Unbalanced Random Matching Markets

Itai Ashlagi   Yashodhan Kanoria   Jacob D. Leshno*

October 10, 2013

Abstract

We characterize the core in random matching markets with unequal numbers of men and women. Preference lists are drawn uniformly at random, independently across agents. We find that, with high probability, a vanishing fraction of agents have multiple stable partners, i.e., the core is small. Further, we find that being on the short side of the market confers a large advantage. For each agent, we assigned a rank of 1 to the agent's most preferred partner, rank of 2 to the next most preferred partner and so forth. If there are n men and $n+1$ women then, we show that, with high probability, in each stable matching, men's average rank of wives is no more than $3 \log n$, whereas the women's average rank of husbands is at least $n/(3 \log n)$. If there are n men and $(1 + \lambda)n$ women for $\lambda > 0$, with high probability, in any stable matching, men's average rank of wives is as small as $O(1)$, whereas the women's average rank of husbands is as large as $\Omega(n)$. Simulations show that these features are observed even in small markets. These findings for unbalanced markets contrast sharply with known results for balanced markets: in particular, our results show that the large core in the balanced case is a knife-edge phenomenon that breaks with the slightest imbalance.

1 Introduction

A two-sided marriage market comprises men and women each of whom has preferences over potential spouses. Equilibrium outcomes are stable matchings, which are matchings where no man and woman prefer each other over their own partners. Stable matchings are exactly the core allocations in this market. This matching model (and variations of it) has been applied in successful market designs, such as school choice programs in Boston (Abdulkadiroglu et al. (2006)) and NYC (Abdulkadiroğlu et al. (2005)), and the National Resident Matching Program (Roth and Peranson (1999)). These markets use a clearinghouse that outputs a stable matching with respect to agents reported preferences.

The set of stable matchings is not necessarily small. Pittel (1989b) characterized balanced random matching markets, in which there are n men and n women, each with independently drawn

uniformly random, complete preferences over the other side of the market, and found that the core is likely to be large. Assigning a rank of 1 when an agent is matched to his first preference, 2 for his second preference, etc., Pittel (1989b) showed that with high probability the men's average rank of wives can be as good as $\log n$, or as bad as $n/(\log n)$, depending on the stable matching. The fraction of agents who are matched to different partners in different stable matchings tends to 1 as the market grows large (Roth and Peranson, 1999; Pittel, 1989b).

These findings are in sharp contrast to empirical evidence. For instance, Roth and Peranson (1999) found that the set of stable matchings in the NRMP is very small, with all stable matchings resulting in almost the same assignment. This observation rendered inconsequential a political debate regarding witch stable mechanism the NRMP should implement. We are not aware of any matching dataset that has shown a sizeable core.

In this paper we characterize the core of random matching markets, as studied by Pittel (1989b), but with an *unequal* number of men and women. We find that even the slightest imbalance in the market leads to harsh competition on the long side and has a striking effect on the size of the core, and on the average rankings in all stable matchings. In a matching market with uniformly random and independent preferences, but with n men and $n + 1$ women, with high probability, under any stable matching, the mens average rank of wives is at least as good as $3 \log n$ while the womens average rank of husbands is at least as bad as $n/(3 \log n)$. Even when there is just one extra woman the core is small; all stable matchings give approximately the same average rank for men and women, and most agents have the same partner in all stable matchings. We prove our results asymptotically in n, and verify through simulations that these features occur in small markets as well (see Figure 2).

For comparison, the men's average rank of wives in every stable matching is approximately the same as it would have been if we were to ignore womens preferences and just let men choose sequentially in a random order[1]. On the other hand, women only do a factor $O(\log n)$ better than they would have had they been matched uniformly at random. Loosely speaking, in any stable matching, the men get to choose and the women are chosen.

Our results resemble the shrinking core phenomenon in a competitive buyer-seller market. In a market with n identical buyers who have unit demand for a value of 1 and n sellers who have unit supply with reservation value of 0, every price between 0 and 1 gives a core allocation. But with $n + 1$ sellers there is a unique clearing price of 0, since any buyer has an "outside option of buying from the unmatched seller who will sell for any positive price. In our setting, however, while an additional unmatched woman is willing to match with every man, many men prefer their partners over the unmatched woman. To understand why such men still benefit from having an additional woman, we need to understand how the allocations of different men are connected. Loosely speaking, some men rank highly the unmatched woman, and must therefore be matched to a highly ranked woman, even when some of these women are not matched to highly ranked partners. These unsatisfied women provide a new "outside option" for other men, who in turn must be matched to a highly ranked women. This renders a sequence of "outside options" making

---

[1]This will be the Random Serial Dictatorship (RSD) mechanism for the men (see e.g. Abdulkadiroğlu and Sönmez (1998)).

all men better off.

To prove the results we introduce a new matching algorithm, which builds on the algorithm used by McVitie and Wilson (1971) and Immorlica and Mahdian (2005). The algorithm exploits the lattice structure of the set of stable matchings: there is a women-optimal stable matching (WOSM) which is the least preferred by all men and a men-optimal stable matching (MOSM) which is the least preferred by all women, and by calculating them both we bound the set of stable matchings. The algorithm first calculates the MOSM by running the men-proposing Deferred Acceptance (Gale and Shapley (1962)). The algorithm then progresses to calculate the WOSM by sequentially checking for each woman to see whether she has a better stable husband. To find whether a woman has a better stable partner, she triggers a rejection chain by divorcing her current husband who proposes to his most preferred woman who has yet to reject him till some woman accepts him. The latest acceptance by a women may displace another man, who in proposes in a similar fashion, and so forth. The rejection chain ends in one of two ways: first, the woman receives a proposal from a man she prefers over the man she divorced, resulting a in new stable matching. Second, when the chain ends in a proposal to an unmatched woman, implying that there is no better stable partner for the woman under the circumstance. Here, the chain uncovers the ripple effect of the unmatched woman, who is an attractive outside option for the last man in the chain. Thus the original man cannot be rejected and the original husband is the woman's best stable partner. The proof analyzes the run of the algorithm on a random matching market and shows that we are likely to find that most chains end up with an application to the unmatched women, and that very few agents participate in chains that make the original woman better off. We conclude that the MOSM and WOSM are likely to be close.

Our results have an important implication for strategic behavior in matching markets. In a marriage market an agent can manipulate only if he has multiple stable partners. For random markets, we show that unless the market is balanced, a diminishing number of agents have multiple stable partners, implying that there is a limited scope for manipulation.

The behavior of random matching markets with a balanced number of men and women has been extensively studied by Pittel (1989a) and Knuth et al. (1990), who show that the core is likely to be large. Coles and Shorrer (2012) and Lee (2011) study manipulation in asymptotically large balanced matching markets, making different assumptions about the utility functions of agents. Coles and Shorrer (2012) define agents' utilities to be equal to the rank of their spouse (varying between 0 and n), which grow linearly with the number of agents in the market. They show that women are likely to gain substantially from manipulating the men-proposing DA by truncating their preferences. Lee (2011) allows for correlation in preferences but assumes that utilities are kept bounded as the market grows large. He shows that in a large enough market most agents cannot gain much utility from manipulating the men-proposing DA. The difference between the two conclusions can be reconciled as a result of the different utility parametrization. In the balanced market, agents are likely to have stable partners of rank ranging from $\log n$ to $n/\log n$, and this difference in rank can be small or large in terms of utilities. In contrast, in the unbalanced market agents are likely to have a unique stable partner.

Previous papers showed assumptions on the structure of the market which imply that the set

of stable matchings is small. In the many-to-one setting Azevedo and Leshno (2012) look at large markets with a constant number of schools and an increasing number of students. They show that the limit can be described by a model consisting of a continuum of students, which has a unique stable matching under broad conditions. Holzman and Samet (2013) shows that the core is small in markets with highly correlated preferences. Immorlica and Mahdian (2005) and Kojima and Pathak (2009) study matching markets in which one side has short, randomly drawn preference lists, and show that in large markets the set of stable matchings is small[2]. Kojima et al. (2010) and Ashlagi et al. (2010) study matching between residents and hospitals when there are couples who need two positions. They essentially assume there are linearly more hospital positions and find that for large markets the core is nonempty and small if the number of couples grows at a slower rate than the number of singles. However, in all of these models a large fraction of popular agents remains unmatched. In our model, we find the set of stable matchings to be small even in small markets, when just one agent remains unmatched, and preferences are heterogenous and uncorrelated. This suggests that in practice matching markets will generally have a small core.

The rest of the paper is organized as follows. Section 2 presents our model and results. In Section 3 we present a matching algorithm which is the basis of our analysis. In Section 4 we prove our main theorem. Section 5 presents numerical simulations which show that our results hold for small markets and that they are robust to the addition of a small amount of correlation in preferences. Finally, in Section 6 we give some final remarks and discuss the limitations of our model.

2 Model and results

2.1 Random matching markets

In a two-sided matching market there is a set of men $\mathcal{M} = \{1, \ldots, n\}$ and set of women $\mathcal{W} = \{1, \ldots, k\}$. Each man m has a complete strict preference list $\succ_m$ over women, and each woman w has a complete strict preference list $\succ_w$ over the set of men. A *matching* is a mapping μ from $\mathcal{M} \cup \mathcal{W}$ to itself such that for every $m \in \mathcal{M}$, $\mu(m) \in \mathcal{W} \cup \{m\}$, and for every $w \in \mathcal{W}$, $\mu(w) \in \mathcal{M} \cup \{w\}$, and for every $m, w \in \mathcal{M} \cup \mathcal{W}$, $\mu(m) = w$ implies $\mu(w) = m$.

A matching μ is *unstable* if there is a man m and a woman w such that $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$. A matching is *stable* if it is not unstable. We say that m is stable for w (and vice versa) if there is a stable matching such that m is matched to w. It is well-known that the core of matching market is the set of stable matchings.

Our focus in this paper is to characterize the stable matchings of randomly generated matching markets. A *random matching market* is generated by drawing a complete preference list for each man and each woman independently, uniformly at random.[3] In section 5 we provide simulation results for more general distributions over preferences, which allow for correlations.

A stable matching always exists, and can be found using the Deferred Acceptance (DA) algorithm by Gale and Shapley (1962). In the DA algorithm one side, either men or women, propose.

---

[2]This was conjectured by Roth and Peranson (1999).

[3]That is, we draw for each man m a complete ranking $\succ_m$ uniformly at random from the $|\mathcal{W}|!$ possible rankings.

4

The men-proposing DA (MPDA) algorithm repeatedly selects an unassigned man m who in turn proposes to his most preferred woman who has not rejected him yet. If m has been rejected by all women the algorithm assigns m to be unmatched. If has more than one proposal, she rejects all but her most preferred one, leaving the rejected men unassigned. When all men are either assigned (to a woman or unmatched) the algorithm terminates and outputs the matching. Any order of proposals by men will produce the same matching.

Gale and Shapley (1962) showed that the men proposing DA finds the *men-optimal stable matching* (MOSM), in which every man is matched to his most preferred stable woman. The MOSM matches every woman to her least preferred stable man. Likewise, the women proposing DA produces the women-optimal stable matching (WOSM) with symmetric properties. We will be interested in the size of the core as well as the average rankings of partners each side of the market is matched to.

Denote the rank of w in the preference list $\succ_m$ of man m by $\mathrm{Rank}_m(w) = |w' : w' \succeq_m w|$, where m's most preferred woman has a rank of 1. Symmetrically denote the rank of m in the preference list of woman w by $\mathrm{Rank}_w(m)$.

Definition 2.1. *Given a matching μ, we define* men's average rank of wives *as*

$$R_{\mathrm{MEN}}(\mu) = \frac{1}{|\mathcal{M} \backslash \bar{\mathcal{M}}|} \sum_{m \in \mathcal{M} \backslash \bar{\mathcal{M}}} \mathrm{Rank}_m(\mu(m)),$$

where $\bar{\mathcal{M}}$ is the set of men who are unmatched under μ.

Similarly, we define the women's average rank of husbands *as*

$$R_{\mathrm{WOMEN}}(\mu) = \frac{1}{|\mathcal{W} \backslash \bar{\mathcal{W}}|} \sum_{m \in \mathcal{W} \backslash \bar{\mathcal{W}}} \mathrm{Rank}_w(\mu(w)),$$

where $\bar{\mathcal{W}}$ is the set of women who are unmatched under μ.

2.2 Main results

The previous literature analyzed balanced random matching markets which have an equal number of men and women. We start by citing a central result on the structure of stable matching in balanced markets:

Theorem [Pittel (1989a)]. *Consider a sequence of random matching markets with n men and n women, with $n \to \infty$. Then*

$$\frac{R_{\mathrm{MEN}}(\mathrm{MOSM})}{\log n} \xrightarrow{p} 1,$$

$$\frac{R_{\mathrm{MEN}}(\mathrm{WOSM})}{n/\log n} \xrightarrow{p} 1$$

where $\xrightarrow{p}$ denotes convergence in probability. Furthermore, as $n \to \infty$ the fraction of agents that have more than one stable partner converges to 1.

5

So in balanced random markets there are many stable matchings (Pittel (1989b) and Knuth et al. (1990)) and there is advantageous to be on the proposing side. We show that these properties do not extended to unbalanced markets; An unequal numbers of men and women (even a difference of one) strongly favors the short side of the market. We show that in unbalanced markets, in all stable matchings, the men's average rank of wives is low and the women's average rank of husbands is high. Furthermore, we bound the difference between the MOSM and WOSM, showing that all stable matchings are "close".

The following theorem summarizes these results, omitting some quantifications for the sake of readability. In the theorem we characterize a typical realization of the unbalanced matching market. Our statements hold true for a given realization of a random market with a probability that goes to 1 as the size of the market grows large. While our analytical results are for asymptotically large markets, we find that the same features hold for small markets as well. Section 5 presents computational simulations of markets of varying sizes. In section 4 we state and prove a detailed quantitative version of the theorem.

