

Unbalanced Random Matching Markets: The Stark Effect of Competition

Itai Ashlagi

Stanford University

Yash Kanoria

Columbia University

Jacob D. Leshno

Columbia University

We study competition in matching markets with random heterogeneous preferences and an unequal number of agents on either side. First, we show that even the slightest imbalance yields an essentially unique stable matching. Second, we give a tight description of stable outcomes, showing that matching markets are extremely competitive. Each agent on the short side of the market is matched with one of his top choices, and each agent on the long side either is unmatched or does almost no better than being matched with a random partner. Our results suggest that any matching market is likely to have a small core, explaining why small cores are empirically ubiquitous.

I. Introduction

Stable matching theory has been instrumental in the study and design of numerous two-sided markets. Two-sided markets are described by two dis-

Part of the work was done when the last two authors were postdoctoral researchers at Microsoft Research New England. We thank Nittai Bergman, Gabriel Carroll, Yeon-Koo Che, Ben Golub, Uriel Feige, Ali Hortaçsu, Nicole Immorlica, Fuhito Kojima, SangMok

Electronically published December 16, 2016
[*Journal of Political Economy*, 2017, vol. 125, no. 1]
© 2017 by The University of Chicago. All rights reserved. 0022-3808/2017/12501-0005\$10.00

joint sets of men and women, where each agent has preferences over potential partners. Examples include entry-level labor markets, dating, and college admissions. The core of a two-sided market is the set of stable matchings, where a matching is stable if there are no man and woman who both prefer each other over their assigned partners. Stability is a useful equilibrium concept for these markets since it predicts observed outcomes in decentralized markets and since stability is a critical requirement for the success of centralized clearinghouses.¹

This paper analyzes competition in matching markets to address two fundamental issues. First, we address the long-standing issue of multiplicity of stable matchings.² Previous studies show that the core is small only under restrictive assumptions on market structure, suggesting that the core is generally large. In contrast, matching markets have an essentially unique stable matching in practice.³ Second, relatively little is known about how the structure of stable outcomes is determined by market characteristics. For example, it is known that increasing the number of agents on one side makes agents on the other side weakly better off (Crawford 1991), but little is known about the magnitude of this effect.

We address these issues by looking at randomly drawn matching markets, allowing for competition arising from an unequal number of agents on either side. Influential works by Pittel (1989b) and Roth and Peranson (1999) study the same model with an equal number of agents on both sides and find that the core is typically large. Our first contribution is showing that this is a knife-edge case: the competition resulting from even the slightest imbalance yields an essentially unique stable matching. Our results, which hold for both small and large markets, suggest that any matching market is likely to have a small core, thereby providing an explanation as to why small cores are empirically ubiquitous.

Lee, Mihai Manea, Markus Mobius, Michael Ostrovsky, Parag Pathak, Assaf Romm, Al Roth, Ran Shorrer, two anonymous referees, and multiple seminar participants for helpful discussions and comments. We thank anonymous referees for helpful comments and suggestions. Ashlagi thanks the research support of National Science Foundation grant SES-1254768. Data are provided as supplementary material online.

¹ Hitsch, Hortaçsu, and Ariely (2010) and Banerjee et al. (2013) use stable matchings to predict matching patterns in online dating and the Indian marriage market, respectively. Stable matching models have been successfully adopted in market design contexts such as school choice (Abdulkadiroğlu, Pathak, and Roth 2005; Abdulkadiroğlu et al. 2006) and resident matching programs (Roth and Peranson 1999). Roth and Xing (1994) and Roth (2002) show that stability is important for the success of centralized clearinghouses.

² The potential multiplicity of stable matchings is a central issue in the literature and has led to many studies about the structure of the core (Knuth 1976), which stable matching to implement (Schwarz and Yenmez 2011), and strategic behavior (Dubins and Freedman 1981; Roth 1982).

³ A unique stable matching was reported in the National Resident Matching Program (NRMP; Roth and Peranson 1999), Boston school choice (Pathak and Sönmez 2008), online dating (Hitsch et al. 2010), and the Indian marriage market (Banerjee et al. 2013). We are not aware of any evidence of a large core in a matching market.

Our second contribution is that the essentially unique stable outcome can be almost fully characterized using only the distribution of preferences and the number of agents on each side. Roughly speaking, under any stable matching, each agent on the short side of the market is matched with one of his top choices. Each agent on the long side either is unmatched or does almost no better than being matched with a random partner. Thus, we find that matching markets are extremely competitive, with even the slightest imbalance greatly benefiting the short side. We present simulation results showing that the short side's advantage is robust to small changes in the model.

Formally, we consider a matching market with n men and $n + 1$ women. For each agent we independently draw a complete preference list uniformly at random. We show that with high probability, (i) the core is small in that almost all agents have a unique stable partner (i.e., they are matched with the same partner in all stable matchings), and (ii) under any stable matching men are, on average, matched with their $\log n$ th most preferred woman, while matched women are, on average, matched with their $n/(\log n)$ th most preferred man. Thus agents on the short side rank their partners, on average, as they would if they were to choose partners in sequence.⁴ Matched agents on the long side, on average, rank their partners approximately the same as if they each chose their match only from a limited, randomly drawn set of $\log n$ potential partners. For example, in a market with 1,000 men and 1,001 women, men are matched on average with their seventh ($\approx \log 1,000$) most preferred woman, while women are matched on average with their 145th ($\approx 1,000/\log 1,000$) most preferred man. We further show that the benefit to the short side is amplified when the imbalance is greater.

Given that the smallest imbalance leads to a small core, even when preferences are heterogeneous and uncorrelated, we expect that matching markets will generally have a small core. In real settings preferences are likely to be correlated, but it is generally thought that correlation reduces the size of the core (see, e.g., Roth and Peranson 1999), and this notion is supported by simulation results in Section IV. We simulated markets with varying correlation structures, market sizes, and list lengths, and for all of them we found that imbalance leads to a small core.

