On the analytical bounds for average rank in one-to-one two-sided matching markets

Justin Hadad*

This Draft: November 30, 2023

* Department of Economics, University of Oxford, Oxford OX1 3UQ

Abstract

I show that the current average rank bounds in the one-to-one two-sided matching literature are loose in the limit, enough so that known comparative static results cannot be recovered. I construct a motivating problem to demonstrate this looseness, modelled after the result that there is some amount of increased competition that agents prefer to choosing their optimal mechanism. These results are tied to the literature via discussion about the size of the (asymptotic) core, as well as on the effects of competition.

1 Introduction

The analysis of rank distribution in stable matchings for one-to-one two-sided matching markets has been a topic of interest in economics and operations research since the seminal work of Gale and Shapley [1962]. There has been a pursuit since then to characterize the average rank obtained by agents, which indicates how satisfactory the match is on average for each side. This analysis is generally conducted in models with uniform preferences, and this is the setup with which we are concerned.

A notable result within the literature has been that the fraction of agents who have multiple partners across stable matchings tend to zero as $n \to \infty$. This leads to an important result: that the average rank of agents in the worst-case stable match and the best-case stable match is the same in the limit. Bounds which can be derived from the literature, specifically from Pittel [2019] and borrowing some results from Ashlagi et al. [2017], are broad enough that this result cannot be easily identified. Using their bounds for total rank, we obtain bounds for average rank which, while indeed well-defined, are loose in the limit. Thus to find more analytically precise asymptotic bounds is still an open problem.

I construct a motivating problem to demonstrate the insufficiency of the bounds, which asks the following: in a very unbalanced market, how many agents have to be added to the other side of the market to outweigh the benefits of securing the optimal match? The fact about the small core in the limit implies that, with very high probability, the average rank for the short side is going to be the same when proposing and accepting. If the average rank is indeed the same, then there should be a finite number of agents that a social planner could add to the long side of the market such that the average rank for the short side improves, regardless of which mechanism is chosen. But I derive analytically that the present bounds are loose enough that this result cannot be retrieved. I show via simulation that this motivating problem does indeed have a positive and finite solution for a given market size.

The paper proceeds as follows: In Section 2, I review the literature. Section 3 contains the derivation of the bounds for average rank. In Section 4 I simulate large unbalanced markets and show that the difference between men's average rank in men-proposing deferred acceptance and women-proposing deferred acceptance (thus in the range of stable matchings) as the market grows large goes to zero. In Section 5 I show any analytic solution to the motivating problem must be negative, and simulate to show that the analytic solution should indeed be positive and finite. Section 6 concludes, and the appendix contains the proof for the key proposition in Section 5.

2 Literature Review

To the best of my knowledge, Pittel [2019] contains the most recent and robust analysis of average rank in large unbalanced two-sided matching markets. He provided the probability with which total rank falls within a given interval; in the following section, I acquire bounds for average rank, additionally using results from the asymptotic analysis of Ashlagi et al. [2017]. Of note is that the results of Pittel [2019] are agreeable with recent findings in Nikzad [2022], which deployed graph theory to show that the average rank for men in the men-optimal stable matching is bounded from above by a constant.

There is some relation to the literature on the size of the core in the limit. A notable finding in matching theory has been that the size of the core—analogous to the number of unique stable matchings—shrinks as the markets grow more unbalanced. This implies the fraction of agents who have multiple unique partners within a stable matching goes to zero, and thus there is zero benefit of proposing into the limit. The first result about the core was found by Pittel [1989] and Knuth et al. [1990], who showed that the fraction of agents in balanced markets with multiple stable partners approaches one as n grows large. Immorlica and Mahdian [2005] showed that in sufficiently large markets there is essentially one unique stable matching in the limit, and that truth-telling dominates for both sides. Kojima and Pathak [2009] found an analogue for the same result, but in many-to-one, sufficiently thick markets. Later, Ashlagi et al. [2017] showed that even adding one agent to just one side of a balanced market can make the core collapse.

There has also been research on the different model conditions that lead to the same small-core result. For example, Roth and Peranson [1999] showed that a limited acceptability assumption (where some agents deem some number of agents unacceptable, and which none of Pittel [2019], Ashlagi et al. [2017], or I impose), also generates the result, that the fraction of people with multiple stable partners tends to zero as the market grows large. Liu and Pycia [2016] proved that regular mechanisms that are asymptotically efficient, symmetric, and strategyproof are asymptotically equivalent, so the choice of mechanism is actually not as relevant to improving outcomes. Lee [2016] showed that agents cannot profitably manipulate the mechanism even when preferences have a common-value component. Knuth [1996], one of the first to consider the notion of average rank in two-sided matching markets, deployed uniform preferences. Holzman and Samet [2014], for example, demonstrated how common values lead to assortative matching, thus the average rank is intuitive from the size of the market (and also has a small core).