Theorem 1. *Consider a sequence of random matching markets, indexed by n, with n men and $n + k$ women, for arbitrary $1 \leq k = k(n) \leq n/2$. With high probability[4], we have*

(i) In every stable matching μ:

$$R_{\text{MEN}}(\mu) \leq 3 \log(n/k)$$

$$R_{\text{WOMEN}}(\mu) \geq \frac{n}{3 \log(n/k)} \, .$$

(ii) The men are almost as well of under the WOSM as under the MOSM:

$$R_{\text{MEN}}(\text{WOSM}) \leq (1 + o(1)) R_{\text{MEN}}(\text{MOSM}) \, .$$

(iii) The women are almost as badly off under the WOSM as under the MOSM:

$$R_{\text{WOMEN}}(\text{WOSM}) \geq (1 - o(1)) R_{\text{WOMEN}}(\text{MOSM}) \, .$$

(iv) The fraction of men (and fraction of women) who have multiple stable partners is $o(1)$.

Theorem 1 implies that any mechanism, which implements a stable matching, assigns the agents on the short side one of their top choices, while the agents on the long side are either unmatched or assigned to a partner that is ranked not much better than a random partner. Furthermore, the small core property implies that it makes little difference which side is proposing in an unbalanced market. To highlight two particular cases of interest we present the following two immediate corollaries. We first focus on markets with the minimal possible imbalance, where there is only one extra woman.

---

[4]Given a sequence of events $\{\mathcal{E}_n\}$, we say that this sequence occurs *with high probability* (whp) if $\lim_{n \to \infty} \mathbb{P}(\mathcal{E}_n) = 1$.

Corollary 2.2. *Consider a sequence of random matching markets with n men and $n + 1$ women. Then, with high probability, in every stable matching the men's average rank of wives is no more than $3 \log n$ and the women's average rank of husbands is at least $\frac{n}{3 \log n}$. Furthermore, the fraction of men (women) who have multiple stable partners converges to 0 as n grows large.*

To get another point of comparison, consider replacing stable matching with the assignments of the men's random serial dictatorship mechanism (RSD). In the RSD mechanism men are ordered at random, and each man chooses his favorite woman that has yet to be chosen. Thus RSD ignores womens' preferences. Under RSD the men's average rank of wives is $\Theta(\log n)$, and women's average rank of husbands is $\Theta(n/2)$. Under any stable matching men's average rank is of the same order as under RSD, and women's average rank is only better by a factor of $\log n/2$.

Remark 1. *By following our proof method for $k = 1$, one can obtain tighter bounds: whp under all stable matchings the men's average rank of wives is at most $(1 + o(1)) \log n$ and the women's average rank of husbands is at least $(1 - o(1))\frac{n}{\log n}$.*

The next case of interest is a random matching market with a large imbalance between the sizes of the two sides of the markets. We look at $k = \lambda n$ for fixed λ, that is, a matching market with linearly more women than men.

Corollary 2.3. *Let $\lambda > 0$ be any positive constant. Consider a sequence of random matching markets with $|\mathcal{M}| = n$, $|\mathcal{W}| = (1 + \lambda)n$. Then for the constant $\kappa = 5 \log(1/\lambda)$ we have that, with high probability, in every stable matching the average rank of wives is at most κ and the average rank of husbands is at least n/κ. Furthermore, the fraction of men (women) who have multiple stable partners converges to 0 as n grows large.*

When there is a substantial imbalance in the market, the allocation is largely driven by the preferences of the men. The average men gets one of his most preferred women, even under the WOSM. The women on the other hand, are able to exercise little choice. Hypothetically, if a woman was assigned to a uniformly random partner, her average rank of her husband would be $(n + 1)/2$. When there is a large imbalance she can do better only by a factor of less than $\kappa/2$ on average, even under the WOSM. Thus, the men are roughly able to choose, whereas the women are roughly chosen.

To complement our analytical analysis, we conducted simulations on varying market sizes. Figures 1 and 2 illustrate the sharp effect of imbalance in markets with 40 women and a varying number of men, from 20 to 80. In Figure 1 we report the fraction of men who have multiple stable partners, both on average across all markets and the top and bottom 2.5% percentile of draws. For each of 10,000 market realizations we calculate the men's average rankings of wives under MOSM and WOSM and report the average of these in Figure 2. Observe that that even in such small markets, the balanced market with 40 men and 40 women is a knife-edge case. For further simulations see Section 5.

Figure 1: Percent of men with multiple stable partners. The main lines indicate the average over 10,000 realizations. The dotted lines indicate the top and bottom 2.5% of the 10,000 realizations.
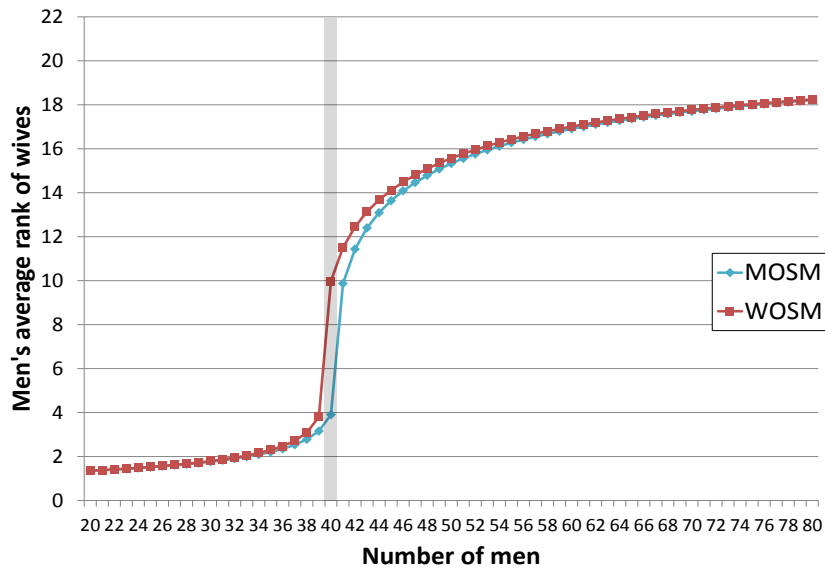

Figure 2: Men's average rank of wives under MOSM and WOSM in a random market with 40 women.

2.3 Intuition and proof idea

Rough intuition for the large advantage of the short side can be gained through the Rural Hospital Theorem (Roth, 1986). Suppose there are $n+1$ women. By the rural hospital theorem the same woman w is unmatched in all stable matchings, and therefore every stable matching must also induce a stable matching for the balanced market (by dropping the unmatched woman w). A stable matching of the balanced market without w remains stable after we add w if none of the men prefers w over his current match. Therefore the average rank of men must be low in all stable matchings. While this is a rough intuition for our proof, we have used a more constructive approach.

We calculate the average rank of men and women through Algorithm 4, which finds the MOSM and WOSM through a series of proposals by men (the algorithm is described in detail in Section 3). The first part of the algorithm runs MPDA to find the MOSM. The second part of the algorithm progresses in a sequence of divorces by women till the WOSM is found. Consider a woman w who is matched to a stable partner m. By divorcing m, w triggers a rejection chain (in which m proposes his most preferred woman whom he hasn't proposed before, possibly displacing another man who proposes in in a similar way). This chain, which we call a *phase*, can be one of two kinds depending on how it terminates: (a) an *improvement phase* that terminates with w accepting a proposal (from a man she prefers over m), and (b) a *terminal phase* that ends with a proposal to an unmatched woman. In an improvement phase a new stable matching results which matches w to a stable partner she prefers over m. On the other hand, a terminal phase implies that m is w's most preferred stable partner.

The intuition for markets of n men and $(1 + \lambda)n$ women is simple. In this case an arbitrary phase beginning with any w is likely to be terminal: the probability that a man in the chain will propose to an unmatched woman is roughly $\frac{\lambda}{1+\lambda}$, while the probability he will propose w is of order $\frac{1}{n}$. Since improvement phases are rare, most women will be matched to the same man under the MOSM and WOSM.

The intuition for markets of n men and $n+1$ women requires some more work. Let S be the set of women for which the algorithm found their best stable men. After finding the MOSM, we can initialize S to contain all unmatched women, since they are unmatched in all stable matchings (by the Rural Hospital Theorem). We show that if a chain reaches any woman $w' \in S$ the phase must be terminal (so there is no need to continue the entire chain till it reaches an unmatched woman), otherwise w' would not have been matched to her best stable partner. We also show that if each woman in the chain of a terminal phase appears only once, every women in this chain is matched to her most preferred stable man, and thus can be added to S. If the chain includes a sub-chain which begins and ends with the same woman w', we call that sub-chain an Internal Improvement Cycle (IIC). An IIC finds an improved stable matching for w', and we show that we can implement the improved stable matching and cut out the sub-chain from the original terminal phase. After removing all IICs, the chain of every terminal phase includes each woman only once, and all women in the chain can be added to S.

Consider starting the first phase terminal in the algorithm with the woman w who received the most proposals. The probability that a proposal will reach w is approximately $\frac{1}{n}$ and the probability

9

that a proposal will reach the unmatched woman is also approximately $\frac{1}{n}$. The phase is likely to be terminal, because the unmatched woman accepts all proposals, while w receives at least $\log(n)$ proposals and will only accept a man which she prefers to all of them. The total length of the chain is likely to be of order $\Theta(n)$, and after we exclude internal improvement cycles, a chain of length $\sim \sqrt{n}$ is left. By the end of the first terminal phase S is already of size $\sim \sqrt{n}$. The probability that a proposal will reach w is still approximately $\frac{1}{n}$, but the probability that a proposal will reach S is much larger. Once S is large almost every phase will be a terminal phase, and almost all women are already matched to their most preferred stable husband.

2.4 Manipulation in unbalanced random matching markets

In this section we consider the implications of our results for strategic agents in matching markets. A *matching mechanism* is a function that takes reported preferences of men and women and produces a matching of men to women. A *stable matching mechanism* is a matching mechanism which produces a matching that is stable with respect to the reported preferences. A matching mechanism induces a direct revelation game, in which each agent reports a preference ranking and receives utility from being assigned to his/her assigned partner.

A mechanism is said to be strategy-proof for men (women) if for every man (woman) it is a dominant strategy to report their true preferences. The men-proposing DA (MPDA) is strategy-proof for men, but some women may benefit from misreporting their preferences. Roth (1982) shows that no stable matching mechanism is strategy-proof for both sides of the market. Demange et al. (1987) show that the scope of strategic behavior is limited - in a stable matching mechanism a woman can never achieve a man she prefers over her husband in the WOSM. In particular, a woman cannot manipulate a stable matching mechanism if she has a unique stable partner (which is her match under both the MSOM and WOSM). In a random, unbalanced matching market most women will have a unique stable partner, and therefore are unable to manipulate their outcome based on their preferences.

Consider the following direct revelation game induced by a stable matching mechanism. Each man m independently draws utilities over matching with every woman $u_m(w) \sim F$. Symmetrically every woman w draws utilities over matching with every woman $u_w(m) \sim G$. We assume that F, G are continuous probability distributions with a finite support $[0, \bar{u}]$. Each agent privately learns his/her own preferences, submits a ranking to the matching mechanism, and receives utility of matching to his/her assigned partner. We say that an agent reports truthfully if he/she submits a ranking list of all possible spouses in order of their utilities.

Theorem 2.4. *Consider a stable matching mechanism, and let $\varepsilon > 0$ and $k > 0$. There exists n_0 such that for any unbalanced market with n men and $n + k$ women for $n > n_0$ it is an ε-Bayes-Nash equilibrium for all agents to report truthfully.*

Note that it is a ε-Bayes-Nash equilibrium for all agents to report truthfully to any stable matching mechanism, regardless of the stable matching selected. While the theorem proves an asymptotic statement, the simulations in Section 5 show that there is little scope for manipulation even in small markets.

10

Proof. From Theorem 1 we know that the expected fraction of women who have multiple stable partners converges to 0 as n tends to infinity. All women are ex ante symmetric, and since preferences are drawn independently and uniformly at random, the women are still symmetric after we reveal the preference list of one woman. Therefore, the interim probability that a woman has multiple stable partners, conditional on her realized preferences, is equal to the expected fraction of woman who have multiple stable partners, and tends to 0 as n tends to infinity. Taking n_0 such that this probability is less than $\varepsilon/\bar{u}$ guarantees that any woman can gain at most ε by misreporting her true preferences. The argument for men is identical. □

A woman w can manipulate a stable matching mechanism by reporting her assigned partner to be unacceptable. This forces the mechanism to match w to a different stable partner, if one exists, or leaving w unmatched, if she has a unique stable partner. Our analysis in section 4 bounds the probability that another stable partner exists, giving us the result above. We believe that by following the analysis in Section 4 more delicately, one can produce tighter bounds on the benefits of manipulation and extend the above results to utilities with unbounded support.

3   Matching algorithm

The goal of this section is to present Algorithm 4 which is the basis of our analysis. This algorithm allows us to calculate the WOSM by a process of successive proposals by men. It first finds the men optimal stable matching using the men-proposing DA algorithm (Algorithm 2 or MPDA), and then progresses to the women optimal stable matching through a series of divorces of matched women followed by men proposals. At the end of this section we show how the run of the algorithm on a random matching market is equivalent to a randomized algorithm. In Section 4 we analyze the randomized algorithm to prove Theorem 1. Throughout our analysis and this section we assume that there are strictly more women than men, that is $|\mathcal{W}| > |\mathcal{M}|$.

For completeness we first state the men proposing DA algorithm, which outputs the MOSM.