A small core implies that there is limited scope for strategic behavior. Suppose that agents report preferences to a central mechanism, which implements a matching that is stable with respect to reported preferences. Demange, Gale, and Sotomayor (1987) consider the induced full

⁴ The corresponding mechanism would be random serial dictatorship, under which men are randomly ordered and each man is assigned in turn to his most preferred woman who was not previously assigned (see, e.g., Abdulkadiroğlu and Sönmez 1998).

information game and show that agents who have a unique stable partner are unable to gain from misreporting their preferences. It follows that in unbalanced markets with incomplete information, truthful reporting is an ε -Bayes-Nash equilibrium.⁵ This suggests that agents should report their preferences truthfully to stable matching mechanisms and may help explain the practical success of stable matching mechanisms.

To gain some intuition for the effect of imbalance in matching markets, it is useful to compare our setting with a competitive, homogeneous buyer-seller market. In a market with 100 homogeneous sellers who have unit supply and a reservation value of zero and 100 identical buyers who have unit demand and value the good at one, every price between zero and one gives a core allocation. However, when there are 101 sellers, competition among sellers implies a unique clearing price of zero, since any buyer has an outside option of buying from the unmatched seller who will sell for any positive price.

Our results show a sharp phenomenon in random matching markets similar to the one in the homogeneous buyer-seller market, despite heterogeneous preferences and the lack of transfers. In particular, the direct outside option argument (used above in the buyer-seller market) does not explain the strong effect of a single additional woman: while the unmatched woman is willing to be matched with any man, she creates a useful outside option only for a few men who rank her favorably. However, this outside option makes these men better off, making their spouses worse off. The men who like these spouses must, in turn, be made better off. This effect ripples through the entire market, making most men better off.

Our proof requires developing some technical tools that may be of independent interest. We build on the work of McVitie and Wilson (1971) and Immorlica and Mahdian (2005) to construct an algorithm that calculates the women-optimal stable matching through a series of rejection chains. Previous analysis of rejection chains by Immorlica and Mahdian (2005) and Kojima and Pathak (2009) analyzed each rejection chain independently. Our new algorithm accounts for the interdependence between different rejection chains, allowing us to analyze the ripple effect generated by a small imbalance. The progress of the modified algorithm on random preferences can be captured by a tractable stochastic process, whose analysis reveals that different chains are likely to be highly connected.

⁵ Furthermore, agents with a unique stable partner will be unaffected by profitable manipulations by other agents.

A. *Related Literature*

Most relevant to our work are the studies by Pittel (1989a) and Knuth, Motwani, and Pittel (1990), who extensively analyze balanced random matching markets. They characterize the set of stable matchings for a random matching market with n men and n women, showing that the men's average rank of wives ranges from $\log n$ to $n/\log n$ in different stable matchings and that the fraction of agents with multiple stable partners approaches one as n grows large. Roughly, our results show that the addition of a single woman makes the core collapse, leaving only the stable matching that is most favorable for men.

Several papers study the size of the core to understand incentives to misreport preferences, analyzing matching markets in which one side has short, randomly drawn preference lists.⁶ Immorlica and Mahdian (2005) show that women cannot manipulate in a one-to-one marriage market, and Kojima and Pathak (2009) show that schools cannot manipulate in many-to-one matching markets. Our results differ in two ways. First, these papers are limited to studying manipulation and the size of the core, while we characterize outcomes and show that competition benefits the short side of the market. Second, their analysis relies on a specific market structure, which generates a large number of unmatched agents.⁷ This market structure was necessary for the analysis in these papers, leading them to conclude that the core is likely to be small only under restrictive assumptions.

Coles and Shorrer (2014) and Lee (forthcoming) study manipulation in asymptotically large balanced matching markets, making different assumptions about the utility functions of agents. Coles and Shorrer (2014) define agents' utilities to be equal to the rank of their spouse (varying between one and n), which grows linearly with the number of agents in the market. They show that women can profitably manipulate the men-proposing deferred acceptance mechanism. Lee (forthcoming) allows for correlation in preferences but assumes that utilities are bounded and the market grows large. He shows that in large markets most agents cannot profitably manipulate the men-proposing deferred acceptance mechanism. The different results stem in part from the different utility parameterizations. In a balanced market an agent's rank of his or her spouse

⁶ The random, short list assumption is motivated by the limited number of interviews as well as the NRMP restriction that allows medical students to submit a rank-ordered list of up to only 30 programs (<http://www.nrmp.org/match-process/create-and-certify-rol-applicants/>) and by the limited number of interviews.

⁷ Their proof requires that randomly drawn preference lists are short enough for agents to have a significant probability of remaining unmatched. Thus, though they assume an equal number of seats and students, their market behaves like a highly unbalanced market. Unless being unmatched is an attractive option, this implies that agents are not submitting long enough lists.

can range from $\log n$ to $n/\log n$, and this difference in rank can be small or large in terms of utilities.⁸

The literature on matching markets with transferable utility provides theoretical predictions on who is matched with whom based on market characteristics. Rao (1993) and Abramitzky, Delavande, and Vasconcelos (2011) exploit random variation in female-male balance to show that outcomes in the marriage market favor the short side. Also related to this paper is the well-known phenomenon that increasing the relative number of agents of one type (“buyers”) benefits agents of the other type (“sellers”) (Shapley and Shubik 1971; Becker 1973). Our paper provides the first quantification of this effect in matching markets without transfers.

Several papers studied markets with strong correlation in preferences. It is well known that if all men have the same preferences over women, there is a unique stable matching. Holzman and Samet (2014) generalize this observation, showing that if the distance between any two men’s preference lists is small, the set of stable matchings is small. Azevedo and Leshno (2016) look at large many-to-one markets with a constant number of schools and an increasing number of students; they find that the set of stable matchings generically converges to the unique stable matching of a continuum model. The core of these markets can equivalently be described as the core of a one-to-one matching market between students and seats in schools, where all seats in a school have identical preferences over students.