I generate the motivating problem as an extension of results from Crawford [1991] (an analogue to which was available in Kelso and Crawford [1982]), and in the auctions setting of Bulow and Klemperer [1996]. Crawford [1991] proved that adding women to the other side of the market weakly improves the men-optimal stable matching for all men. Later, Bulow and Klemperer [1996] showed that, in a setting of private values, expected revenue from an absolute English auction with $n+1$ bidders exceeds expected revenue from an English auction with n bidders followed by a take-it-or-leave-it offer to the last remaining bidder.

3 Theoretical Results

Let $\mathcal{M}_{n_1,n_2}$ be the set of stable matchings for a market with size n_1 men and n_2 women where $n_1 < n_2$. For any matching $M_{n_1,n_2} \in \mathcal{M}_{n_1,n_2}$, let $Q(M_{n_1,n_2})$ be the total rank of men for their partners. When the size of the market is clear, let $\mathcal{M}$ and M be shorthand respectively for $\mathcal{M}_{n_1,n_2}$ and M_{n_1,n_2}. Define $\mathbf{n}$ as the vector (n_1, n_2) and $s(\mathbf{n}) := \log \frac{n_2}{n_2-n_1}$, and consider that $s(\mathbf{n})$ is bounded.[1] Note that lower rank is preferred.

The notion of a *concentration point* is useful for determining the interval in which a value

[1]We note this because the interested reader may visit Pittel [2019] and note there are two cases: one where $s(\mathbf{n})$ is bounded, and another where $s(\mathbf{n})$ is unbounded. The unbounded case does not apply in very large markets; note that as $n_2 \to \infty$, the denominator sends the logarithm to zero.

is expected to lie.

Definition 1. *The concentration point* q_α *for total rank is defined as the smallest* q *such that* $\mathrm{P}(|Q(M) - q_\alpha| \leq \delta) \geq \alpha$ *for all matchings* $M \in \mathcal{M}$ *for some positive* δ.

From Pittel [2019] we get bounds on q_α for men's total rank $Q(M)$.

Theorem 1 (Pittel [2019])**.** *For* $n_2 > n_1$ *and* n_1 *sufficiently large,*

$$\mathrm{P}\left(\max_{\mathcal{M}}\left|\frac{Q(\mathcal{M})}{n_2 s(\mathbf{n})} - 1\right| \geq \delta(\mathbf{n})\right) \leq \mathcal{P}(\mathbf{n}) \tag{1}$$

where $s(\mathbf{n}) = \log \frac{n_2}{n_2 - n_1}, \delta(\mathbf{n}) := n_1^{-a}, \mathcal{P}(\mathbf{n}) := \exp\left(-n_1^{1-2a}\right)$ *for some* $a < \frac{1}{2}$.

I rearrange this statement to describe the concentration point around average rank.

Lemma 1. *The concentration point around the average rank is described by*

$$\mathrm{P}\left(\max_{\mathcal{M}}\left|\frac{Q\left(\mathcal{M}\right)}{n_1} - \frac{n_2}{n_1} s\left(\mathbf{n}\right)\right| \geq \frac{n_2}{n_1}\delta\left(\mathbf{n}\right) s\left(\mathbf{n}\right)\right) \leq \mathcal{P}\left(\mathbf{n}\right) \tag{2}$$

for $n_2 > n_1$ *and* n_1 *sufficiently large, and where* $s(\mathbf{n}) = \log \frac{n_2}{n_2 - n_1}$, $\delta(\mathbf{n}) = n_1^{-a}, \mathcal{P}(\mathbf{n}) = \exp\left(-n_1^{1-2a}\right)$ *for some* $a < \frac{1}{2}$.

The upper bound of the implied distribution corresponds to the worst-case ranking for men (as in women-proposing deferred acceptance, henceforth "WPDA") and the lower bound corresponds to the best-case ranking (as in men-proposing deferred acceptance, henceforth "MPDA"). From this concentration point, I construct bounds for average rank, which requires leveraging additional results from Ashlagi et al. [2017] for the upper-bound average rank.