Algorithm 2. Men-proposing deferred acceptance

- Input: *Preferences for n men and n + k women.*

- Initialization: *No woman got any proposals, and no man was rejected by any woman.*

- Iterate:

  1. *Proposals: Every unmatched man proposes to his most preferred woman who did not reject him yet (if he prefers her to being unmatched). If no new proposal has been made, output the current matching.*

  2. *Response to proposals: Every woman considers all the men who are proposing to her, temporarily keeping her most preferred man and rejecting all others.*

Before presenting Algorithm 4 we first give a simplified version. The following algorithm (adapted from McVitie and Wilson (1971) and Immorlica and Mahdian (2005)) produces the WOSM from the MOSM by finding each woman's most preferred stable match. It maintains a set S of women whose most preferred stable match has been found. Denote the set of women who are unmatched under the MOSM by $\bar{\mathcal{W}}$.[5] Initialize $S = \bar{\mathcal{W}}$, as by the rural hospital theorem these women are unmatched under any stable matching.

Algorithm 3. MOSM to WOSM (simplified)

- Input: *A matching market with n men and $n + k$ women.*

- Initialization: *Run the men-proposing deferred acceptance to get the men-optimal stable matching μ. Set $S = \bar{\mathcal{W}}$ to be the set of women unmatched under μ. Index all women in $\mathcal{W}$ according to some order. Set $\hat{w}$ to be the women with the lowest index in $\mathcal{W} \backslash S$.*

- New phase:

    1. *Set $\tilde{\mu} \leftarrow \mu$.*
    2. Divorce: *Set $m \leftarrow \mu(\hat{w})$ and have $\hat{w}$ reject m.*
    3. Proposal: *Man m proposes to his most preferred woman w whom he has not yet proposed.[6]*
    4. *w's Decision:*
        (a) *If $w \neq \hat{w}$ prefers her current match $\mu(m)$ to m, or if $w = \hat{w}$ and she prefers $\tilde{\mu}(\hat{w})$ to m, she rejects m. Go to step 3.*
        (b) *If $w \notin \{\hat{w}\} \cup \bar{\mathcal{W}}$, and w prefers m to $\mu(m)$, then w rejects her current partner and accepts m ($m' \leftarrow \mu(w)$, $\mu(w) \leftarrow m$ and $m \leftarrow m'$). Go to step 3.*
        (c) New stable matching: *If $w = \hat{w}$ and $\hat{w}$ prefers m to all her previous proposals, w accepts m and a stable matching is found. Select $\hat{w} \in \mathcal{W} \backslash S$ and start a new phase from step 1.*
        (d) End of terminal phase: *If $w \in \bar{\mathcal{W}}$, restore $\mu \leftarrow \tilde{\mu}$, restore $R(m)$ and add $\hat{w}$ to S. Select $\hat{w} \in \mathcal{W} \backslash S$ if the set is not empty, otherwise terminate and output μ. Start a new phase from step 1.*

It will be convenient to use the following terminology. A **phase** is a sequence of proposals made by the algorithm between visits to step 1. An **improvement phase** is a phase that terminates in step 4(c) (a new stable matching is found). A **terminal phase** is a phase that terminates in step 4(d) (there is no better stable husband for $\hat{w}$). We refer to the sequence of women who reject their husbands in a phase as the **rejection chain**.

---

[5]Since there are more women than men $\bar{\mathcal{W}} \neq \phi$.

[6]Since there are more women than men, such a woman always exists. To account for a man who prefers to be unmatched, add the "empty woman" $\phi \in \bar{\mathcal{W}}$ to represent being unmatched. A man who chooses to be unmatched over any woman he has yet to propose simply chooses $w = \phi \in \bar{\mathcal{W}}$, and the algorithm continues directly to step 4(d).

In each phase the algorithm tries to find a more preferred husband for $\hat{w}$, which requires divorcing $m = \tilde{\mu}(\hat{w})$ and assigning him to another woman. In improvement phases the algorithm finds a more preferred stable husband for $\hat{w}$. In terminal phases the algorithm finds that $\hat{w}$ cannot be assigned a man she prefers over m without creating a blocking pair, and therefore m is $\hat{w}$'s most preferred stable parter.

Proposition 3.1. *Algorithm 3 outputs the women optimal stable matching.*

Proof. The algorithm terminates as each man can make only a finite number of proposals, and there is at most one terminal phase (that is rolled back) per woman. Consider a phase which begins with a stable $\tilde{\mu}$. Immorlica and Mahdian (2005) show that if the phase ends in step 4(c), the matching $\tilde{\mu}$ at the end of the phase is stable as well. By induction, every $\tilde{\mu}$ is a stable matching. Immorlica and Mahdian (2005) also show that if a phase ends in step 4(d) then $\tilde{\mu}(\hat{w})$ is $\hat{w}$'s most preferred stable man. Any subsequent matching $\tilde{\mu}$ is a stable matching that does not $\hat{w}$ worse off, and therefore matches $\hat{w}$ to her most preferred stable man. Finally, any woman in $\bar{\mathcal{W}}$ is unmatched under the WOSM by the rural hospital theorem (Roth (1986)). Thus the algorithm terminates with $\tilde{\mu}$ being the WOSM which matches all women to their most preferred stable husband. $\square$

We follow to refine the algorithm. Since proposals made during terminal phases are rolled back, we can end a phase (and roll back to $\tilde{\mu}$) as soon as we learn that the phase is a terminal phase. During the run of Algorithm 3 each woman in S is matched to her most preferred stable husband. Therefore we can terminate the phase (and roll back) if a woman in S accepts a proposal, as this can only happen in terminal phases. Furthermore, suppose that in a terminal phase the rejection chain includes each woman at most once. We show that every woman in the rejection chain is matched under $\tilde{\mu}$ to her most preferred stable husband, and can therefore be added to S. When the rejection chain includes a woman more than once there are improvement cycles in the chain. We can identify these improvement cycles and implement them as an **Internal Improvement Cycle** (IIC). Specifically, whenever a woman in the chain receives a new proposal we check if she prefers the proposing man over her best stable partner we found so far. If she prefers the proposing man, the part of the rejection chain between this proposal and her best stable partner so far forms an improvement cycle. We implement the IIC by recording the stable matching we found in $\tilde{\mu}$ and removing the cycle from the rejection chain. By removing these cycles from the rejection chain we are left a rejection chain that includes each woman at most once, and can be added to S when the phase is terminal.

Applying these modifications to Algorithm 3 gives us Algorithm 4. It keeps track of the women in the current rejection chain as an ordered set $V = (v_1, v_2, .., v_J)$, and adds all of them to S if the phase is terminal. It also ends a phase as terminal when a woman in S accepts a proposal. The algorithm keeps track of $\nu(w)$ - the current number of proposals received by women w, and $R(m)$ - the set of women who rejected m so far. μ keeps track of the current proposals and $\tilde{\mu}$ records the women-optimal stable matching we found so far.

Algorithm 4. MOSM to WOSM

- Input: *A matching market with n men and $n + k$ women.*

- Initialization: *Run the men-proposing deferred acceptance to get the men-optimal stable matching μ and set $R(m)$ and $\nu(w)$ accordingly. Set $S = \bar{\mathcal{W}}$ to be the set of women unmatched under μ. Index all women in $\mathcal{W}$ according to some order. Set $\hat{w}$ to be the woman in $\mathcal{W} \backslash S$ who received maximum number of proposals so far. Set t to be the total number of proposals made so far.*

- New phase:

   1. *Set $\tilde{\mu} \leftarrow \mu$. Set $V = (v_1) = (\hat{w})$.*

   2. Divorce: *Set $m \leftarrow \mu(\hat{w})$ and have $\hat{w}$ reject m (add $\hat{w}$ to $R(m)$).*

   3. Proposal: *Man m proposes w, his most preferred woman who hasn't rejected him yet.[7] Increment $\nu(w)$ and proposal number t.*

   4. *w's Decision:*

      (a) *If $w \notin V$ prefers $\mu(w)$ to m, or if $w \in V$ and w prefers $\tilde{\mu}(w)$ to m, then reject m (add w to $R(m)$) and return to step 3.*

      (b) *If $w \notin S \cup V$ and w prefers m to $\mu(w)$ then w rejects her current partner. Do $m' \leftarrow \mu(w)$, $\mu(w) \leftarrow m$. Add w to $R(m')$, append w to the end of V. Set $m \leftarrow m'$ and return to step 3.*

      (c) New stable matching: *If $w \in V$ and w prefers m to $\tilde{\mu}(w)$ then we found a stable matching. If $w = \hat{w} = v_1$ do $\mu(\hat{w}) \leftarrow m$. Select $\hat{w} \in \mathcal{W} \backslash S$ and start a new phase from step 1.*
      If $w = v_\ell$ for $\ell > 1$, record her current husband $m' \leftarrow \mu(w)$. Call all proposals made since, and including, m' proposed w an internal improvement cycle (IIC). *Update $\tilde{\mu}$ for the women in the loop by setting $\tilde{\mu}(v_j) \leftarrow \mu(v_j)$ for $j = \ell + 1, ..., J$ and set $\tilde{\mu}(w) \leftarrow m$ and $\mu(w) \leftarrow m$. Remove $v_\ell, .., v_J$ from V, set the proposer $m \leftarrow m'$, decrement $\nu(w)$, decrement t, and return to step 3 in which m will apply to w.*

      (d) End of terminal phase: *If $w \in S$ and w prefers m to $\mu(w)$ then restore $\mu \leftarrow \tilde{\mu}$ and add all the women in V to S. If $S = \mathcal{W}$, terminate and output $\tilde{\mu}$. Otherwise, set $\hat{w}$ to be the woman with smallest index in $\mathcal{W} \backslash S$ and begin a new phase by returning to 1.*

In step 4(c) we found a new stable matching. If the rejection chain cycles back to the original woman, we have an improvement phase. If the rejection chain cycles back to a woman w_ℓ in the middle of the chain we implement the IIC (Implementing the IIC is equivalent to an improvement phase that begins with w_ℓ). Update the best stable matching $\tilde{\mu}$ for all the women in the cycle, and make it the current assignment. Then take m' and make him propose again to w_ℓ, as we changed $\mu(w_\ell)$. Decrement t and $\nu(w_\ell)$ in order not to count this proposal twice.

Proposition 3.2. *Algorithm 4 outputs the women optimal stable matching.*

---

[7]See footnote 6.

Proof. We show that there is a sequence of selections of $\hat{w}$ such that Algorithm 4 is equivalent to Algorithm 3, and therefore both of them output the WOSM. After the initialization step both algorithms are in an identical state. To produce the sequence of selections of $\hat{w}$ we start from an arbitrary sequence, for example index the women and select the one with the lowest index to be $\hat{w}$, as in Algorithm 3. We follow both algorithms in parallel, except that Algorithm 4 skips some of the phases that Algorithm 3 performs. Through the run of the algorithms we find a sequence such that the two algorithms will hold identical $\mu, \tilde{\mu}$ at the end of every phase.

Both algorithms perform the same actions when Algorithm 3 performs steps 4(a), 4(b) and step 4(c) when $w = \hat{w}$. Assume that the algorithms are identical up to proposal $t-1$ and suppose that in proposal t Algorithm 4 performs step 4(c) with $w \neq \hat{w}$, that is, we found an IIC with $w = v_\ell$. We change the sequence of selections of $\hat{w}$ so that v_ℓ is chosen to be $\hat{w}$ just before the current phase. That makes the previous phase an improvement phase for $\hat{w} = v_\ell$ in which the cycle of the IIC is implemented. At the end of the (updated) previous phase both algorithms are identical. In the the current phase both algorithms will continue to be identical until woman v_ℓ receives a proposal, which will happen in step t. In that step the two algorithms will also be identical since $v_\ell \notin V \cup S$ at that time. Thus under the revised ordering the two algorithms are identical up to step t.

Assume that the algorithms are identical up to proposal $t-1$ and suppose that in proposal t Algorithm 4 performs step 4(d), that is, a woman $w \in S$ accepted the proposal. First we show that Algorithm 3 also finds a terminal phase, and thus at the end of the phase both algorithms revert to the same $\tilde{\mu}$. Since $w \in S$ we either have that $w \in \bar{\mathcal{W}}$ or that w was added to S in a previous phase. If $w \in \bar{\mathcal{W}}$ both algorithms will declare a terminal phase and revert to $\tilde{\mu}$. Otherwise, w was added to S in some previous phase as a part of a rejection chain that ended in a woman already in S. Recursively build a rejection chain $C = \{w_1, w_2, ..., w_q\} \subset S$ where $w_1 = w$, $w_q \in \bar{\mathcal{W}}$ such that in the terminal phase in which woman w_j was added to S woman w_j rejected her husband m_j, triggering a series of proposals by m_j that resulted in m_j making a proposal which is accepted by w_{j+1}. At the time of proposal t under Algorithm 4 we have that $\mu(w_j) = m_j$, since at the end of that terminal phase each woman in the rejection chain is reverted to her match under $\tilde{\mu}$, and this man must be the man she rejected during the phase (recall that a woman appears in a rejection chain at most once). Under the induction assumption this is true also under Algorithm 3.

Consider the continuation of the run of Algorithm 3. When $w = w_1$ rejects her husband m_1, he will make proposals in the same order as he did in the terminal phase in which w_1 was added to S under Algorithm 4. All the women which m_1 prefers over w_2 rejected him back then, and since throughout the algorithm, women can only improve to whom they are matched to, they will reject him again in the current phase. Therefore, m_1 will end up proposing to w_2. Since at that point $\mu(w_2) = m_2$ the proposal will be accepted by w_2, making m_2 the new proposer. By induction, m_j will make an accepted proposal to w_{j+1}, until we reach a proposal to $w_q \in \bar{\mathcal{W}}$. At that step Algorithm 3 declares a terminal phase and rolls back to $\tilde{\mu}$, at which point the phase ends with the same $\mu, \tilde{\mu}$ as Algorithm 4.

Next, any phase that is skipped by Algorithm 4 starts with a $\hat{w}$ that is already in S. We have shown before that when a woman in S rejects her husband the resulting rejection chain ends in S, and therefore all such phases are terminal phases and skipping them does not change μ or $\tilde{\mu}$.