B. Organization of the Paper

Section II presents our model and results. Section III provides intuition for the results, outlines our proof, and presents the new matching algorithm that is the basis for our proof. Section IV presents simulation results, which show that the same features occur in small markets, and verifies the robustness of our results. Section V gives some final remarks and discusses the limitations of our model.

Online appendix A proves the correctness of our new matching algorithm. The proof of our main results is in online appendix B. In online appendix C, we discuss how our results may be extended to many-to-one random matching markets.

⁸ Both $\log n/n$ and $1/\log n$ converge to zero, meaning that as the market gets large, both men and women are matched with a partner in the top percentile of their preferences, although the convergence for women is very slow. For example, $\log n/n = 1$ percent for $n \approx 600$, but $1/\log n = 1$ percent for $n \approx 2.68 \times 10^{45}$.

II. Model and Results

A. Random Matching Markets

A two-sided matching market is composed of a set of men $\mathcal{M} = \{1, \dots, n\}$ and a set of women $\mathcal{W} = \{1, \dots, n + k\}$. Each man m has a complete strict preference list \succ_m over the set of women,⁹ 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 $m \in \mathcal{M}$, $\mu(m) \in \mathcal{W} \cup \{m\}$; 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 $\mu(w) = m$. We use $\mu(w) = w$ to denote that woman w is unmatched under μ .

A matching μ is *unstable* if there are 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. It is well known that the core of a matching market is the set of stable matchings. We say that m is a *stable partner* for w (and vice versa) if there is a stable matching in which m is matched with w .

A *random matching market* is generated by drawing a complete preference list for each man and each woman independently and uniformly at random. Thus, for each man m , we draw a complete ranking \succ_m from a uniform distribution over the $|\mathcal{W}|!$ possible rankings.

A stable matching always exists and can be found using the deferred acceptance (DA) algorithm by Gale and Shapley (1962).¹⁰ They show that the men-proposing DA finds the *men-optimal stable matching* (MOSM), in which every man is matched with his most preferred stable woman. The MOSM matches every woman with her least preferred stable man. Likewise, the women-proposing DA produces the *women-optimal stable matching* (WOSM) with symmetric properties.

We are interested in the size of the core, as well as how matched agents rank their assigned partners. Denote the rank of woman w in the preference list \succ_m of man m by $\text{Rank}_m(w) \equiv |\{w' : w' \succ_m w\}|$. A smaller rank is better, and m 's most preferred woman has a rank of 1. Symmetrically, denote the rank of m in the preference list of w by $\text{Rank}_w(m)$.

DEFINITION 1. Given a matching μ , the *men's average rank of wives* is given by

$$R_{\text{MEN}}(\mu) = \frac{1}{|\mathcal{M} \setminus \overline{\mathcal{M}}|} \sum_{m \in \mathcal{M} \setminus \overline{\mathcal{M}}} \text{Rank}_m(\mu(m)),$$

where $\overline{\mathcal{M}}$ is the set of men who are unmatched under μ . Similarly, the *women's average rank of husbands* is given by

⁹ That is, each man prefers any woman over being unmatched.

¹⁰ We describe the DA algorithm in Sec. III.

$$R_{\text{WOMEN}}(\mu) = \frac{1}{|\mathcal{W} \setminus \overline{\mathcal{W}}|} \sum_{w \in \mathcal{W} \setminus \overline{\mathcal{W}}} \text{Rank}_w(\mu(w)),$$

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

We use two metrics for the size of the core. First, we consider the fraction of agents who have multiple stable partners. Second, we consider the difference between the men's average rank of wives under the MOSM and under the WOSM.

B. Previous Results

Previous literature has analyzed balanced random matching markets, which have an equal number of men and women. We start by citing a key result on the structure of stable matchings in balanced markets.

THEOREM (Pittel 1989a). In a random matching market with n men and n women, the fraction of agents who have multiple stable partners converges to one as $n \rightarrow \infty$. Furthermore,

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

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

where \xrightarrow{p} denotes convergence in probability.

This result shows that in a balanced market, the core is large under both measures: most agents have multiple stable partners, and the men's average ranks of wives under the MOSM and the WOSM are significantly different. We find that this does not extend to unbalanced markets.

C. The Size of the Core in Unbalanced Markets

In our main result, we show that in a typical realization of an unbalanced market, almost all agents have a unique stable partner, and the men's average rank of wives and the women's average rank of husbands are almost the same under all stable matchings. We omit quantifications for the sake of readability and give a stronger version of the theorem in online appendix B.

THEOREM 1. Consider a sequence of random matching markets, indexed by n , with n men and $n + k$ women, for arbitrary $k = k(n) \geq 1$. Fix any $\varepsilon > 0$. With high probability,¹¹ we have that

¹¹ Given a sequence of events $\{\mathcal{E}_n\}$, we say that this sequence occurs with high probability if $\lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{E}_n) = 1$.

- i. the fraction of men and the fraction of women who have multiple stable partners are each no more than ε , and
- ii. the men's average rank of wives is almost the same under all stable matchings, as is the women's average rank of husbands:¹²

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

$$R_{\text{WOMEN}}(\text{WOSM}) \geq (1 + \varepsilon)R_{\text{WOMEN}}(\text{MOSM}).$$

For centralized unbalanced markets, this result implies that the choice of the proposing side in DA makes little difference. For decentralized markets, this implies that stability gives an almost unique prediction.

Figures 1 and 2 illustrate the results and show that the same features hold in small unbalanced markets. Figure 1 reports the fraction of men who have multiple stable partners in random markets with 40 women and 20–60 men. Figure 2 plots the men's average rank of wives under MOSM and WOSM. Observe that, even in such small markets, the large core of the balanced market (40 men and 40 women) is a knife-edge case.

D. Characterization of Stable Outcomes

The next theorem shows the advantage of the short side. When there are more women than men, the men's average rank of wives is small, meaning that most men are matched with one of their top choices. On the other hand, the women's average rank of husbands is not much better than that resulting from random assignment. The theorem states the result for a general imbalanced market, and we give simplified expressions for special cases of interest in the next subsection. We give a stronger version of the theorem in online appendix B.