Proposition 1. *The expectation of the lower bound for average rank in a market with* n_1 *agents on the shorter side and* n_2 *agents on the longer side is*

$$\mathbb{E}\frac{Q(\mathcal{M}_{n_1,n_2})}{n_1} \geq (1 - \mathcal{P}(\mathbf{n}))\left(\frac{n_2}{n_1}s(\mathbf{n}) - \frac{n_2}{n_1}\delta(\mathbf{n})s(\mathbf{n})\right) + \mathcal{P}(\mathbf{n})1$$

where $s(\mathbf{n}) = \log \frac{n_2}{n_2 - n_1}$, $\delta(\mathbf{n}) = n_1^{-a}, \mathcal{P}(\mathbf{n}) = \exp\left(-n_1^{1-2a}\right)$*, and for some* $a < \frac{1}{2}$.

Proof. Begin with the concentration point bound from Pittel [2019]:

$$\mathrm{P}\left(\max_{\mathcal{M}}\left|\frac{Q(\mathcal{M})}{n_2 s(\mathbf{n})} - 1\right| \geq \delta(\mathbf{n})\right) \leq \mathcal{P}(\mathbf{n}).$$

With probability $\mathcal{P}(\mathbf{n})$ it holds that $\max_{\mathcal{M}} \left| \frac{Q(\mathcal{M})}{n_2 s(\mathbf{n})} - 1 \right| \geq \delta(\mathbf{n})$. Multiplying through by $s(\mathbf{n})$ within $\mathrm{P}(\cdot)$, I arrive at the normalized total men rank:

$$\mathrm{P}\left(\max_{\mathcal{M}} \left| \frac{Q(\mathcal{M})}{n_2} - s(\mathbf{n}) \right| \geq \delta(\mathbf{n}) s(\mathbf{n}) \right) \leq \mathcal{P}(\mathbf{n}).$$

So for at least $1 - \mathcal{P}(\mathbf{n})$ ratio of preference profiles, the normalized total men rank $\frac{Q(\mathcal{M})}{n_2}$ is within $\delta(\mathbf{n}) s(\mathbf{n})$ of $s(\mathbf{n})$. Further manipulation (dividing through by n_1 and removing normalization) generates the inequality

$$\mathrm{P}\left(\max_{\mathcal{M}} \left| \frac{Q(\mathcal{M})}{n_1} - \frac{n_2}{n_1} s(\mathbf{n}) \right| \geq \frac{n_2}{n_1} \delta(\mathbf{n}) s(\mathbf{n}) \right) \leq \mathcal{P}(\mathbf{n}).$$

This means that for at least $1 - \mathcal{P}(\mathbf{n})$ ratio of preference profiles, the unscaled average men rank $\frac{Q(\mathcal{M})}{n_1}$ is within $\frac{n_2}{n_1} \delta(\mathbf{n}) s(\mathbf{n})$ of $\frac{n_2}{n_1} s(\mathbf{n}) = \frac{n_2}{n_1} \log \frac{n_2}{n_2 - n_1}$. So, for $1 - \mathcal{P}(\mathbf{n})$ of preference profiles, $\frac{Q(\mathcal{M})}{n_1} \leq \frac{n_2}{n_1} \delta(\mathbf{n}) s(\mathbf{n}) + \frac{n_2}{n_1} s(\mathbf{n})$. For $\mathcal{P}(\mathbf{n})$ of preference profiles we can merely assert that $\frac{Q(\mathcal{M})}{n_1} \geq 1$, hence our bound from below is

$$\mathbb{E} \frac{Q(\mathcal{M})}{n_1} \geq (1 - \mathcal{P}(\mathbf{n})) \left(\frac{n_2}{n_1} s(\mathbf{n}) - \frac{n_2}{n_1} \delta(\mathbf{n}) s(\mathbf{n}) \right) + \mathcal{P}(\mathbf{n}) 1.$$

□

For the upper bound of the average rank I leverage the following result from Ashlagi et al. [2017].

