end{aligned} \tag{34}$$

Applying (34) with $x = \lambda$ under $q = 1$ and $x = 2\lambda$ under $q = 2$ gives

$$\begin{aligned} s_{1,2} &= \frac{2}{\lambda} - \left(1 + \frac{2}{\lambda}\right)e^{-\lambda}, \\ s_{2,2} &= \frac{1}{\lambda} - \left(1 + \frac{1}{\lambda}\right)e^{-2\lambda}. \end{aligned}$$

Since $\pi_{2,2} = 2s_{2,2} - s_{2,2}^2$, substitution yields

$$\begin{aligned} s_{1,2} - \pi_{2,2} = \frac{1}{\lambda^2} - \left(1 + \frac{2}{\lambda}\right)e^{-\lambda} + 2\left(1 - \frac{1}{\lambda^2}\right)e^{-2\lambda} \\ + \left(1 + \frac{1}{\lambda}\right)^2 e^{-4\lambda}. \end{aligned}$$

We elide the verification that this is positive. Therefore $\pi_{1,2} = s_{1,2} > \pi_{2,2}$.

Combining the three preceding cases, we have collectively that $\pi_{1,c} > \pi_{q,c}$ for every $c \geq 2, q \geq 2$. Thus $q = 1$ maximizes the match rate. Applicant welfare is trivially maximized at $q = 1$, hence we have shown that $q = 1$ dominates all other quotas. □

F Multiple Rounds

F.1 Proposition 3

The number of unmatched applicants when applicants propose at round t is $A - \mathbb{E}[\text{number of matched applicants}]$. Equivalently, we can substitute and rewrite this as $\lambda F - F h_t^F = F(\lambda - h_t^F)$. Each unmatched applicant then proposes to one of the F firms; an unmatched firm receives at least one proposal with probability $1 - e^{-(\lambda - h_t^A)}$. So we can write the recursive form for the fraction of matched firms as

$$h_{t+1}^A = h_t^A + (1 - h_t^A)\left(1 - e^{-(\lambda - h_t^A)}\right).$$

Following the same construction, the fraction of matched firms when firms propose

$$h_{t+1}^F = h_t^F + (\lambda - h_t^F)\left(1 - e^{-(1 - h_t^F)/\lambda}\right).$$

Both initial conditions are equal at $t = 0$ and are increasing on $[0, 1]$, and for any current match rate, the applicant-proposing dynamics produce a weakly larger match rate than the firm-proposing

rule whenever $\lambda \geq 1$.[23] Therefore if $h_t^A \geq h_t^F$, $h_{t+1}^A \geq h_{t+1}^F$, and the result follows by induction.

Because the match rate is higher and because applicants apply in the order of their preferences, applicant welfare is also higher; dominance follows. □

F.2 Theorem 3

Recall we denote by h_2^F the match probability of a firm in a 2-round firm proposing game. We wish to show that $h_2^F \geq m_q$ for $q \geq 2$:

$$h_2^F = \lambda \left[1 - \exp\left(1 - e^{-1/\lambda} - \frac{2}{\lambda}\right)\right] \geq \lambda \left[1 - \left(1 - \frac{1 - e^{-\lambda q}}{\lambda q}\right)^q\right] = m_q.$$

We know it is sufficient to consider $q \in \{2, 3\}$. Then it suffices to define and prove the following inequality:

$$\Psi_q(\lambda) = \ln\left(\left(1 - \frac{1 - e^{-\lambda q}}{\lambda q}\right)^q\right) - \left(1 - e^{-1/\lambda} - \frac{2}{\lambda}\right) \geq 0$$

where we simply took logarithms. We then need the following for $q = \{2, 3\}$: to show $\Psi_q'(\lambda) \leq 0$ on $[1, \infty)$, and in the limit the function goes to zero: $\lim_{\lambda \to \infty} \Psi_q(\lambda) = 0$. Define $B(x) = x \cdot \frac{1-(x+1)e^{-x}}{x-1+e^{-x}}$ for $x > 0$, which appears in the derivative below:

$$\Psi_q'(\lambda) = \frac{1}{\lambda^2}\Big(B(\lambda q) - (2 - e^{-1/\lambda})\Big).$$

Observe $B(\lambda q) \leq 2 - e^{-1/\lambda}$ holds for $q = 2, 3$. Hence we have $\Psi_q'(\lambda) \leq 0$ for all $\lambda \geq 1$. Now we show that $\lim_{\lambda \to \infty} \Psi(\lambda) = 0$. As $\lambda \to \infty$, we have $e^{-\lambda q} \to 0$, hence

$$\left(1 - \frac{1 - e^{-\lambda q}}{\lambda q}\right)^q \to 1.$$

Then since $1/\lambda \to 0$ we also have $e^{-1/\lambda} \to 1$ and $2/\lambda \to 0$, so

$$1 - e^{-1/\lambda} - \frac{2}{\lambda} \to 0.$$

Therefore $\lim_{\lambda \to \infty} \Psi_q(\lambda) = \ln(1) - 0 = 0$. □

[23] Define $\Phi_A(x) = x + (1 - x)\left(1 - e^{-(\lambda - x)}\right) = 1 - (1 - x)e^{x - \lambda}$ and $\Phi_F(x) = x + (\lambda - x)\left(1 - e^{-(1-x)/\lambda}\right) = \lambda - (\lambda - x)e^{(x-1)/\lambda}$, note $\Phi_A'(x) = xe^{x-\lambda} \geq 0$ and $\Phi_F'(x) = \frac{x}{\lambda}e^{(x-1)/\lambda} \geq 0$, and then compute the sign of the difference of the original functions.