THEOREM 2. Consider a sequence of random matching markets, indexed by n , with n men and $n + k$ women, for arbitrary $k = k(n) \geq 1$. Fix any $\varepsilon > 0$. With high probability, the following hold for every stable matching μ :

$$R_{\text{MEN}}(\mu) \leq (1 + \varepsilon) \left(\frac{n + k}{n} \right) \log \left(\frac{n + k}{k} \right),$$

$$R_{\text{WOMEN}}(\mu) \geq n / \left[1 + (1 + \varepsilon) \left(\frac{n + k}{n} \right) \log \left(\frac{n + k}{k} \right) \right].$$

For comparison, consider the assignments generated by the men's random serial dictatorship (RSD) mechanism. In RSD, men are ordered at random, and each man chooses his favorite woman who has yet to be cho-

¹² For any stable matching μ it follows from the properties of the MOSM and WOSM that $R_{\text{MEN}}(\text{MOSM}) \leq R_{\text{MEN}}(\mu) \leq R_{\text{MEN}}(\text{WOSM})$.

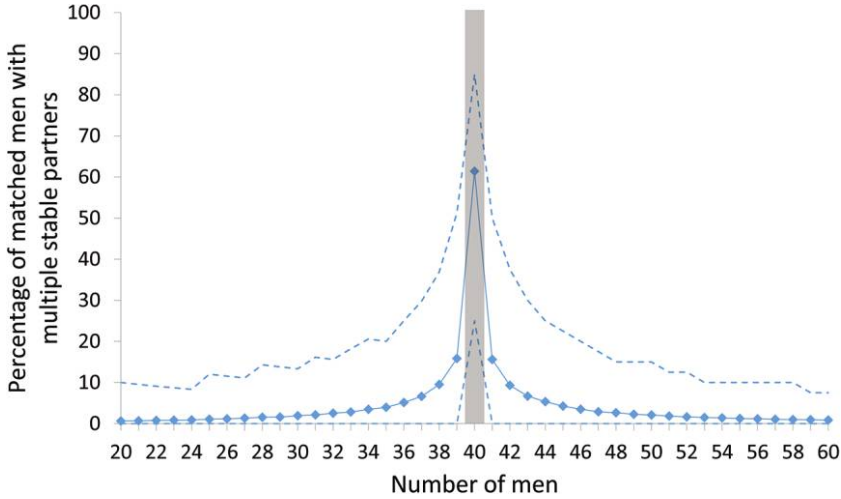


FIG. 1.—Percentage of men with multiple stable partners, in random markets with 40 women and a varying number of men. The main line indicates the average over 10,000 realizations. The dotted lines indicate the top and bottom 2.5th percentiles.

sen (ignoring women's preferences). The men's average rank of wives under RSD is approximately $[(n+k)/n]\log[(n+k)/k]$.¹³ Thus, under any stable matching, the men's average rank of wives would be almost the same as under RSD. The women's average rank of husbands under any stable matching is better than getting a random husband by only a small factor of at most $\log n$. Thus, roughly speaking, in any stable matching, the short side "chooses" while the long side is "chosen."

Figure 2 illustrates the advantage of the short side. When men are on the short side (there are fewer than 40 men), they are matched, on average, with one of their top choices. When men are on the long side, they either are unmatched or rank their partner only slightly better than a random match. See figure 3 below for a comparison with the men's average rank of wives under RSD. Section IV provides simulation results indicat-

¹³ Following an analysis very similar to the proof of lemma B.4(i) in the online appendix, we can show that the men's average rank of wives under RSD is, with high probability, at least

$$(1 - \varepsilon) \left(\frac{n+k}{n} \right) \log \left(\frac{n+k}{k} \right)$$

and at most

$$(1 + \varepsilon) \left(\frac{n+k}{n} \right) \log \left(\frac{n+k}{k} \right).$$

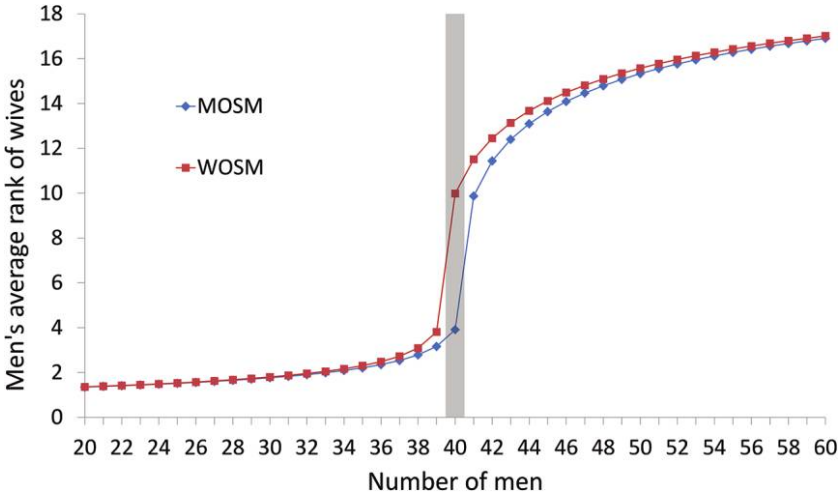


FIG. 2.—Men’s average rank of wives under MOSM and WOSM in random markets with 40 women and a varying number of men. The lines indicate the average over 10,000 realizations.

ing that theorem 2 gives a good approximation for small markets and that the advantage of the short side persists under correlated preferences.

E. Special Cases: Small and Large Imbalances

To highlight two particular cases of interest, we present the following two immediate corollaries. We first focus on markets with minimal imbalance, where there is only one extra woman.