Theorem 2 (Ashlagi et al. [2017]). *Let $K = \lambda n_1$ and $\lambda > 0$ be any positive constant. Consider a sequence of random matching markets with n_1 men and $(1+\lambda) n_1$ women. Define the constant $\kappa = 1.01(1+\lambda)\log(1+1/\lambda)$. With high probability, in every stable matching, the average rank of wives is at most κ, the average rank of husbands is at least $n_1/(1+\kappa)$, and the fractions of men and women who have multiple stable partners converge to 0 as $n_1 \to \infty$.*

Proposition 2. *Let $n_2 = \lambda n_1$ and $\lambda > 0$ be any positive constant. The expectation of the upper bound for average rank in a market with n_1 agents on the shorter side and n_2 agents on the longer side is*

$$\mathbb{E} \frac{Q(\mathcal{M}_{n_1,K})}{n_1} \leq (1 - \mathcal{P}(\mathbf{n})) \left(\frac{n_2}{n_1} s(\mathbf{n}) + \frac{n_2}{n_1} \delta(\mathbf{n}) s(\mathbf{n}) \right) + \mathcal{P}(\mathbf{n}) \left(1.01(1+\lambda)\log(1+1/\lambda)\right)$$

where $s(\mathbf{n}) = \log \frac{n_2}{n_2 - n_1}$, $\delta(\mathbf{n}) = n_1^{-a}, \mathcal{P}(\mathbf{n}) = \exp\left(-n_1^{1-2a}\right)$, and for some $a < \frac{1}{2}$.

The derivation is equivalent to the above except the worst-case average rank is bounded from above by the constant $\kappa = 1.01(1+\lambda)\log(1+1/\lambda)$.

In Figure 1, I plot the average rank, as well as the upper and lower bounds, for a given a value. Note the high-probability event and low-probability event constitute a bounded area. Recognize in most cases that, because of the small probability of $\mathcal{P}\left(\mathbf{n}\right)$, average rank coincides with the higher probability event.