15

To complete the proof observe that once we move a woman to be selected as $\hat{w}$ we never change the selected order before and including her. By the end of the run we therefore found a sequence of selections such that the two algorithms hold identical $\mu, \tilde{\mu}$ at the end of every phase, and that every phase that Algorithm 4 skips is a terminal phase. Thus when the Algorithm 4 terminates, it outputs the same matching as Algorithm 3, which is the WOSM by Proposition 3.1. □

The following lemma shows how Algorithm 4 allows us to compare the WOSM and MOSM.

Lemma 3.3. *The difference between the sum of mens' rank of wives under WOSM and the sum of mens' rank of wives under MOSM is equal to the number of proposals in improvement phases and IICs during Algorithm 4.*

Proof. Note that at the end of each terminal phase (Step 4(d)) we roll back all proposals made and return to $\tilde{\mu}$, the matching from the previous phase. Therefor, we can consider only improvement phases and IIC, in which each proposal increases the rank of the proposing man by one. □

3.1 Randomized algorithm

As we are interested in the behavior of Algorithm 4 on a random matching market, we transform the deterministic algorithm on random input into a randomized algorithm which will be easier to analyze. The randomized, or coin flipping, version of the algorithm does not receive preferences as input, but draws them through the process of the algorithm.[8] This is often called the *principle of deferred decisions.*

The algorithm reads the next woman in the preference of a man in step 3 and whether a woman prefers a man over her current proposal in step 4. Since the algorithm ends a phase immediately when a woman $w \in S$ accepts a proposal, no man applies twice to the same woman during the algorithm, and therefore the algorithm never reads previously revealed preferences. When preferences are drawn independently and uniformly at random the distribution of both can be calculated from $R(m)$, the set of women who rejected m, and $\nu(w)$, the number of proposals woman w received so far. In step 3 the randomized algorithm selects the woman w uniformly at random from $\mathcal{W} \backslash R(m)$. In step 4 the probability that w prefers m over her current match can be given directly from $\nu(w)$ for $w \notin S$ or bounded for $w \in S$.[9] Table 1 describes the probabilities for the possible decisions of w.

Note that the event Step 4(a) is the complement of the union of the events in Table 1.

4 Proof of Theorem 1

We will prove the following quantitative version of the main theorem:

---

[8] The initialization step of the randomized version of Algorithm 4 calls the randomized version of Algorithm 2.

[9] The probability that a woman $w \in S$ accepts a man m can be calculated from the number of proposals she received during improvement phases or MPDA and the number of proposals she received during terminal phases. Since the bound on the acceptance probability we calculate from $\nu(w)$ is sufficient for our analysis we omit the additional counters from the algorithm.

Step	Event	Probability
4(b)	$w \notin S \cup \{\hat{w}\}$ prefers m to $\mu(w)$	$\frac{1}{\nu(w)+1}$
4(c)	$\hat{w}$ prefers to $\tilde{\mu}(\hat{w})$	$\frac{1}{\nu(\hat{w})+1}$
4(d)	$w \in S \backslash \bar{\mathcal{W}}$ prefers m to $\mu(w)$	at least $\frac{1}{\nu(w)+1}$

Table 1: Probabilities in a run of Algorithm 4 on a random matching market.

Theorem 5. *Consider a sequence of random matching markets, indexed by n, with n men and $n + k$ women, for arbitrary $1 \leq k = k(n) \leq n/2$. There exists $n_0 < \infty$ such that for all $n > n_0$, with probability at least $1 - \exp\{-(\log n)^{0.4}\}$, we have*

(i) In every stable matching μ:

$$R_{\text{MEN}}(\mu) \leq 3\log(n/k)$$

$$R_{\text{WOMEN}}(\mu) \geq \frac{n}{3\log(n/k)}\,.$$

(ii) The men are almost as well of under the WOSM as under the MOSM:

$$\frac{R_{\text{MEN}}(\text{WOSM})}{R_{\text{MEN}}(\text{MOSM})} \leq 1 + (\log n)^{-0.4}\,.$$

(iii) The women are almost as badly off under the WOSM as under the MOSM:

$$\frac{R_{\text{WOMEN}}(\text{WOSM})}{R_{\text{WOMEN}}(\text{MOSM})} \geq 1 - (\log n)^{-0.4}\,.$$

(iv) Less than $n/(\log n)^{0.4}$ men, and less than $n/(\log n)^{0.4}$ women have multiple stable partners.

Remark 2. *The main theorem states bounds that hold with high probability, but does not bound expectations, e.g., $\mathbb{E}[R_{\text{MEN}}(\text{MOSM})]$ and $\mathbb{E}[\frac{R_{\text{WOMEN}}(\text{WOSM})}{R_{\text{WOMEN}}(\text{MOSM})}]$. Extension of our analysis should yield similar bounds on these quantities as well.*

Definition 4.1. *Given a sequence of events $\{\mathcal{E}_n\}$, we say that this sequence occurs* with very high probability (wvhp) *if*

$$\lim_{n\to\infty} (1 - \mathbb{P}(\mathcal{E}_n))\exp\{(\log n)^{0.4}\} = 0\,.$$

Clearly, it suffices to show that (i)-(iv) in Theorem 5 hold wvhp.

To prove Theorem 5, we analyze the number of proposals in Algorithm 2 followed by Algorithm 4, which will provide us the average rank of wives in the women optimal stable match. We partition the run leading to the WOSM into three parts.

1. **Part I is the run of DA (Algorithm 2)**, which by an analysis similar to that in (Pittel (1989a)), takes no more than $3n\log(n/k)$ proposals wvhp.

2. **Part II are the proposals in Algorithm 4 that take place before the end of first terminal phase.** We show that wvhp,

   - Part II takes no more than $n(\log(n/k))^{0.45}$ proposals.
   - When part II ends the set S contains at least $n^{(1-\varepsilon)/2}$ elements.

3. **Part III are the proposals in Algorithm 4 after Part II that take place until $|S| \geq n^{0.7}$. Thus, this part ends at the end of a terminal phase when $|S|$ exceeds $n^{0.7}$ for the first time.** We show that, wvhp, part III requires $O(n^{0.47})$ phases, and $o(n)$ proposals.

4. **Finally, Part IV includes the remaining proposals from the end of part III until Algorithm 4 terminates or $50n \log n$ total proposals have occurred (including proposals made in Parts I and II), whichever occurs earlier.** Because the set S is large, most phases are terminal phases containing no IICs, and most acceptances lead to eventual inclusion in S. We show that, wvhp, part IV ends with termination of the algorithm, and that the number of proposals in improvement phases and IICs is $o(n)$. But the increase in sum of men's rank of wives from the MOSM to the WOSM is exactly the number of proposals in improvement phases and IICs, yielding the result.

Throughout this section, we consider the preferences on both sides of the market as being revealed sequentially as the algorithm proceeds. We reason about the probabilities of various outcomes at each proposal, conditioned on the history of the algorithm, i.e., the preferences revealed so far. When a proposal occurs to a woman $w \in V$, we reveal the ordering of $\tilde{\mu}(w), m, \mu(w)$ in the preference list of w. There are two possibilities $\mu(w) \succ_w m \succ_w \tilde{\mu}(w)$ and $m \succ_w \mu(w) \succ_w \tilde{\mu}(w)$. Each of these possibilities occurs with probability $1/(\nu(w) + 1)$, where $\nu(w)$ is the number of proposals received by w before the current proposal.

Lemma 4.2. *Consider a man m, who is proposing at step 3 of Algorithm 4. Consider a subset of women $\mathcal{A} \subseteq \mathcal{W} \backslash R(m)$. Let $\nu(\mathcal{A}) = \frac{1}{|\mathcal{A}|} \sum_{\tilde{w} \in \mathcal{A}} \nu(\tilde{w})$ be the average number of proposals received by women in $\mathcal{A}$. The man m proposes to some woman w in the current step. Conditional on $w \in \mathcal{A}$ and all preferences revealed so far, the probability that m is the most preferred man who proposed w so far, is at least $\frac{1}{\nu(\mathcal{A})+1}$.*

Proof. For any woman $\tilde{w} \notin R(m)$ the probability that m is the most preferred man who applied to w so far is $\frac{1}{\nu(\tilde{w})+1}$. Conditional on $w \in \mathcal{A}$, the probability that m is the most preferred man who applied to w so far is at least

$$\frac{1}{|\mathcal{A}|} \sum_{\tilde{w} \in \mathcal{A}} \frac{1}{\nu(\tilde{w}) + 1} \geq \frac{1}{\nu(\mathcal{A}) + 1}$$

by Jensen's inequality. $\qquad\square$

The following lemma will be convenient and its proof is trivial:

Lemma 4.3. *If all men have lists of length $|\mathcal{W}|$ and $|\mathcal{W}| > |\mathcal{M}|$, then no man ever reaches the end of their list in Algorithm 2 or Algorithm 4.*

4.1 Part I

For the analysis in this section we consider the following equivalent version of deferred acceptance:

Algorithm 6. *Index the men $\mathcal{M}$. Initialize $S_{\mathcal{M}} = \phi$, $\bar{\mathcal{W}} = \mathcal{W}$, $\nu(w) = 0 \; \forall w \in \mathcal{W}$ and $R(m) = \phi \; \forall m \in \mathcal{M}$.*

1. *If $\mathcal{M} \backslash S_{\mathcal{M}}$ is empty then terminate. Else, let m be the man with the smallest index in $\mathcal{M} \backslash S_{\mathcal{M}}$. Add m to $S_{\mathcal{M}}$.*

2. *Man m proposes to his most preferred woman w whom he has not yet applied to (increment $\nu(w)$). If he is at the end of his list, go to Step 1.*

3. *Decision of w:*

   - *If $w \in \bar{\mathcal{W}}$, i.e, w is unmatched then she accepts m. Remove w from $\bar{\mathcal{W}}$. Go to Step 1.*
   - *If w is currently matched, she accepts the better of her current match and m and rejects the other. Set m to be the rejected man, add w to $R(m)$ and continue at Step 2.*

Note that the output of Algorithm 6 is the same as the output of Algorithm 2, i.e., it is the man optimal stable match (we have just reordered the proposals). The output of Algorithm 6 is given as an input to Algorithm 4. Again, we think of preferences as being revealed as the algorithm proceeds, with the women only revealing preferences among the set of men who have proposed them so far.

The next lemma establishes upper bounds on the average and maximum rank of men of over women and a lower bound on the women's average rank over men. The upper bound for the worst possible rank of men is due to Pittel (1989a) who obtained this bound in a balanced market (by adding more women to the market men are only becoming better off).

Lemma 4.4. *Fix any $\epsilon > 0$. Let μ be the man optimal stable matching. The following hold wvhp:*

(i) *The men's average rank of wives in μ is at most $2.9 \log(n/k)$ and is at least $0.99((n+k)/n) \log((n+k)/k)$.*

(ii) *$\max_{m \in \mathcal{M}} \mathrm{Rank}_m(\mu(m)) \le 3(\log n)^2$.*

(iii) *The women's average rank of husbands in μ is at least $n/(2.95 \log(n/k))$.*

Proof. Consider a random matching market with n men and n women. Pittel (1989a) showed that the average of men's rank of wives in the men-optimal stable matching is at most $(1 + \epsilon)n \log n$ wvhp and Pittel (1989b, Theorem 6.1) showed that the maximal man's rank of his wife is at most $3(\log n)^2$ wvhp. Crawford (1991) showed that in any (deterministic) matching market all men are weakly better off when a woman is added to the market. Together, we have that in a random matching with n men and at least n women, wvhp, the maximal men's rank of his wife is at most $3(\log n)^2$, establishing (ii).

Also, tracking Algorithm 6, similar to Pittel (1989a), we claim that, wvhp, the sum of men's rank of wives is at most $(1+\epsilon)(n+k) \log((n+k)/k) \le 2.9n \log(n/k)$ for $k \le n/2$ and small enough

$\epsilon > 0$. This claim immediately implies the stated bound (i) on men's average rank of wives. To prove the claim, we use the fact that the number of proposals is stochastically dominated by the number of draws in the coupon collector problem, when there are $n + k$ coupons and we need to draw n distinct coupons. This latter quantity is a sum of Geometric$((n+k-i+1)/(n+k))$ random variables for $i = 1, 2, \ldots, n$. The mean is

$$\sum_{i=1}^{n} \frac{n+k}{n+k-i+1} = (n+k)\Big(\frac{1}{k+1} + \frac{1}{k+2} + \ldots + \frac{1}{n+k}\Big) = (n+k)\log((n+k)/k) + o(n+k).$$

Routine arguments (e.g., see Durrett (2010)) can be used to show that, in fact, this sum exceeds $(1 + \epsilon)(n + k)\log((n + k)/k)$ with probability $\exp(-\Theta(n))$. We can also establish that the sum of men's rank of wives is at least $(1 - \epsilon)(n + k)\log((n + k)/k) \geq (1 - \epsilon)n\log(n/k)$ using similar arguments; Recall that wvhp, no man makes more than $3(\log n)^2$ proposals. It follows that for each proposal that occurs during the search for the i-th unmatched woman, the probability that an unmatched woman is found is at most

$$p_i = \frac{n+k-i+1}{n+k-\min(3(\log n)^2, i-1)}.$$

It follows that the number of proposals needed to find the i-th woman stochastically dominates Geometric(p_i), conditioned on everything so far. It follows that the mean total number of proposals is at least

$$\sum_{i=1}^{n} \frac{n+k-3(\log n)^2}{n+k-i+1} = \big(n+k-3(\log n)^2\big)\Big(\frac{1}{k+1} + \frac{1}{k+2} + \ldots + \frac{1}{n+k}\Big)$$
$$= (n+k)\log((n+k)/k) - o(n+k).$$

Again, routine arguments (e.g., see Durrett (2010)) can be used to show that, in fact, a sum exceed of independent Geometric(p_i) random variables for $i = 1, 2, \ldots, n$ is less than $(1-\epsilon)(n+k)\log((n+k)/k)$ with probability $\exp(-\Theta(n))$. This establishes the lower bound on men's average rank of wives.