COROLLARY 1. Consider a sequence of random matching markets with n men and $n + 1$ women. Fix any $\epsilon > 0$. With high probability, in every stable matching, the men’s average rank of wives is no more than $(1 + \epsilon) \log n$, the women’s average rank of husbands is at least $n/(1 + \epsilon) \log n$, and the fractions of men and women who have multiple stable partners are each no more than ϵ .

The next case of interest is a random matching market with a large imbalance, taking $k = \lambda n$ for fixed λ .

COROLLARY 2. For $\lambda > 0$, consider a sequence of random matching markets with $|\mathcal{M}| = n, |\mathcal{W}| = (1 + \lambda)n$. Fix any $\epsilon > 0$. Define the constant $\kappa = (1 + \epsilon)(1 + \lambda) \log(1 + 1/\lambda)$. We have that with high probability, in every stable matching, the men’s average rank of wives is at most κ , the women’s average rank of husbands is at least $n/(1 + \kappa)$, and the fractions of men and women who have multiple stable partners are each no more than ϵ .

When there is a substantial imbalance in the market, the allocation is largely driven by men's preferences. For example, in a market with 5 percent extra women, men will be matched, on average, with roughly their third most preferred woman. The women's average rank of husbands under WOSM is only a factor of $(1 + \kappa)/2 = 2.1$ better than being matched with a random man. Thus, the benefit of being on the short side becomes more extreme when the imbalance is bigger.¹⁴

We can examine welfare through fractiles, in addition to rank. In markets with n men and $n + 1$ women, on average, men receive a wife who is at the $(\log n/n)$ th fractile of their preference list, and women receive a husband who is at the $(1/\log n)$ th fractile of their preference list. Note that both sides of the market asymptotically receive their top fractile as n grows large; the rates of convergence are very different (see n. 8). In markets with n men and $(1 + \lambda)n$ women, even as n grows, women do not receive their top fractile. For examination of welfare in matching markets, see subsequent work by Che and Tercieux (2014) and Yariv and Lee (2014).

F. Implications for Strategic Behavior

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

Men-proposing DA (MPDA) is strategy-proof for men, but some women may benefit from misreporting their preferences.¹⁵ Demange et al. (1987) show that a woman cannot profitably manipulate a stable matching mechanism if she has a unique stable partner. In a random unbalanced matching market, most women will have a unique stable partner and therefore cannot gain from misreporting their preferences.

Formally, we consider the following direct revelation game induced by a stable matching mechanism. Each man m independently draws utilities $u_m(w) \sim F$ for matching with each woman w . Symmetrically, each woman w draws utilities $u_w(m) \sim G$ for matching with each man m . We assume that F, G are nonatomic probability distributions with finite support $[0, \bar{u}]$. Each agent privately learns his or her own preferences, submits a

¹⁴ As we would expect, $\kappa = \kappa(\lambda)$ is monotone decreasing in λ and $\lim_{\lambda \rightarrow \infty} \kappa(\lambda) = 1 + \varepsilon$.

¹⁵ A mechanism is said to be strategy-proof for men (women) if it is a dominant strategy for every man (woman) to report preferences truthfully. Roth (1982) shows that no stable matching mechanism is strategy-proof for both sides of the market.

ranking to the matching mechanism, and receives the utility of being matched with his or her assigned partner. We say that an agent reports truthfully if he or she submits a ranked list of spouses in the order of his or her utility of being matched with them.

THEOREM 3. Consider any stable matching mechanism. Let there be a sequence of random matching markets, indexed by n , with n men and $n + k$ women, for arbitrary $k = k(n) \geq 1$. For any $\varepsilon > 0$, there exists n_0 such that for any $n > n_0$, it is an ε -Bayes-Nash equilibrium for all agents to report truthfully.

Note that theorem 3 applies to any stable matching mechanism, regardless of the stable matching it selects. While the theorem is stated for large n , the simulations in Section IV show that there is little scope for manipulation even in small markets.

Proof. For any $\delta > 0$, theorem 1 tells us that the expected fraction of women with multiple stable partners is no more than δ with probability at least $1 - \delta$, for large enough n . Thus, the expected fraction of women with multiple stable partners is bounded by 2δ for large enough n . 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 women with multiple stable partners and thus bounded by 2δ . Choosing $\delta \leq \varepsilon/(2\bar{u})$ guarantees that any woman can gain at most ε by misreporting her true preferences for large enough n . The argument for men is identical. QED

Using the stronger version of theorem 1 (theorem 1 in online app. B), one can produce tighter bounds on the gains from manipulation and extend the above result to utilities with unbounded support.

III. Proof Idea and Algorithm

This section presents the intuition for theorems 1 and 2. We first provide intuition for the advantage of the short side in Section III.A. Our proof (provided in online apps. A and B) follows a different and more constructive approach, which uses a new matching algorithm to trace the structure of the set of stable matchings. We present the main ideas of the proof in Section III.B and the new matching algorithm in Section III.C.

A. Intuition for the Small Core and Advantage of the Short Side

Suppose that there are $n + 1$ women. By the rural hospital theorem (Roth 1986), the same woman \bar{w} is unmatched in all stable matchings. Thus, every stable matching must also induce a stable matching for the balanced

market that results from dropping \bar{w} . A stable matching of the balanced market remains stable after we add \bar{w} only if all men prefer their assigned match over \bar{w} . Intuitively, a matching is more likely to satisfy this constraint when the men are better off. Typically only a few stable matchings (that are all close to the MOSM of the balanced market) satisfy this constraint, and the core of the original market is small.

B. Sketch of the Proof

We prove our results by calculating both the MOSM and WOSM through a sequence of proposals by men, as specified in algorithm 2 below. Since men receive a low average rank, the run of the algorithm on a randomly drawn market is a short and tractable stochastic process. By analyzing this stochastic process we uncover the desired properties of the core.