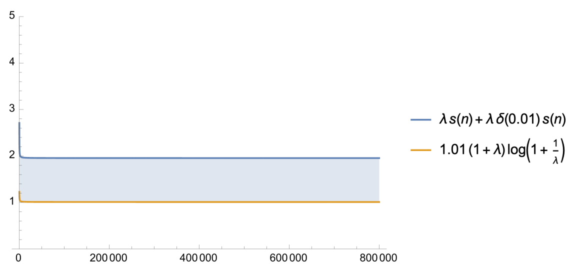

(a) Upper Bound for $a = 0.01$

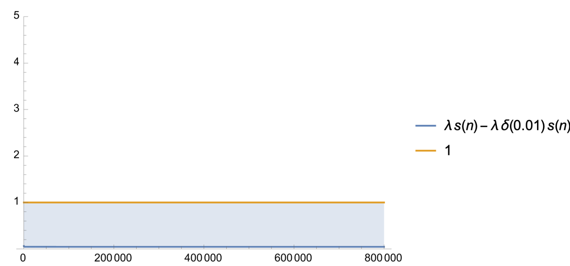

(b) Lower Bound for $a = 0.01$

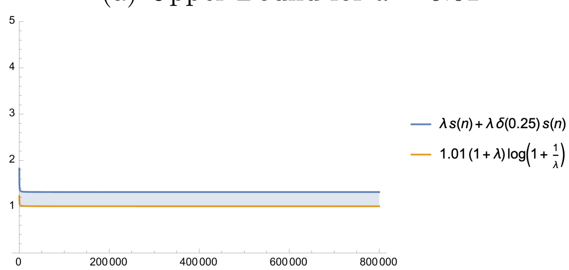

(c) Upper Bound for $a = 0.25$

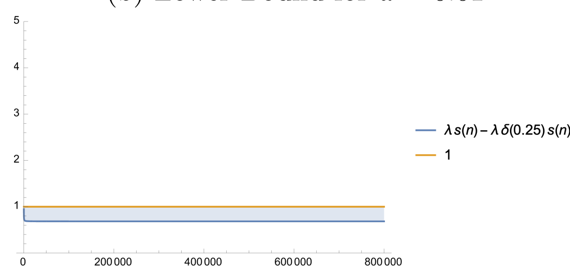

(d) Lower Bound for $a = 0.25$

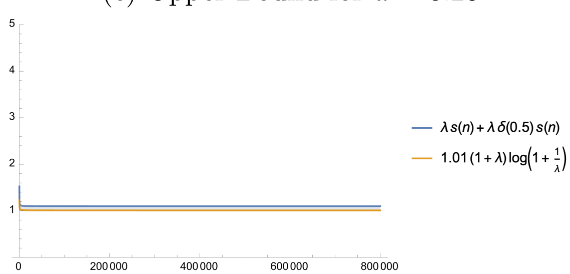

(e) Upper Bound for $a = 0.50$

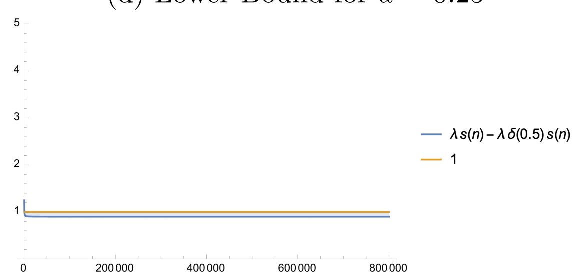

(f) Lower Bound for $a = 0.50$

Figure 1: Upper and Lower Bounds of Theoretical Average Rank

If the bounds were sufficiently tight, the gap between the inequalities in the above propositions would be quite small, so long as $n_2 > n_1$. But that is not upheld. Consider alone the "high" probability event, that is the term in each bound with the coefficient $(1 - \mathcal{P}(\mathbf{n}))$. The gap between the upper bound and lower bound instance is

$$\left(\frac{n_2}{n_1} s\left(\mathbf{n}\right) + \frac{n_2}{n_1}\delta\left(\mathbf{n}\right)\right) - \left(\frac{n_2}{n_1} s(\mathbf{n}) - \frac{n_2}{n_1}\delta(\mathbf{n}) s(\mathbf{n})\right) = 2\frac{n_2}{n_1^{1+a}}$$

which is obviously unbounded as $n_2 \to \infty$.

4 Simulations

In the investigation of large unbalanced matching markets, I employ a computational approach to examine the difference in average rank outcomes of participants.

I run simulations for a given market instance where the number of men is strictly less than the number of women. Each participant possesses a preference ordering independently and identically drawn over the members of the opposite side, with no ties in preferences. Specifically, a permutation of the set $\{0, 1, 2, ..., n-1\}$ is generated for each participant, where n is the count of agents on the opposite side. The deferred acceptance algorithm that generates a stable matching is shown in Algorithm 1.

Algorithm 1 Deferred Acceptance Algorithm

Require: Proposing preferences matrix P, Accepting preferences matrix A
Ensure: Matching M

Initialize all proposers and acceptees as unmatched
Create empty proposal lists for all acceptees
while there exists an unmatched proposer p with a non-empty preference list **do**
  Let a be the highest-ranked acceptee in p's preference list
  Remove a from p's preference list
  if a is unmatched **then**
    Match p with a
  else if a prefers p to her current match p' **then**
    Unmatch p' from a
    Match p with a
  end if
end while
return Matching M

The algorithm converges to a stable matching where no pair of participants could mutually benefit from deviating from their current match (there is no "blocking pair"). The focal outcome metric is the average rank of the partners to whom participants are matched. The script generates a market instance and preferences, runs MPDA and WPDA, and computes average rank for each. Averages are calculated across 50 unique market instances. I plot the average rank for men in MPDA and WPDA in Figure 2.

Figure 3 contains the difference in average rank across best- and worst-case stable matchings for a given market arrangement. From the figure it is clear that MPDA weakly dominates WPDA. Note when comparing difference in rank that there is a fair number of simulations for which the average rank difference is close to zero. This implies that the utility difference between the best-case and worst-case stable matching is marginal (if not zero) and thus that there is a unique stable matching for most market arrangements. The gap approaches zero