Now consider the women's ranks of husbands. For a woman w, who has received $\nu(w)$ proposals in Part I, the rank of her husband is a random variable that depends only on $\nu(w)$, and not anything else revealed so far. We have
$$\mathbb{E}[\mathrm{Rank}_w(\mu(w))] = \frac{n+1}{\nu(w)+1}.$$

Define
$$M \equiv \mathbb{E}\Bigg[\sum_{w \in \mathcal{W} \setminus \bar{\mathcal{W}}} \mathrm{Rank}_w(\mu(w))\Bigg] = (n+1)\sum_{w \in \mathcal{W} \setminus \bar{\mathcal{W}}} \frac{1}{\nu(w)+1}.$$
Using Azuma's inequality (see Durrett (2010)), we have

$$\mathbb{P}((1/n)\sum_{w \in \mathcal{W} \setminus \bar{\mathcal{W}}} \mathrm{Rank}_w(\mu(w)) \geq \frac{M}{n} - \Delta) \leq \exp\Big\{-\frac{n\Delta^2}{2n^2}\Big\},$$

since $\text{Rank}_w(\mu(w)) \in [0, n]$. Plugging in $\Delta = n^{3/4}$ yields

$$\mathbb{P}\Big((1/n) \sum_{w \in \mathcal{W} \setminus \bar{\mathcal{W}}} \text{Rank}_w(\mu(w)) \geq \frac{M}{n} - n^{3/4}\Big) \leq \exp\left\{-\frac{n^{1/2}}{2}\right\}. \tag{1}$$

Using Jensen's inequality in the definition of M, we have

$$M \geq (n+1)n \cdot \frac{1}{1 + (1/n)\sum_{w \in \mathcal{W} \setminus \bar{\mathcal{W}}} \nu(w)}$$

$$\geq \frac{n^2}{2.9 \log(n/k)} \quad \text{wvhp}, \tag{2}$$

where we used (i), i.e., $\sum_{w \in \mathcal{W} \setminus \bar{\mathcal{W}}} \nu(w) \leq 2.9n \log(n/k)$ wvhp.

Combining Eqs. (1) and (2) we obtain (iii). $\qquad\square$

Lemma 4.5. *Suppose $k \leq n^{0.1}$. Then, wvhp, there are fewer than $n^{0.99}$ women who each receive less than $(1/2)\log n$ proposals.*

Proof. Now consider a woman w'. For each proposal, it goes to w' with probability at least $1/(n+k)$, unless the proposing man has already proposed w'. Suppose w' receives fewer than $(\log n)/2 \leq (2/3)(1 - 2\epsilon)\log(n+k)/k$ proposals, where, for instance, we can define $\epsilon = 0.01$. Since, wvhp, each man makes at most $3(\log n)^2$ proposals (Lemma 4.4(ii)), wvhp there are at most $(3/2)(\log n)^3$ proposals by men who have already proposed w'. Using Lemma 4.4(i), wvhp, there are at least $(1-\epsilon)(n+k)\log((n+k)/k)$ proposals in total. It follows that, wvhp, there are at least $(1-\epsilon)(n+k)\log((n+k)/k) - (3/2)(\log n)^3 \geq (1-2\epsilon)(n+k)\log((n+k)/k)$ proposals by men who have not yet proposed w'. Using Fact 6.1 (i), the probability that fewer than $(2/3)(1-2\epsilon)\log((n+k)/k)$ of these proposals go to w' is

$$\mathbb{P}\big(\text{Binomial}((1-2\epsilon)(n+k)\log((n+k)/k), 1/(n+k)) < (2/3)(1-2\epsilon)\log(n+k)/k)\big)$$
$$\leq 2\exp\{-(1-2\epsilon)\log((n+k)/k)/27\} \leq ((n+k)/k)^{-1/28} \leq n^{-0.02},$$

for $k \leq n^{0.1}$. It follows that the expected number of women who receive fewer than $(1/2)\log n$ proposals is no more than $(n+k)n^{-0.02} \leq 2n^{0.98}$. By Markov's inequality, the number of women who receive fewer than $(1/2)\log n$ proposals is no more than $n^{0.99}$, with probability at least $1 - 2n^{0.98}/n^{0.99} = 1 - o(\exp\{-(\log n)^{0.4}\})$. $\qquad\square$

4.2 Part II

Lemma 4.7 below shows that by the end of Part II, the number of proposals by each man is "small", the number of proposals received by each woman is "small", and that the set S is "large". Since S will be large at the end of this part, Part III will terminate "quickly". The next lemma says that wvhp, there will not be too many proposals in Part II (this bound will be assumed in the proof of Lemma 4.7).

Lemma 4.6. *Part II completes in no more than $(n/k)(\log n)^{0.45} \leq n(\log(n/k))^{0.45}$ proposals wvhp.*

Note: Recall $1 \le k \le n/2$, and notice that $f(x) = x/(\log x)^{0.45}$ in monotone increasing for $x \ge 2$. It follows that $f(n/k) \le f(n)$, i.e.,

$$(n/k)/(\log(n/k))^{0.45} \le n/(\log n)^{0.45}$$
$$\Rightarrow \quad (n/k)(\log n)^{0.45} \le n(\log(n/k))^{0.45}$$

Proof. For each proposal (Step 3) in Part II, the probability of Step 4(d), which will end Part II, is the probability that the man m proposes to an unmatched woman $\frac{k}{|\mathcal{W} \setminus R(m)|} \ge \frac{k}{n+k} \ge \frac{2k}{3n}$. Therefore the probability that the number of proposals in part II exceeds $(n/k)(\log n)^{0.45}$ is at most $\left(1 - \frac{2k}{3n}\right)^{(n/k)(\log n)^{0.45}} \le \exp(-(2/3)(\log n)^{0.45}) = o(\exp\{-(\log n)^{0.4}\})$. $\qquad\square$

Lemma 4.7. *Fix any $\varepsilon > 0$. At the end of Part II, the following hold wvhp:*

(i) *No man has applied to a lot of women:*

$$\max_{m \in \mathcal{M}} |R(m)| < n^{\varepsilon}. \tag{3}$$

(ii) *The set S is large: $|S| \ge n^{(1-\varepsilon)/2}$.*

(iii) *No woman received many proposals*

$$\max_{w \in \mathcal{W}} \nu(w) < n^{\varepsilon}. \tag{4}$$

Proof. Using Lemma 4.4, we know that wvhp, there have been no more than $3n \log(n/k)$ proposals and no man has proposed more than $3(\log n)^2$ women in Part I. Assume that these two conditions hold for the rest of the proof.

We begin with (i). We say that a man m starts a **run of proposals** when m is rejected by a woman at step 4(b) or is divorced from $\hat{w}$ at step 2. We say that a *failure* occurs if a man starts more than $(\log n)^2$ runs or if the length of any run exceeds $(\log n)^3$ proposals. We associate a failure with a particular proposal t, when for the first time, a man starts his $(\log n)^2 + 1$-th run, or the proposal is the $(\log n)^3 + 1$-th proposal in the current run.

Consider the number of runs of a given man m. Man m starts at most one run at step 2. The other runs start when the proposing man $m' \ne m$ proposes to the women m is currently matched to and m' is accepted. At any proposal the probability that m' proposes to any particular woman is no more than the probability that he proposes to $\bar{\mathcal{W}}$. Now if the latter happens, Part II ends. Therefore, it follows that the number of runs man m has in part II is stochastically dominated[10] by $1 + \text{Geometric}(1/2)$. Hence, the probability that a man has more than $(\log n)^2$ runs is bounded by $\left(\frac{1}{2}\right)^{(\log n)^2 - 1} \le 1/n^2$, showing that man m has fewer than $(\log n)^2$ runs in Part II with probability at least $1 - 1/n^2$. It follows from a union bound over all men $m \in \mathcal{M}$ that wvhp failure due to number of runs does not occur.

---

[10] A real values random variable with cumulative distribution F_1 is said to be stochastically dominated by another r.v. with cumulative distribution F_2 if $F_2(x) \le F_1(x)$ for all $x \in \mathbb{R}$.

Assume failure did not occur before or at the beginning of a run of man m. The number of proposals man m accumulates until either the run ends or a failure occurs is bounded by

$$(\log n)^2 \cdot (\log n)^3 \le n/2$$

for sufficiently large n. In each proposal in the run before failure, man m proposes to a uniformly random woman in $\mathcal{W} \backslash R(m)$. Since there were at most $4n \log n$ proposals so far, we have that

$$\nu\left(\mathcal{W}\backslash R(m)\right) \le \frac{4n \log n}{n/2} = 8 \log n \,.$$

From Lemma 4.2, we have that the probability of acceptance at each proposal is at least $\frac{1}{\nu(\mathcal{W}\backslash R(m))+1} \ge \frac{1}{8\log n+1}$. Therefore the probability of man m making $(\log n)^3$ proposals without being accepted is bounded by[11]

$$\left(\frac{1}{8\log n + 1}\right)^{(\log n)^3} \le \frac{1}{n^3} \,.$$

Thus, the run has length no more than $(\log n)^3$ with probability at least $1 - 1/n^3$. Now the number of runs is bounded by n^2, so we conclude that wvhp failure due to number of runs does not occur. Finally, assuming no failure,

$$|R(m)| \le (\log n)^2 \cdot (\log n)^3 < n^\varepsilon$$

establishing (i).

We now prove (ii). If $k \ge n^{(1-\varepsilon)/2}$, the set S is already large enough at the beginning of Part II, and there is nothing to prove. Suppose $k < n^{(1-\varepsilon)/2}$. Consider the evolution of $|V|$ during Part II. We first provide some intuition. Part II contains about $n/k = \omega(n^{1/2})$ proposals before it ends. We start with $|V| = 0$, and $|V|$ initially builds up without any new stable matches found. We can estimate the size of $|V|$ when Step 4(c) (new stable match) occurs as follows[12]: Suppose we reach $|V| \sim N$. Consider the next accepted proposal. Ignoring factors of $\log n$, the probability that the woman who accepts is in V is $\sim |V|/n \sim N/n$. Thus, for one of the next N accepted proposals to include a woman in V, we need $N \cdot N/n \sim 1$, i.e., $N \sim \sqrt{n}$. Thus, when $|V|$ reaches a size of about $\sqrt{n}$, then an IIC forms over the next $\sim \sqrt{n}$ proposals, reducing the size of $|V|$. This occurs repeatedly, with $|V|$ converging to an 'equilibrium' distribution with mean of order $\sqrt{n}$, and this distribution has a light tail. Thus, when the phase ends, we expect $|V| \sim \sqrt{n}$.

We now formalize this intuition. Whenever $|V| < n^{(1-\varepsilon)/2}$, for the next proposal, the probability that:

- The proposal goes to a woman in V is less than $2/n^{(1+\varepsilon)/2}$. Such a proposal is necessary to creating an IIC.

- The proposal goes to a woman in $S = \bar{\mathcal{W}}$, is at most $2k/(n+k) \le 2k/n \le 2/(n^{(1+\varepsilon)/2})$. Such a proposal would terminate the phase.

---

[11] Again this inequality holds for large enough n. We omit to explicitly specify the condition "for large enough n" when such inequalities appear subsequently in Section 4.

[12] This analysis is analogous to that of the birthday paradox.

- The proposal goes to a woman in $\mathcal{W}\backslash(S\cup V)$, who accepts it, is at least $1/(5\log n)$, using the fact that there have been no more than $4n\log n$ proposals so far and Lemma 4.2.

Suppose we start with any $|V| < n^{(1-\varepsilon)/2}$, for instance we have $|V| = 0$ at the start of the phase, we claim that with probability at least $1 - 3k/\sqrt{n}$, we reach $|V| = n^{(1-\varepsilon)/2}$ (call this an 'escape') before there is a proposal to S and before $\sqrt{n}$ proposals occur. We prove this claim as follows: There is a proposal to S among the next $\sqrt{n}$ proposals with probability no more than $2k/\sqrt{n}$. Suppose that $|V|$ stays less than $n^{(1-\varepsilon)/2}$. Then there are $n^{\varepsilon/4}$ or more proposals to women in $|V|$ among $\sqrt{n}$ total proposals with probability no more than $2^{-n^{\varepsilon/4}}$ using Fact 6.1 (ii) on Binomial$(\sqrt{n}, 2/n^{(1+\varepsilon)/2})$. Also, the probability that there are less than $n^{1/2-\varepsilon/8}$ proposals accepted by women in $\mathcal{W}\backslash(V\cup S)$ (these women are added to V) is at most $2^{-n^{1/2-\varepsilon/8}}$, using Fact 6.1 (ii) on Binomial$(\sqrt{n}, 1/(5\log n))$, since at each proposal, such a woman is added with probability at least $(1/(5\log n))$. But if there are

- less than $n^{\varepsilon/4}$ proposals to V, each such proposal reducing $|V|$ by at most $n^{1/2-\varepsilon/2}$,

- no proposal to S, and

- at least $n^{1/2-\varepsilon/8}$ women added to V,

then we must reach $|V| = n^{(1-\varepsilon)/2}$. Thus, the overall probability of *not* reaching $|V| = n^{(1-\varepsilon)/2}$ before there is a proposal to S and before $\sqrt{n}$ proposals occur is at most $2k/\sqrt{n} + 2^{-n^{\varepsilon/2}} + 2^{-n^{1/2-\varepsilon/4}} \leq 3k/\sqrt{n}$. In particular, the probability of a failed escape is at most $3k/\sqrt{n}$.