Algorithm 2 calculates the WOSM through the use of rejection chains. The algorithm maintains a stable matching, initiated to be the MOSM, and in each phase attempts to move to a stable matching that is more preferred by women. A phase starts by selecting a woman \hat{w} and having her divorce her husband m . This triggers a rejection chain in which man m continues to propose, possibly displacing other men who propose in turn. We name the phase according to how the chain ends: (a) an *improvement phase* ends with \hat{w} accepting a proposal from a man she prefers over m , and (b) a *terminal phase* ends with a proposal to an unmatched woman. An improvement phase finds a new stable matching that matches \hat{w} with a partner she prefers over m . A terminal phase implies that m is \hat{w} 's most preferred stable husband (and hence we roll it back before proceeding). Once we have found a terminal phase for every woman, the algorithm is at the WOSM and terminates.

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

The analysis for markets of n men and $n + 1$ women is more involved and requires us to use links between different rejection chains.¹⁶ Denote by S the set of women for whom the algorithm has already found their most preferred stable husbands. We initialize S to contain all unmatched women (by the rural hospital theorem, they are unmatched in all stable matchings). When a terminal phase visits each woman at most once, each

¹⁶ In our proof, we simultaneously consider all possible values of k . However, we discuss the special case of $k = 1$ here for simplicity of exposition.

woman in the chain can be added to S , as divorcing their husbands will result in a subchain that is also terminal.¹⁷ When a new chain reaches a woman $w' \in S$, that new chain must be terminal, since reaching $w' \in S$ implies that the new chain merges with a previous terminal chain. Thus, the set S allows us to track rejection chains and how they merge together.

We first track the progress of the algorithm until the end of the first terminal phase, which occurs with the first proposal to the unmatched woman. Each proposal by a man has about $1/n$ chance to go to the unmatched woman, and therefore, the first proposal to the unmatched woman occurs after order n proposals. That terminal phase will involve a chain of order \sqrt{n} distinct women, who are all added to S .¹⁸ After this, the likelihood that a proposal goes to a woman in S is substantially increased, meaning that most proposals are part of terminal phases, and the set S rapidly grows larger. Once S is large, almost every phase is terminal, and almost all women are already matched with their most preferred stable husband. The algorithm terminates when $S = \mathcal{W}$ is reached.

We calculate the men's average rank of wives under the MOSM using methods similar to those in Pittel (1989a), showing that the number of proposals by men is roughly equal to the solution of the coupon collector's problem. The men's average rank under the WOSM is calculated by adding the proposals during improvement phases of algorithm 2. We deduce bounds on the women's average rank of husbands from the number of proposals each woman receives.

C. Matching Algorithm

This section presents algorithm 2, which calculates the WOSM from the MOSM through a sequence of proposals by men. This algorithm can be used to quickly calculate the WOSM when women are on the long side. All algorithms in this section assume that there are strictly more women than men (i.e., $|\mathcal{W}| > |\mathcal{M}|$). For proofs and further details, refer to online appendix A.

For completeness, we first present the MPDA algorithm that outputs the MOSM.

ALGORITHM 1 (Men-proposing deferred acceptance [MPDA]).

1. Every unmatched man proposes to his most preferred woman who has not already rejected him. If no new proposal is made, output the current matching.

¹⁷ If a chain includes a subchain that begins and ends with the same woman w' , we call that subchain an internal improvement cycle (IIC). An IIC is equivalent to a separate improvement phase for w' , which can be cut out of the original chain. After all IICs are removed, the chain of every terminal phase includes each woman at most once.

¹⁸ These are the women who remain after removal of IICs; cf. n. 17.

2. If a woman has multiple proposals, she tentatively keeps her most preferred man and rejects the rest. Go to step 1.

We now present algorithm 2, which calculates the WOSM. It maintains the most recent stable matching $\tilde{\mu}$ and a set S of women whose most preferred stable match has already been found. In each phase a rejection chain is initiated to check whether there is another stable matching preferred by women. The variable m is used to hold the proposing man and μ is used to hold the temporary assignment. Each rejection chain is tracked using an ordered set $V = (v_1, v_2, \dots, v_j)$ of women. If a proposal is accepted by a woman in V , the algorithm updates $\tilde{\mu}$ to a new stable matching that makes women better off. When a proposal is accepted by a woman in S , the phase is terminal, and all the women in V are added to S . We use the notation $x \leftarrow y$ for the operation of copying the value of variable y to variable x .

ALGORITHM 2 (MOSM to WOSM).

- Input: A matching market with n men and $n + k$ women.
- Initialization: Run the MPDA algorithm to get the men-optimal stable matching μ . Initialize S to be the set of women unmatched under μ . Select any $\hat{w} \in \mathcal{W} \setminus S$.
- New phase:
 1. Set $\tilde{\mu} \leftarrow \mu$. Set $v_1 \leftarrow \hat{w}$ and $V \leftarrow (\hat{w})$.
 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 who has not already rejected him.¹⁹
 4. Woman w 's decision:
 - a. If $w \notin V$ and w prefers $\mu(w)$ over m , or if $w \in V$ and w prefers $\tilde{\mu}(w)$ over m , then w rejects m . Go to step 3.
 - b. If $w \notin S \cup V$ and w prefers m over $\mu(w)$, then w rejects her current partner. Set $m' \leftarrow \mu(w)$, $\mu(w) \leftarrow m$. Append w to the end of V . Set $m \leftarrow m'$ and go to step 3.
 - c. New stable matching: If $w \in V$ and w prefers m over $\tilde{\mu}(w)$, then we have found a stable matching. If $w = \hat{w} = v_1$, set $\mu(\hat{w}) \leftarrow m$. Select $\hat{w} \in \mathcal{W} \setminus S$ and start a new phase from step 1. If $w = v_\ell$ for $\ell > 1$, record her current husband as $m' \leftarrow \mu(w)$. Call the set of all proposals made after the proposal of m' to w an IIC. Set $\mu(w) \leftarrow m$ and update $\tilde{\mu}$ for the women in the loop by setting $\tilde{\mu}(v_j) \leftarrow \mu(v_j)$ for $j = \ell, \ell + 1, \dots, J$. Remove v_ℓ, \dots, v_j from V , set

¹⁹ In markets with complete preferences and strictly more women than men, such a woman always exists. For general markets, if m prefers to be unmatched over any woman who has not already rejected him, the algorithm continues directly to step 4(d).