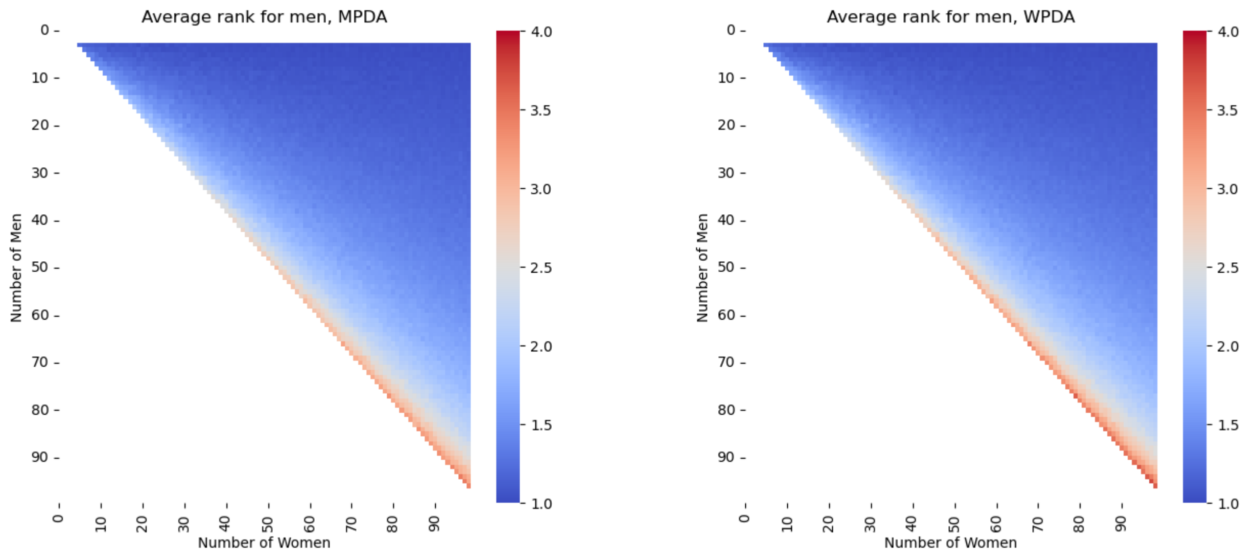

Figure 2: Simulated average rank for men in MPDA and WPDA

as market grows more unbalanced.

5 Motivating Problem

I now discuss the motivating problem, which takes inspiration from Crawford [1991] and Bulow and Klemperer [1996]. I seek the following: what number of agents have to be added to the long side of the market to make the average rank from that new market *certainly lower* than the average rank from proposing? I first discuss how the bounds constructed from Pittel [2019] and Ashlagi et al. [2017] imply a negative solution, and include a proof in the appendix. I then use simulations to show that a positive, finite solution does exist, and that the derived bounds are indeed too loose.

Numerically, I seek the value k for which the upper bound of $\frac{Q(M_{n_1,K})}{n_1}$ is less than the lower bound of $\frac{Q(M_{n_1,n_2})}{n_1}$. Because, in practice, we know the fraction of agents who have multiple stable partners goes to 0 as $n \to \infty$, we expect the gap between these bounds to go to zero.

Proposition 3. *Using the bounds from Pittel [2019] and Ashlagi et al. [2017], there is no positive $k = K - n_2$ for which the following inequality holds:*

$$\mathbb{E}\frac{Q(\mathcal{M}_{n_1,K})}{n_1} \leq \mathbb{E}\frac{Q(\mathcal{M}_{n_1,n_2})}{n_1}.$$

A positive solution to this would give the number of agents who have to be added to the larger side of the market for the smaller side to benefit equally from simply proposing in the

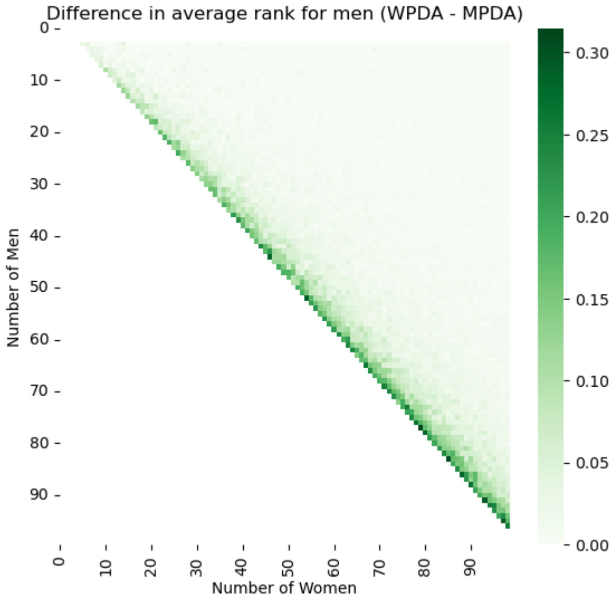

Figure 3: Difference between simulated average rank for men in MPDA and WPDA

original market. Note that this is solving the following:

$$(1-\mathcal{P}(\mathbf{n}_K))\left(\frac{K}{n_1}s\left(\mathbf{n}_K\right)+\frac{K}{n_1}\delta\left(\mathbf{n}_K\right)s\left(\mathbf{n}_K\right)\right)+\mathcal{P}\left(\mathbf{n}_K\right)(1.01(1+\lambda)\log(1+1/\lambda)) \tag{3}$$

$$\leq(1-\mathcal{P}(\mathbf{n}))\left(\frac{n_2}{n_1}s(\mathbf{n})-\frac{n_2}{n_1}\delta(\mathbf{n})s(\mathbf{n})\right)+\mathcal{P}(\mathbf{n})1 \tag{4}$$

where $s(\mathbf{n})=\log\frac{n_2}{n_2-n_1}$, $\delta(\mathbf{n})=n_1^{-a}, \mathcal{P}(\mathbf{n})=\exp\left(-n_1^{1-2a}\right)$, and for some $a<\frac{1}{2}$. The proof is in the appendix.