We now bound the number of times $|V|$ reduces from a value larger than $n^{(1-\varepsilon)/2}$ to a value smaller than $n^{(1-\varepsilon)/2}$. Suppose $|V| \geq n^{(1-\varepsilon)/2}$. The probability that a proposal goes to S is at least $k/(n+k) \geq k/(2n)$. The probability that a proposal goes to one of the first $n^{(1-\varepsilon)/2}$ women in $|V|$ is at most $2n^{(1-\varepsilon)/2}/n \leq 2/n^{(1+\varepsilon)/2}$. Thus, the number of times the latter occurs is stochastically dominated by Geometric$(k/(4n^{(1-\varepsilon)/2})) - 1$. Thus, the total number of escapes needed to ensure $|V| \geq n^{(1-\varepsilon)/2}$, including the one at the start of the phase, is stochastically dominated by Geometric$(k/(4n^{(1-\varepsilon)/2}))$, which exceeds $n^{1/2-\varepsilon/4}/k$ with probability at most

$$\left(1 - k/(4n^{(1-\varepsilon)/2})\right)^{n^{1/2-\varepsilon/4}/k} \leq \exp(-n^{\varepsilon/4}/4).$$

Assuming no more than $n^{1/2-\varepsilon/4}/k$ escapes are needed, one of these escapes fails with probability at most $(n^{1/2-\varepsilon/4}/k) \cdot (3k/\sqrt{n}) = 3n^{-\varepsilon/4}$. Thus, the overall probability of $|V| < n^{(1-\varepsilon)/2}$ when the phase ends is bounded by $\exp(-n^{\varepsilon/4}/4) + n^{-\varepsilon/4} = o(\exp\{-(\log n)^{0.4}\})$. Thus, wvhp, $|V| \geq n^{(1-\varepsilon)/2}$ for all phases in Part II, including the terminal phase. This establishes (ii).

Finally, we establish (iii). Again we assume in our proof that Parts I and II end in no more than $4n\log n$ proposals in total, and that (i) holds (if not, we abandon our attempt to establish (iii), but this does not happen wvhp). Fix a woman w. For each proposal, the probability that *she* receives the proposal is no more than $2/n$, using (i). Thus, the total number of proposals she receives is no more than Binomial$(2n\log n, 2/n)$ which is less than n^{ε}, except with probability $2^{-n^{\varepsilon}/(4\log n)}$ by Chernoff bound (see Fact 6.1 in the Appendix). Union bound over the women gives us that (iii) holds wvhp. $\qquad\square$

Lemma 4.8. *Wvhp, the number of accepted proposals in Part II is no more than $n/(2\sqrt{\log n})$ and the improvement in sum of women's ranks of husbands during Part II is no more than $n^2/(2(\log n)^{3/2})$.*

Proof. If $k \geq n^{0.1}$, then we already know that wvhp the number of proposals is no more than $n^{0.95}$ using Lemma 4.6.

If $k < n^{0.1}$, then we know from Lemma 4.5, that, wvhp, fewer than $n^{0.99}$ women each received fewer than $\log n/2$ proposals in Part I. Further, from Lemma 4.7, wvhp, no man has proposed to more than n^ε women in Parts I and II. It follows that for each proposal in Part II, it goes to a woman who has already received $\log n/2$ or more proposals with probability at least $1 - n^{-0.01}/2$. Hence, the probability that the proposal is accepted is at most $2.5/\log n$. But the total number of proposals in Part II, wvhp, is less than $n(\log n)^{0.45}$ from Lemma 4.6. It follows using Fact 6.1 that, wvhp, fewer than $3n/(\log n)^{0.55} \leq n/(10\sqrt{\log n})$ proposals are accepted in Part II.

We now bound the improvement in the sum of women's ranks of husbands. Using Markov's inequality, there are, wvhp, at most $n^{0.995}$ proposals to women who have received fewer than $\log n/2$ proposals so far. The maximum possible improvement in rank from these proposals is $(n+k)n^{0.995} \leq 2n^{1.995}$. The number of proposals accepted by women who have received at least $\log n/2$ proposals so far is, wvhp, at most $n/(10\sqrt{\log n})$, as we showed above. For such a proposal accepted by a woman w' who has received $\nu(w') \geq (\log n)/2$ previous proposals, the expected improvement in rank is

$$\frac{n - \nu(w')}{\nu(w') + 1} - \frac{n - \nu(w') - 1}{\nu(w') + 2} \leq n/((\log n)/2)^2 = 4n/\log n \,.$$

Further the improvement in rank is in the interval $[1, n - 1]$. Thus, the total improvement in rank is stochastically dominated by a sum of independent X_i, for $i = 1, 2, \ldots, n/(10\sqrt{\log n})$, with $\mathbb{E}[X_i] \leq 4n/\log n$ and $1 \leq X_i \leq n - 1$. It follows using Azuma's inequality (see, e.g., Durrett (2010)) that this sum exceeds $n/(10\sqrt{\log n}) \cdot 4n/\log n + n^{1.6} \leq n^2/(2.1(\log n)^{3/2})$ with probability at most

$$2 \exp\left\{ -\frac{(n^{1.6})^2}{2 \cdot n/(10\sqrt{\log n}) \cdot n^2} \right\} = \exp\{-n^{0.1}\} = o(\exp\{-(\log n)^{0.4}\}) \,.$$

Thus, the total improvement in the sum of women's ranks of husbands is, wvhp, no more than $2n^{1.995} + n^2/(2.1(\log n)^{3/2}) \leq n^2/(2(\log n)^{3/2})$. □

4.3 Part III

Let S_{II} be the set S at the end of part II.

The next lemma provides upper bounds (that are achieved wvhp) on the number of proposals each man makes and the number of proposals each woman receives throughout Parts III and IV.

Let $\mathcal{E}_t$ be the event that until proposal t, no man has applied to more than $n^{0.6}$ women in total or to more than $n^{3\varepsilon}$ women in S_{II}, and no woman has received $n^{2\varepsilon}$ or more proposals. Let $\mathcal{E}_\infty$ be the event that these same conditions hold when Part IV ends.

Lemma 4.9. *The event $\mathcal{E}_\infty$ occurs wvhp.*

Proof. By Lemma 4.7, we know that at the end of Part II, no man has made more than n^ε proposals, that $|S_{\mathrm{II}}| \geq n^{(1-\varepsilon)/2}$, and that no woman has received more than n^ε proposals, wvhp. We assume that all these conditions hold.

Fix a man m. We argue that if m makes a successful proposal to a woman in $S_{\mathrm{II}} \cup \backslash \bar{\mathcal{W}}$, then he makes no further proposals in Algorithm 4: If m makes a successful proposal to a woman in S, this ends the phase making the phase a terminal one, man m goes back to the woman to whom he was matched at the beginning of the phase, and this woman becomes a member of S. Thus, if a m makes a successful proposal to a woman in S, he makes no further proposals. In particular, if m makes a successful proposal to a woman in $S_{\mathrm{II}} \backslash \bar{\mathcal{W}}$, then he makes no further proposals.

Suppose man m is proposing in proposal t and that $\mathcal{E}_t$ holds. Then m has not yet applied to at least $3n^{(1-\varepsilon)/2}/4$ women in S_{II}. Hence the probability of applying to a woman in S_{II} is at least $n^{(1-\varepsilon)/2}/2$. Further, since no woman has received $n^{2\varepsilon}$ or more proposals, the probability of the proposal being accepted is at least $1/n^{2\varepsilon}$. Hence, the probability of the proposal going to a woman in S_{II} and being accepted is at least $n^{-3\varepsilon-1/2}$. Hence, the man makes fewer than $n^{0.6}/2$ proposals in Part IV, and proposes to fewer than $n^{3\varepsilon}/2$ additional women in S_{II}, except with probability $\exp(-n^\varepsilon/2)$. Using a union bound over the men, wvhp, no man has applied to more than $n^{0.6}$ women in total or to more than $n^{3\varepsilon}$ women in S_{II} until the end of Part IV.

Fix a woman w. Each time a proposal occurs, since no man has proposed to more than $n^{0.6}$ women (assuming $\mathcal{E}_t$ holds), the probability of the proposal going to w is less than $2/n$. Since there are at most $50n \log n$ proposals in total, the number of proposals received by w in Part III is more than n^ε with probability less than $\mathbb{P}(\mathrm{Binomial}(50n \log n, 2/n) \geq n^\varepsilon) \leq 2^{-n^\varepsilon}$, using Chernoff bounds (see Fact 6.1(ii) in the Appendix). Using a union bound over the women, wvhp, no woman has received more than n^ε proposals until the end of Part IV.

The result follows combining the analyses in the two paragraphs above. □

We now focus on Part III. We show (Lemma 4.10) that for every phase in Part III, whp:

- the phase is a terminal phase, and

- that $|V|$ at the end of the phase is at least $n^{0.25}$.

For each such phase, $|S|$ increases by at least $n^{0.25}$. In addition, we show that phases are short, with the expected length of a phase being $O(n^{1/2+3\varepsilon})$. We infer that, wvhp, we reach $|S| \geq n^{0.7}$, i.e., the end of Part III, in $o(n^{0.47})$ phases, containing $o(n)$ proposals. Lemma 4.11 below formalizes this.

Lemma 4.10. *Assume $|S_{\mathrm{II}}| \geq n^{(1-\varepsilon)/2}$, cf. Lemma 4.7. Consider a phase during Part III. Suppose $\mathbb{I}(\mathcal{E}_t) = 1$ at the start of the phase. Then, whp, either $\mathbb{I}(\mathcal{E}_{t'}) = 0$ at the end of the phase, or we have:*

- *The phase is a terminal phase.*

- *At the end of the phase is at least $|V| \geq n^{0.25}$.*

Proof. Assume $\mathbb{I}(\mathcal{E}_\tau) = 1$ throughout the phase (otherwise there is nothing to prove). Since we are considering a phase during Part III, we know that $|S| < n^{0.7}$. Also, $|S| \geq |S_{\mathrm{II}}| \geq n^{(1-\varepsilon)/2}$ by assumption. For each proposal, there is a probability of at least $|S|n^{-2\varepsilon}/(2n) \geq n^{-1/2-3\varepsilon}$ and at most $2|S|/n \leq 2n^{-0.3}$, that the proposal is to a woman in S and is accepted. It follows that, whp, the phase is a terminal phase, and the number of proposals in the phase is in $[n^{0.28}, n^{0.52}]$. It is easy to see that with probability at least $(1 - 2n^{0.28}/n)^{n^{0.28}} = 1 - o(1)$, all of the first $n^{0.28}$ proposals in the phase are to distinct women, meaning that there are no IICs. For each proposal, the probability of acceptance is at least $1/(1 + n^{2\varepsilon})$, since no woman has received $n^{2\varepsilon}$ proposals, so whp, there are at least $n^{0.25}$ accepted proposals among the first $n^{0.28}$ proposals, using Fact 6.1 (ii) on $\mathrm{Binomial}(\mathrm{n}^{0.26}, 1/(1 + \mathrm{n}^{2\varepsilon}))$. Now, consider the first $n^{0.25}$ women in V. These women receive no further proposals during the phase with a probability at least $(1 - 2n^{0.25}/n)^{n^{0.52}} = 1 - o(1)$. Hence, whp, these women are part of V at the end of the phase, establishing $|V| \geq n^{0.25}$ at the end of the phase as needed. $\qquad\square$

Lemma 4.11. *Wvhp, Part III contains less that $n^{0.99}$ proposals.*

Proof. We first show that the next $n^{0.47}$ phases after the end of Part II complete in fewer than $n^{0.99}$ proposals. Since, wvhp, $|S_{\mathrm{II}}| \geq n^{(1-\varepsilon)/2}$, if $\mathbb{I}(\mathcal{E}_t) = 1$ then for proposal t the probability of ending the phase (due to acceptance by a woman in S) is at least $n^{-1/2-3\varepsilon}$. It follows that either $\mathbb{I}(\mathcal{E}_\infty) = 0$ or wvhp, the next $n^{0.47}$ phases after the end of Part II complete in no more than $n^{0.47+1/2+4\varepsilon} \leq n^{0.99}$ proposals, using Fact 6.1 (ii) on $\mathbb{P}(\mathrm{Binomial}(\mathrm{n}^{0.97+4\varepsilon}, \mathrm{n}^{-1/2-3\varepsilon}) \geq \mathrm{n}^{0.47})$.

Now we show that wvhp, Part III contains fewer than $n^{0.47}$ phases. Suppose this is not the case, then, by our definition of Part III, at most $n^{0.45}$ of these phases increase $|S|$ by $n^{0.25}$ or more. But using Lemma 4.10, either $\mathbb{I}(\mathcal{E}_\infty) = 0$, or this occurs with probability at most

$$\mathbb{P}\big(\mathrm{Binomial}(\mathrm{n}^{0.47}, 1-\varepsilon) \leq \mathrm{n}^{0.45}\big) \leq \mathbb{P}\big(\mathrm{Binomial}(\mathrm{n}^{0.47}, 1/2) \leq \mathrm{n}^{0.47}/4\big) \leq 2\exp(-\mathrm{n}^{0.47}/24)\,,$$

using Fact 6.1 (i). In other words, either $\mathbb{I}(\mathcal{E}_\infty) = 0$ or, wvhp, Part III contains fewer than $n^{0.47}$ phases.

But Lemma 4.9 tells us that $\mathbb{I}(\mathcal{E}_\infty) = 1$ wvhp. Combining the above, we deduce that wvhp, Part III contains fewer than $n^{0.47}$ phases and fewer than $n^{0.99}$ proposals. $\qquad\square$

4.4 Part IV

Lemma 4.12. *Suppose we are at Step 3 (time t) of Algorithm 4 during Part IV, we have $\mathbb{I}(\mathcal{E}_t) = 1$, and man m is proposing. Then, for large enough n, the probability that:*

(i) Man m proposes to S and is accepted is at least $n^{-0.31}$.

(ii) Man m proposes to $\mathcal{W}\backslash(R(m) \cup \hat{w})$ and is accepted is at least $0.9n/t$.

Proof. Note that $|S| \geq n^{0.7}$, whereas by definition of $\mathcal{E}_t$ the man has proposed to no more than $n^{0.6}$ women and no woman has received more than $n^{2\varepsilon}$ proposals. (i) follows from Lemma 4.2.