- the proposer $m \leftarrow m'$, and return to step 3, in which m (earlier m') will again propose to w .
- d. End of terminal phase: If $w \in S$ and w prefers m over $\mu(w)$, 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, select $\hat{w} \in \mathcal{W} \setminus S$ and begin a new phase from step 1.

IV. Computational Experiments

This section presents simulation results that complement our theoretical results. We first present simulation results for small markets and test the accuracy of estimates based on our theoretical findings. We then investigate the effects of correlation in preferences and test the robustness of our results. Finally, we present simulation results for unbalanced many-to-one matching markets.

A. Numerical Results for Our Model

The first computational experiment illustrates the sharp effect of imbalance in a small market. We consider markets with 40 women and a varying number of men, from 20 to 60 men. For each market size we simulate 10,000 realizations by drawing uniformly random complete preferences independently for each agent. For each realization we compute the MOSM and WOSM. Figure 1 (Sec. II) reports the fraction of men who have multiple stable partners (averaged across realizations, as well as the top and bottom 2.5th percentiles); this fraction is small in all unbalanced markets.

Figure 3 reports an average across realizations of the men's average rank of wives under the MOSM and WOSM as well as under RSD.²⁰ The results for the balanced market (40 men and 40 women) replicate the previous analysis by Pittel (1989a) and Roth and Peranson (1999). But in any unbalanced market, the men's average rank of wives is almost the same under the MOSM and WOSM. When there are fewer men than women (i.e., fewer than 40 men), the men's average rank of wives under any stable matching is almost the same as under RSD, with most men receiving one of their top choices. When there are more men than women in the market, the men's average rank of wives is not much better than 20.5, which would be the result of a random assignment.

Tables 1 and 2 report simulation results for markets with varying size and imbalance. Table 1 reports the men's average rank of wives under

²⁰ The women's average rank of husbands is symmetrically given by switching the number of men and women. In the RSD mechanism men are ordered at random, and each man chooses his favorite woman who has yet to be chosen (thus RSD ignores women's preferences).

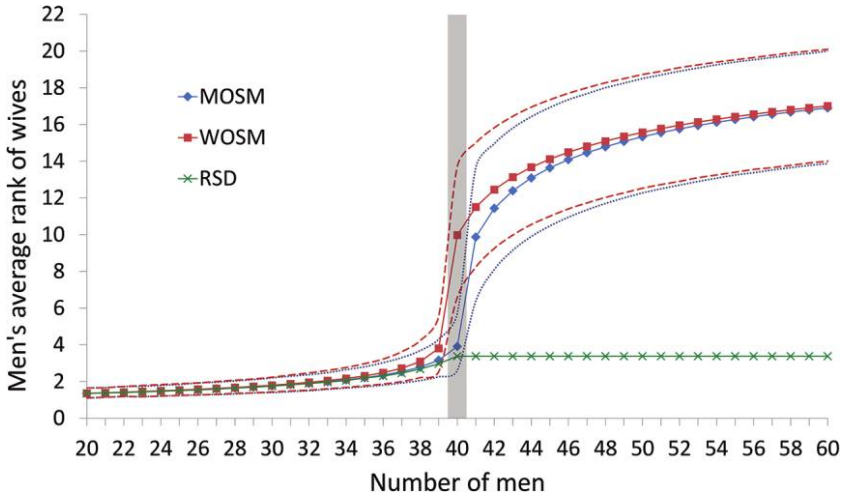


FIG. 3.—Men's average rank of wives under MOSM and WOSM in random markets with 40 women and a varying number of men. The main lines indicate the average over matched men's average rank of wives in all 10,000 realizations. The dotted lines indicate the top and bottom 2.5th percentiles of the 10,000 realizations. The line labeled RSD gives the men's average rank under the RSD mechanism.

the MOSM and WOSM, showing the sharp effect of imbalance across different market sizes. Under both the MOSM and WOSM, the men's average rank of wives is large when there are strictly fewer women (cols. -10 , -5 , -1) and small when there strictly fewer men (cols. 1, 2, 3, 5, 10). In addition, table 1 reports theoretical estimates of the average rank based on the asymptotic results in Section II. The estimate from our theorem is asymptotically accurate for the men's average rank of wives when men are on the short side, while it is only an asymptotic lower bound when men are the long side (see online app. D for details). Observe that when men are on the short side of the market (the right side of the table), our estimate gives a surprisingly good approximation. When men are on the long side (the left side of the table), our estimate is only about 10–25 percent below the true values. Table 2 presents the percentage of men who have multiple stable partners. This percentage is large in balanced markets but is small in all unbalanced markets.

Using the matching algorithm from Section III, we were able to run simulations of rather large unbalanced matching markets. Table 3 provides numerical results of 1,000 draws of matching markets of different sizes. The results show that in markets of any size, a slight imbalance leads to a small core and agents on the short side are typically matched with one of their top choices.