The fact that no positive number exists is contrary to expectation. If the bounds were sufficiently tight, the shaded area between the bounds as is seen in Figure 1 could be viewed analogously to the size of the core, which should vanish in the limit. Indeed, these bounds would suppose a range of stable matchings. But the looseness of the bounds means this result cannot be recovered. I now show in simulation that there is a finite number of agents that can be added to other side of the market to outweigh the benefits of proposing.

Contradiction from simulation

In simulation the motivating problem does have a finite, positive solution for a given market size. This follows intuition: given there is, with very high probability, a unique stable matching in large unbalanced matching markets, securing proposal rites should make essentially zero difference on the average rank, and so, given average rank is normalized and preferences are regenerated for a given market size, the average rank rank should weakly improve as the number of women increases.

I show the number of agents that have to be added to outweigh proposing effects in Figure 4, for market sizes up to 100. Figure 5 shows the same but normalized for the original length of the long side, thus each cell corresponds to a fraction of the long side of the market that flips the inequality.

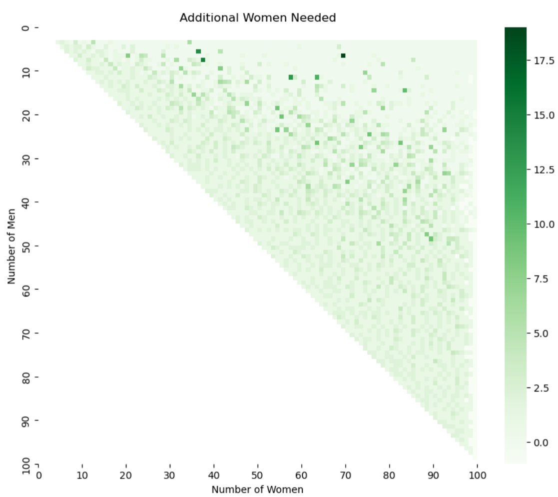

Figure 4: The additional number of women needed to outweigh the benefits of proposing.

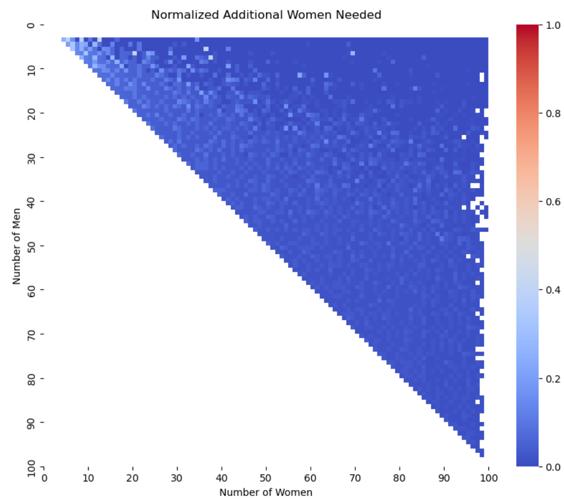

Figure 5: The additional number of women needed to outweigh the benefits of proposing, normalized to the number of women.

6 Conclusion

This paper revisits the problem of rank distribution in stable matchings for one-to-one two-sided matching markets, a topic of enduring interest since the foundational work of Gale and Shapley. I focus on the implications of recent analytical bounds for large markets, particularly those presented by Pittel [2019], and show that these bounds are loose in the limit.

To illustrate this, I introduce a motivating problem that seeks to quantify the number of agents that must be added to the longer side of the market to make the benefit of increased competition outweigh the benefit of securing the optimal match. My analytical results show that no such number exists under the current bounds, thereby highlighting the limitations of existing analytical tools for studying this problem.

Acknowledgements

I owe thanks to Marek Pycia for comments and discussion, as well as Alex Teytelboym and Kyle Woodward for additional review.