Proof of (ii): Since m has not applied to more than $n^{0.6}$ women so far, we know that $|\mathcal{W}\backslash(R(m)\cup \hat{w})| \geq 0.95n$. Also, the total number of proposals so far is $t-1$. Using Lemma 4.2, the probability of applying to $\mathcal{W}\backslash(R(m)\cup\hat{w})$ and being accepted is at least

$$\frac{1}{1+(t-1)/(0.95n)} \geq \frac{0.9n}{t},$$

for large enough n, using $t > n$ since Part I itself requires at least n proposals. $\qquad\square$

Lemma 4.13. *Wvhp, Part IV ends due to termination of the algorithm.*

Proof. Suppose Part IV does not end with termination (if not we are done) and that $\mathcal{E}_\infty$ occurs (Lemma 4.9 guarantees this wvhp). Reveal each proposal sequentially.

For $t \leq 40n\log n$, call proposal t a 'seemingly-good' proposal when acceptance by $w' \in \mathcal{W}\backslash(R(m)\cup\hat{w})$ occurs. Denote the set of seemingly good proposals by $\mathcal{A}$. We use Lemma 4.12. For each proposal t, there is a probability at least $0.9n/t$ of it being a seemingly-good proposal, conditioned on the history so far. Define independent $X_t \sim \text{Bernoulli}(0.9n/t)$ for $t = t_0, t_0+1\ldots, 40n\log n$, where t_0 is the first proposal in Part III. Then we can set up a coupling so that proposal $t \in \mathcal{A}$ whenever $X_t = 1$. Now

$$\sum_{t=4n\log n}^{40n\log n} 1/t \geq (0.99)\ln(10) \geq 2.27\,,$$

$$\Rightarrow \qquad \sum_{t=4n\log n}^{40n\log n} \mathbb{E}[X_t] \qquad \geq 2n$$

Using Fact 6.1, we deduce that

$$\sum_{t=4n\log n}^{40n\log n} X_t \geq 7n/4$$

wvhp, implying

$$|\mathcal{A}| \geq \sum_{t=t_0}^{40n\log n} X_t \geq 7n/4 \tag{5}$$

wvhp, since we know that $t_0 \leq 4n\log n$ wvhp using Lemmas 4.4 and 4.6.

We call a seemingly-good proposal $t \leq 40n\log n$ a 'good' proposal if the following conditions are satisfied:

- During the current phase, there is no proposal to a woman in V. In particular, there are no IICs.

- The phase is a terminal phase that ends during Part IV.

We denote the set of good proposals by $\mathcal{G} \subseteq \mathcal{A}$. We now argue that

$$|S_{\text{II}}| \geq |\mathcal{G}|\,, \tag{6}$$

where S_{II} is the set S at the end of Part IV. If $w' \notin (S \cup \bar{w})$, then w' becomes part of S at the end of the phase if the proposal is a good proposal. For each terminal phase, there is exactly one good proposal to a woman in $S \cup \bar{w}$, which we think of as accounting for $\hat{w}$, which also becomes a part of S. Thus, for every good proposal, one woman joins S during Part IV, establishing Eq. (6).

Now consider any phase in Part IV that starts before the $40n \log n$-th proposal. Call such a phase an *early* phase. Using Lemma 4.12, the phase contains more than $n^{0.32}$ proposals with probability at most $(1 - n^{-0.31})^{n^{0.32}} \leq \exp(-n^{0.01}) \leq 1/n^2$. But the total number of early phases is no more than $40n \log n$. It follows that using a union bound that, wvhp, there is no early phase that contains more $n^{0.32}$ proposals.

Now, the probability of a phase containing fewer than $n^{0.32}$ proposals, and containing a proposal to a woman in V is at most $n^{0.32} \cdot 2n^{0.32}/n \leq n^{-0.35}$, since $|V| \leq n^{0.32}$ throughout such a phase. Further, there are at most $40n \log n \leq n^{1.01}$ early phases. It follows, using Fact 6.1 (ii) on Binomial($n^{1.01}, n^{-0.35}$), that the number of early phases containing a proposal to a woman in V is, wvhp, no more than $n^{0.67}$. It follows that, wvhp, no more than $n^{0.67} \cdot n^{0.32} = n^{0.99}$ proposals occur in early phases containing a proposal to V. But all proposals in $\mathcal{A} \backslash \mathcal{G}$ must occur in such phases. We deduce that, wvhp,

$$|\mathcal{A} \backslash \mathcal{G}| \leq n^{0.99}. \tag{7}$$

Combining Eqs. (5) and (7), we deduce that $|\mathcal{G}| \geq 3n/2 \geq |\mathcal{W}|$ wvhp. Plugging in Eq. (6), we obtain $|S_{\text{II}}| \geq |\mathcal{W}|$ at the end of Part IV wvhp, which we interpret[13] as "With wvhp, our assumption that Part IV does not end with termination was incorrect. In other words, Part IV ends with termination wvhp". $\qquad\square$

Lemma 4.14. *The number of proposals in improvement phases and in IICs in Part IV is no more than $n^{0.99}$ wvhp.*

Proof. In the proof of Lemma 4.13, we in fact showed that whp in Part IV, the number of proposals in phases that include a proposal to a woman in V is no more than $n^{0.99}$. (Actually, we showed this bound for 'early' phases, and also showed that wvhp, the algorithm terminates with an early phase, so that all phases in Part IV are early phases). But improvement phases and phases containing IICs must include a proposal to a woman in V. The result follows. $\qquad\square$

Finally, using these lemmas we give the proof of Theorem 5.

Proof of Theorem 5. Using Lemma 3.3, we can calculate the sum of men's rank of wives by summing up the rank under the MOSM and the number of proposals made during improvement phases and IICs during the run of Algorithm 4. By Lemma 4.6, 4.11 and Lemma 4.14, the total number of proposals that occur in improvement phases and IICs (in Parts II-IV) of Algorithm 4 is, wvhp, no more than $(n/k)(\log(n))^{0.45} + 2n^{0.99}$. Using Lemma 3.3, we get that

Sum of men's ranks of wives(WOSM) $-$ Sum of men's ranks of wives(MOSM)

$$\leq (n/k)(\log(n))^{0.45} + 2n^{0.99} \leq 2n \log(n/k)/\sqrt{\log n} \tag{8}$$

---

[13]Recall our initial assumption that Part IV does not end with termination. Finding that $|S| \geq n$ under this assumption simply means that Part IV did, in fact, end with termination.

using $2n^{0.99} \leq n/\sqrt{\log n} \leq n\log(n/k)/\sqrt{\log n}$ and

$$(n/k)/\log(n/k) \leq 2n/\log n \qquad \text{for } k \leq n/2$$
$$\Rightarrow \quad (n/k)(\log(n))^{0.45} \leq n\log(n/k)/\sqrt{\log n},$$

for large enough n. But sum of men's ranks of wives under MOSM $\geq 0.99n\log(n/k)$ from Lemma 4.4 (i). Theorem 5 (ii) follows.

The only agents whose partner changes in going from the MOSM to the WOSM are the ones who make or receive accepted proposals during improvement phases and IICs. But this number on each side of the market is, wvhp, no more than $n/(2\sqrt{\log n})$ in Part II (Lemma 4.8), and no more than $n^{0.99}$ each in Part III (Lemma 4.11) and Part IV (Lemma 4.14, leading to a bound of $n/(2\sqrt{\log n}) + 2n^{0.99} \leq n/\sqrt{\log n}$ on the total number of agents with multiple stable partners on each side of the market, establishing Theorem 5 (iv).

By Lemma 4.8, wvhp, the improvement in the sum of women's rank of husbands in Part II is at most $n^2/(2(\log n)^{3/2})$. In Parts III and IV, whp there are at most $2n^{0.99}$ women who obtain better husbands (see above), and the rank improves by less than n for each of these women, so the improvement in sum of ranks is less than $2n^{1.99}$. It follows that the total improvement in sum of ranks is, wvhp, less than $n^2/(\log n)^{3/2}$. Combining with Lemma 4.4 (iii) we obtain Theorem 5 (iii). Lemma 4.4 (i) and (iii), combined with Eqs. (8) and the bound above, also gives Theorem 5 (i). □

5 Computational experiments

To complement our asymptotic results derived from our analytical analysis we conducted simulations that allowed us to investigate unbalanced matching markets of varying sizes. The simulations show that our results hold for random markets of any size. We also investigated random markets with correlated preferences and show that our results hold for markets with small amount of correlation in preferences. We note a few interesting phenomena that arise when preferences become increasingly correlated.

5.1 Varying the number of men in small markets

The the first computational experiment illustrates the sharp effect of imbalance in small markets. We simulate markets with uniformly random preferences with 40 women and a varying number of men, from 20 men to 80 men. For each market size we randomly drew 10,000 realizations of a matching market by drawing independently a complete preference list for each agent uniformly at random. Figure 1 (Section 2.2) shows that the fraction of men who have multiple stable partners (averaged over all realizations) is small in all markets except the one one with 40 men and 40 women. For each realization we further calculate the men's average rank of wives[14] (excluding unmatched agents) under the MOSM and WOSM and take the average over all 10,000 markets.

---

[14]Women's average rank of husbands is symmetrically given by switching the number of men and women.

Average rank 1 indicates that all non-single men are matched to their most preferred woman, and rank increases as men are matched to less preferred wives.

The resulting graph, plotted in Figure 3, shows that the balanced market is atypical. We highlighted the results for the balanced market, with 40 men and 40 women, which replicate the previous analysis by Pittel (1989a) and Roth and Peranson (1999). In the balanced market there is a big difference between the MOSM and WOSM and the proposing side greatly benefits. However, when the number of men is not exactly 40 we see that the MOSM and WOSM are close and the average rank depends on whether there are more men or more women. When there are less men than women, that is less than 40 men, the men get their top choices under any stable matching.[15] When there are are more men than women, that is more than 40 men, the men are either unmatched or get an average rank that is not much better than a random woman (note that a random woman will give an average rank of 20.5).
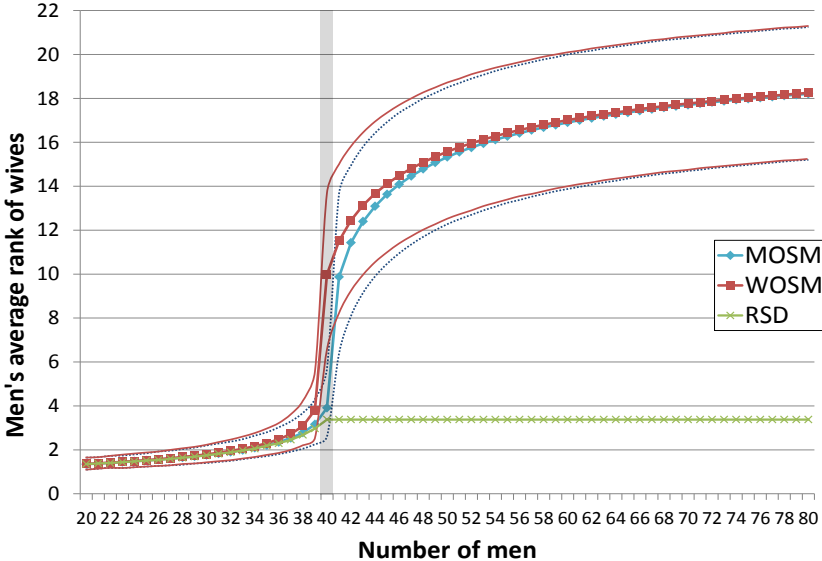
Figure 3: Men's average rank of wives under MOSM and WOSM in a random market with 40 women. The main lines indicate the average over matched mens' average rank of wives in all 10,000 realizations. The dotted lines indicate the top and bottom 2.5% of the 10,000 realizations.

Figure 3 shows that these features hold not only on average, but also in a typical market. To show the concentration of the men's average rank of wives we added dotted lines which trace the 2.5% percentile and 97.5% percentile for both MOSM and WOSM. Figure 3 also includes the men's average rank under random serial dictatorship (RSD). RSD generates a matching by randomly ordering the men and sequentially letting each man choose his most preferred woman out of the women who are left. Note that when there are more women than men, men's average rank of wives

---

[15]We use the fact that the WOSM gives an upper bound on the men's average rank of wives in any stable matching, and the MOSM gives a lower bound.

under all stable matchings is close to their average rank of wives under RSD.

5.2 Larger markets

Table 2 reports simulation results for large markets. Observe that under both WOSM and MOSM the men's average rank of wives is high when there are strictly less women (columns $-10, -5, -1$) and low when there strictly less men (columns $1, 2, 3, 5, 10$). Also note the significant improvement in the average rankings under WOSM when moving from an equal sized market (column 0) to a market where there is one extra woman. (Column 1)

$|\mathcal{M}|$ \ $|\mathcal{W}| - |\mathcal{M}|$		-10	-5	-1	0	1	2	3	5	10
100	MOSM	29.5	27.2	20.3	5.0	4.1	3.7	3.4	3.0	2.6
100	WOSM	30.1	28.2	23.6	20.3	4.9	4.1	3.6	3.2	2.6
200	MOSM	53.6	48.1	35.3	5.7	4.8	4.3	4.1	3.7	3.1
200	WOSM	54.7	49.9	41.0	35.5	5.7	4.7	4.4	3.8	3.2
500	MOSM	115.8	102.6	75.9	6.7	5.7	5.3	5.0	4.5	3.9
500	WOSM	118.0	106.3	86.6	76.2	6.7	5.7	5.3	4.7	4.0
1000	MOSM	203.8	181.4	136.2	7.4	6.4	6.0	5.7	5.2	4.6
1000	WOSM	207.5	187.6	155.1	137.3	7.4	6.5	6.0	5.4	4.7
2000	MOSM	364.5	324.2	249.6	8.1	7.1	6.7	6.3	5.9	5.3
2000	WOSM	370.8	334.4	280.7	249.1	8.1	7.1	6.6	6.1	5.4
5000	MOSM	793.1	713.5	560.0	9.1	8.1	7.6	7.3	6.8	6.2
5000	WOSM	804.7	732.8	622.5	560.2	9.1	8.1	7.6	7.0	6.3

Table 2: Men's average rank of wives in MOSM and WOSM for different market sizes. A man's most preferred wife has rank 1, and higher rank indicates a less preferred wife.

Table 3 presents the fraction of men who have multiple stable partners under varying market sizes.

$|\mathcal{M}|$ \ $|\mathcal{W}| - |\mathcal{M}|$	-10	-5	-1	0	1	2	3	5	10
100	2.1	4.1	15.1	75.3	15.4	9.5	6.5	4.5	2.3
200	2.2	3.8	14.6	83.6	14.6	8.0	6.2	4.1	2.1
500	2.0	3.8	12.6	91.0	13.1	7.1	5.5	3.6	2.0
1000	1.9	3.5	12.3	94.5	12.2	7.2	5.1	3.4	2.0
2000	1.8	3.1	11.1	96.7	11.1	6.1	4.8	2.9	1.7
5000	1.5	2.7	10.1	98.4	10.2	6.0	4.3	2.8	1.5

Table 3: Percentage of men who have multiple stable partners

Table 4 provides simulation results of 1000 draws of large matching markets. Note that even in

a market of a 1000000 men and a 1000001 women the allocation under all stable matchings favors men, who on average get one of their top choices.

N	K=+1			K=+10		
	Men's avg rank under		% Men w. multiple	Men's avg rank under		% Men w. multiple
	MOSM	WOSM	stable partners	MOSM	WOSM	stable partners
10	1.98 (0.45)	2.29 (0.60)	13.84 (18.82)	1.31 (0.20)	1.33 (0.21)	1.19 (5.13)
100	4.09 (0.72)	4.89 (1.08)	15.16 (12.98)	2.55 (0.26)	2.61 (0.27)	2.30 (3.15)
1,000	6.47 (0.79)	7.44 (1.28)	11.9 (10.17)	4.59 (0.30)	4.69 (0.31)	1.95 (2.03)
10,000	8.80 (0.79)	9.80 (1.30)	9.45 (8.30)	6.88 (0.30)	6.98 (0.32)	1.46 (1.47)
100,000	11.11 (0.83)	12.09 (1.31)	7.66 (6.60)	9.16 (0.31)	9.26 (0.32)	1.08 (1.02)
1,000,000	13.4 (0.80)	14.41 (1.27)	6.62 (6.04)	11.46 (0.30)	11.56 (0.32)	0.85 (0.80)

Table 4: Men's average rank in different market sizes with a small imbalance. Standard deviation are given in parentheses.

5.3 Correlated preferences

We ran simulations to investigate matching markets with correlated preferences. We consider a market with 30 men and 40 women and vary the degree of correlation in men's preferences over women. The utility of men m from matching to woman $w \in \{1, ..., 40\}$ is taken to be

$$u_m(w) = \beta \times w + \varepsilon_{mw}$$

where ε_{mw} is an iid draw from the extreme value distribution. We keep woman's preferences to be uniformly random and independent For each value of β we draw a matching market by drawing random preferences for women and ε_{mw} draws for men. We report the men's average rank of wives under the MOSM and WOSM as well as the percent of men with multiple stable partners. For comparison we also report the men's average rank of wives under RSD. The result is report in figure 4.

When $\beta = 0$ we have a uniformly random matching market, in which the average rank under MOSM and WOSM is close, men are almost as well off under any stable matching as they are under RSD and few agents have multiple stable partners. When β is high all men have almost the same preferences over women. The average rank of men is high under all stable matching and under RSD, and few agents have multiple stable partners. For intermediate values of β there is a bigger gap between the men's average rank in stable matchings and their average rank under RSD, but all stable matchings give almost the same men's average rank of wives. Notice that the fraction of men who have multiple stable partners is non monotonic in β. For small values of β an increase in the correlation in preferences increases the expected fraction of men with multiple stable partners. However, this fraction is small (under 5%) in all markets we examined.

To examine whether our results are robust to the addition of small amounts of correlation we highlight the range $\{\beta : 0 \leq \beta * (K + N) \leq 4\}$ in Figure 5. The average rank under RSD can serve both as a benchmark for the average rank and as a measure of the correlation of men's preferences.
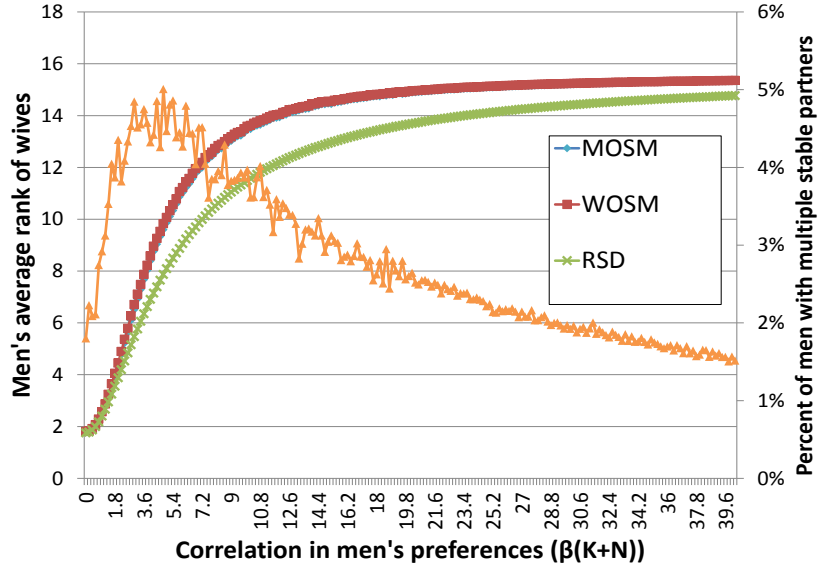
33

Figure 4: Men's average rank of wives and percent of men with multiple stable partners in a market with 30 men 40 women and varying amount of correlation in men's preferences over women. Women's preferences are uniformly random.
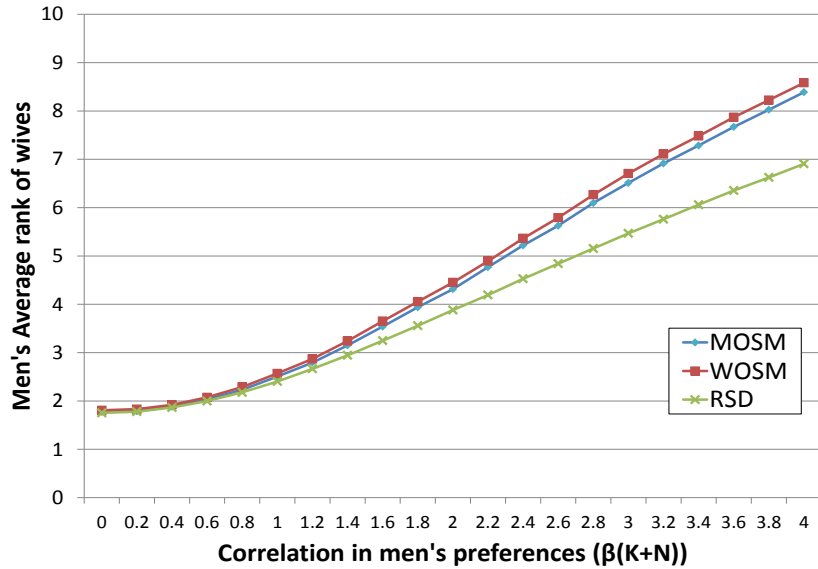

Figure 5: Men's average rank of wives and percent of men with multiple stable partners in a market with 30 men 40 women and varying amount of correlation in men's preferences over women. Women's preferences are uniformly random.

34

We can see that when there is only a slight amount of correlation the MOSM and WOSM are close to each other and both give almost the same average men's rank as RSD. The average percent of men with multiple stable partners (not plotted) is still small as well. This computational experiment suggests our findings are robust to small amounts of correlation.

6 Discussion

We consider random marriage markets in which preferences are drawn independently and uniformly at random. We show that unless the number of men equals the number of women in the market, the men's average ranking of wives is nearly the same across all stable matchings. Consequently, in such markets that adopt the deferred acceptance algorithm, it makes little difference whether men propose (leading to the MOSM) or women propose (leading to the WOSM). The short side of the market has a small average rank of partners whereas the long side of the market has a large average rank of partners, meaning that agents on the short side have a substantial advantage. While we are able to prove our results for large random markets, we also conduct numerical simulations which demonstrate that similar features hold true for small random markets as well. Unlike previous results which showed that the core is small in large markets, our findings are tenable without necessitating a large fraction of unmatched agents and for relatively small markets.

Our results suggest that there is a sharp effect of competition in unbalanced matching markets. This resembles the sharp effect in Bertand competition or other competitive equilibrium models. While this sharp effect may suggest that these models are imperfect, many centralized real matching markets implement the stable mechanisms we analyze. Our model mainly differs from real matching markets in that we analyse a stylized distribution of preferences.

For general distributions over preferences, one cannot identify a "short side" just from the number of men and women in the market. As an example consider a market with 15 men and 20 women in two tiers; 10 women are "top" and 10 women are "middle". Preferences are drawn uniformly at random, except that every man prefers any top woman over any middle woman. It can be shown that every stable matching in this market can be found by selecting a stable matching in the market that comprises the 15 men and the 10 top women, and then a stable matching in the market with the remaining 5 men and the 10 middle women. Our results can be directly applied to both parts, implying that roughly the 10 top women will get to choose, and the remaining 5 men will get to choose from the middle women. Even though there are more men than women, men are not the short side in this market.

Another example is the extreme case in which all men have the same preferences over women. Even if there are n men and $n + k$ women, there is a unique stable matching in which the most preferred women chooses her most preferred man, the second most preferred woman chooses her most preferred man that has yet to be chosen, and so forth. When women's preferences are drawn uniformly at random the (matched) women's average rank of husbands is $O(\log n)$ and the men's average rank of wives is $\Omega(n)$.

We use simulations to explore varying levels of correlation in preferences. We varied the amount of correlation in preferences and found the core to be small. While extreme correlation in preferences

can annul our finding that "the short side chooses", as exemplified above, we find that our findings seem to hold when we add small amounts of correlation to the market. Addressing general random preference structures, which allow for correlation, is left for future work.

To the best of our knowledge, the balanced random matching market was the only known example of a natural model of preferences that leads to a large core. Our results show that this, in fact, is a knife-edge phenomenon. We leave it as a challenging open question to either confirm that matching markets generically have small cores, or to find an example of a natural/realistic matching market that has a large core.

References

A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

A. Abdulkadiroğlu, P.A. Pathak, and A.E. Roth. The new york city high school match. *American Economic Review*, pages 364–367, 2005.

A. Abdulkadiroglu, P.A. Pathak, A.E. Roth, and T. Sonnmez. Changing the boston school choice mechanism. Working paper, 2006.

I. Ashlagi, M. Braverman, and A. Hassidim. Stability in large matching markets with complementarities. Working paper, 2010.

E. D. Azevedo and J. D. Leshno. A Supply and Demand Framework for Two-Sided Matching Markets. Working paper, 2012.

P. Coles and R. Shorrer. Optimal truncation in matching markets. Unpublished, 2012.

V.P. Crawford. Comparative statics in matching markets. *Journal of Economic Theory*, 54(1): 389–400, 1991.

G. Demange, D. Gale, and M. Sotomayor. A further note on the stable matching problem. *Discrete Applied Mathematics*, 16(3):217–222, 1987.

R. Durrett. *Probability: theory and examples*. Cambridge University Press, 2010.

D. Gale and L. L. Shapley. College Admissions and the Stability of Marriage. *American Mathematical Monthly*, 69:9–15, 1962.

R. Holzman and D. Samet. Matching of like rank and the size of the core in the marriage problem. Unpublished, 2013.

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.

D. E. Knuth, R. Motwani, and B. Pittel. Stable husbands. *Proceedings of the first annual ACM-SIAM Symposioum of Discrete Algorithms*, 1:1–14, 1990.

F. Kojima and P. A. Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99:608–627, 2009.

F. Kojima, P. Pathak, and A. E. Roth. Matching with couples: Stability and incentives in large markets. Working paper, 2010.

S. Lee. Incentive Compatibility of Large Centralized Matching Markets. Working paper, 2011.

D.G. McVitie and L.B. Wilson. The stable marriage problem. *Communications of the ACM*, 14(7): 486–490, 1971.

B. Pittel. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, 2(4): 530–549, 1989a.

B. Pittel. On likely solutions of a stable matching problem. In *Proceedings of the 3rd ACM Symposium on Discrete Algorithms*, pages 10–15, 1989b.

A. E. Roth. The Economics of Matching: Stability and Incentives. *Mathematics of Oper- ations Research*, 7:617–628, 1982.

A. E. Roth and E. Peranson. The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design. *American Economic Review*, 89:748–780, 1999.

A.E. Roth. On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica: Journal of the Econometric Society*, pages 425–427, 1986.

Appendix

Fact 6.1 (Chernoff bounds (see Durret (2010))). *Let $X_i \in \{0,1\}$ be independent with $\mathbb{P}[X_i = 1] = \theta_i$ for $1 \leq i \leq n$. Let $X = \sum_{i=1}^{n} X_i$ and $\lambda = \sum_{i=1}^{n} \theta_i$.*

(i) Fix any $\delta \in (0,1)$. Then

$$\mathbb{P}(|X - \lambda| \geq \lambda\delta) \leq 2\exp\{-\delta^2\lambda/3\}. \tag{9}$$

(ii) For any $R \geq 6\lambda$, we have

$$\mathbb{P}(X \geq R) \leq 2^{-R}. \tag{10}$$