Appendix

Proof of Proposition 3. I follow Ashlagi et al. [2017] and recognize that with $K = \lambda n_1$, an unbounded long side implies $\lambda \to \infty$. Taking the limit using L'Hopital's rule we see that as $\lambda \to \infty$,

$$e^{-\theta(n_1^{1-2a})}(1.01(1+\lambda)\log(1+1/\lambda)) \to 0.$$

Now I focus on the first term on each side of the inequality, thus I intend to simplify

$$(1 - e^{-\theta(n_1^{1-2a})})\left(\frac{K}{n_1}\frac{n_1}{K-n_1} + \frac{K}{n_1}\delta\left(\mathbf{n}_K\right)\frac{n_1}{K-n_1}\right) \leq \tag{5}$$

$$(1 - e^{-\theta(n_1^{1-2a})})\left(\frac{n_2}{n_1}\frac{n_1}{n_2-n_1} - \frac{n_2}{n_1}\delta(\mathbf{n})\frac{n_1}{n_2-n_1}\right). \tag{6}$$

Eliminating terms and recognizing that n_1 cancels everywhere:

$$\left(\frac{K}{K-n_1} + \delta\left(\mathbf{n}_K\right)\frac{K}{K-n_1}\right) \leq \left(\frac{n_2}{n_2-n_1} - \delta(\mathbf{n})\frac{n_2}{n_2-n_1}\right).$$

Multiplying both sides by $(K - n_1)$ and plugging in for $\delta(\cdot)$ gives

$$K + n_1^{-a}K \leq \frac{n_2(K-n_1)}{n_2-n_1} - n_1^{-a}\frac{n_2(K-n_1)}{n_2-n_1}.$$

Factoring out and isolating K gives

$$K \leq \frac{n_2}{n_2-n_1}\frac{(1-n_1^{-a})}{(1+n_1^{-a})}(K-n_1).$$

It is useful to define $Z := \frac{n_2}{n_2-n_1}\frac{(1-n_1^{-a})}{(1+n_1^{-a})}$, giving

$$K(1-Z) \leq -n_1 Z,$$

so

$$K \leq \frac{-n_1 Z}{1-Z}.$$

Recall we are seeking $k = K - n_2$, which defines defines the number of agents that have to be added to the long side of the market for the benefit of increased competition on the short side of the market to outweigh the benefits of proposing. It is already immediate that $k = K - n_2$ is negative; nonetheless plugging back in for Z in the above gives

$$K \leq \frac{-n_1 \frac{n_2}{n_2 - n_1} \frac{(1-n_1^{-a})}{(1+n_1^{-a})}}{1 - \frac{n_2}{n_2 - n_1} \frac{(1-n_1^{-a})}{(1+n_1^{-a})}},$$

so

$$k \leq \frac{-n_1 \frac{n_2}{n_2 - n_1} \frac{(1-n_1^{-a})}{(1+n_1^{-a})}}{1 - \frac{n_2}{n_2 - n_1} \frac{(1-n_1^{-a})}{(1+n_1^{-a})}} - n_2.$$

But $k < 0$ for all values n_1 and n_2 such that $n_1 < n_2$.

$\square$

References

I. Ashlagi, Y. Kanoria, and J. Leshno. Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, 2017.

J. Bulow and P. Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180–194, 1996.

V. P. Crawford. Comparative statics in matching markets. *Journal of Economic Theory*, 54 (2):389–400, 1991.

D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 1962.

R. Holzman and D. Samet. Matching of like rank and the size of the core in the marriage problem. *Games and Economic Behavior*, 88:277–285, 2014.

N. Immorlica and M. Mahdian. Marriage, honesty, and stability. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 53–62, 2005.

A. S. Kelso and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

D. E. Knuth. An exact analysis of stable allocation. *Journal of Algorithms*, 20(2):431–442, 1996.

D. E. Knuth, R. Motwani, and B. Pittel. Stable husbands. *Random Structures & Algorithms*, 1(1):1–14, 1990.

F. Kojima and P. A. Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99(3):608–27, June 2009.

S. Lee. Incentive Compatibility of Large Centralized Matching Markets. *The Review of Economic Studies*, 84(1):444–463, 09 2016.

Q. Liu and M. Pycia. Ordinal efficiency, fairness, and incentives in large markets. 09 2016.

A. Nikzad. Rank-optimal assignments in uniform markets. *Theoretical Economics*, 17:25–55, 2022.

B. Pittel. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, 1989.

B. Pittel. On likely solutions of the stable matching problem with unequal numbers of men and women. *Mathematics of Operations Research*, 2019.

A. E. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4):748–780, September 1999.