TABLE 1
MEN'S AVERAGE RANK OF WIVES UNDER THE MOSM AND WOSM
FOR DIFFERENT MARKET SIZES

$ \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
WOSM	30.1	28.2	23.6	20.3	4.9	4.1	3.6	3.2	2.6
EST	25.3	22.9	17.5		4.7	4.0	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
WOSM	54.7	49.9	41.0	35.5	5.7	4.7	4.4	3.8	3.2
EST	45.7	40.8	31.5		5.3	4.7	4.3	3.8	3.2
500:									
MOSM	115.8	102.6	75.9	6.7	5.7	5.3	5.0	4.5	3.9
WOSM	118.0	106.3	86.6	76.2	6.7	5.7	5.3	4.7	4.0
EST	98.2	87.6	69.0		6.2	5.5	5.2	4.7	4.0
1,000:									
MOSM	203.8	181.4	136.2	7.4	6.4	6.0	5.7	5.2	4.6
WOSM	207.5	187.6	155.1	137.3	7.4	6.5	6.0	5.4	4.7
EST	175.2	157.3	126.2		6.9	6.2	5.8	5.3	4.7
2,000:									
MOSM	364.5	324.2	249.6	8.1	7.1	6.7	6.3	5.9	5.3
WOSM	370.8	334.4	280.7	249.1	8.1	7.1	6.6	6.1	5.4
EST	314.6	284.7	232.3		7.6	6.9	6.5	6.0	5.3
5,000:									
MOSM	793.1	713.5	560.0	9.1	8.1	7.6	7.3	6.8	6.2
WOSM	804.7	732.8	622.5	560.2	9.1	8.1	7.6	7.0	6.3
EST	690.5	631.1	525.2		8.5	7.8	7.4	6.9	6.2

NOTE.—The numbers for each market size are averages over 1,000 realizations. A man's most preferred wife has rank 1, and a larger rank indicates a less preferred wife. EST is the theoretical estimate of the men's average rank based on theorem 2 (see online app. D for details).

B. *Size of the Core under Correlated Preferences*

This section presents simulation results to examine the effects of correlation in preferences on the size of the core. We simulated a large class of distributions and found a large core only under balanced markets with

TABLE 2
PERCENTAGE OF MEN WHO HAVE MULTIPLE STABLE PARTNERS
FOR DIFFERENT 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
1,000	1.9	3.5	12.3	94.5	12.2	7.2	5.1	3.4	2.0
2,000	1.8	3.1	11.1	96.7	11.1	6.1	4.8	2.9	1.7
5,000	1.5	2.7	10.1	98.4	10.2	6.0	4.3	2.8	1.5

NOTE.—The numbers for each market size are averages over 1,000 realizations.

TABLE 3
MEN'S AVERAGE RANK IN DIFFERENT MARKET SIZES WITH A SMALL IMBALANCE

\mathcal{M}	\mathcal{W} - \mathcal{M} = +1			\mathcal{W} - \mathcal{M} = +10		
	Men's Average Rank under		% Men with Multiple Stable Partners	Men's Average Rank under		% Men with Multiple Stable Partners
	MOSM	WOSM		MOSM	WOSM	
10	2.0 (.4)	2.3 (.6)	13.8 (18.8)	1.3 (.2)	1.3 (.2)	1.2 (5.1)
100	4.1 (.7)	4.9 (1.1)	15.2 (13.0)	2.5 (.3)	2.6 (.3)	2.3 (3.1)
1,000	6.5 (.8)	7.4 (1.3)	11.9 (10.2)	4.6 (.3)	4.7 (.3)	1.9 (2.0)
10,000	8.8 (.8)	9.8 (1.3)	9.4 (8.3)	6.9 (.3)	7.0 (.3)	1.5 (1.5)
100,000	11.1 (.8)	12.1 (1.3)	7.7 (6.6)	9.1 (.3)	9.3 (.3)	1.1 (1.0)
1,000,000	13.4 (.8)	14.4 (1.3)	6.6 (6.0)	11.5 (.3)	11.6 (.3)	.8 (.8)

NOTE.—The numbers for each market size are averages over 1,000 realizations. Standard deviations are given in parentheses.

little correlation in preferences.²¹ Our simulations suggest that the core generally becomes smaller as preferences become more correlated, although there are examples of the opposite.²²

We present results on preferences generated from the following random utility model, adapted from Hitsch et al. (2010). Each agent ℓ has two characteristics, x_ℓ^A and x_ℓ^D . The utility of agent i for being matched with agent j is given by

$$u_i(j) = \beta x_j^A - \gamma(x_i^D - x_j^D)^2 + \varepsilon_{ij}, \quad (1)$$

where ε_{ij} is an idiosyncratic term for the pair (i, j) independently drawn from the standard logistic distribution. We use x_i^A as a vertical quality that is desirable for all agents, with $x_i^A \sim U[0, 1]$ drawn independently for each agent. We use x_i^D as a location, with $x_i^D \sim U[0, 1]$ drawn independently for each agent. The values of the coefficient γ determine agents' preferences to be close to/far from those of their partner. The same values of the coefficients β , γ are used for both men and women.

When $\beta = \gamma = 0$, preferences are drawn independently and uniformly at random. As β increases (keeping γ fixed), preferences become more correlated; and if $\beta \rightarrow \infty$, all men will have the same preferences over

²¹ We did find somewhat unnatural correlation structures that also lead to a large core in balanced markets, e.g., when generating preferences by drawing ε_{mw} independently for each pair (m, w) and setting $u_w(m) = \varepsilon_{mw} = -u_m(w)$. However, for all such structures we examined the core is small in unbalanced markets.

²² It is possible to construct specific examples such that an increase in correlation increases the size of the core. For example, if there are 41 men and 40 women with independently and identically distributed preferences, we can create a balanced market (and hence a large core) by making all women rank the same man at the bottom of their list.

women (and symmetrically). Taking $\gamma \neq 0$ generates alignment between $u_m(w)$ and $u_w(m)$.²³

Figure 4 shows the size of the core for different levels of imbalance and a range of coefficient values. For each market we simulate 2,000 realizations and report the average percentage of men with multiple stable partners.²⁴ We consider markets with 40 women and either 40, 41, 45, or 60 men, $\beta \in [0, 20]$ and $\gamma \in [-20, 20]$. Observe that correlation tends to reduce the size of the core (see nn. 21 and 22). The only markets that have a large core are balanced markets with low levels of correlation in preferences. We interpret these results as complementary to our theoretical results, giving additional suggestive evidence that general preference distributions are likely to generate a small core.

C. Robustness of the Short Side Advantage

Theorem 2 shows that if there are strictly fewer men than women and preferences are complete and uncorrelated, then men will have a large advantage: men will be matched with one of their top choices, while women will be matched with a husband who is not much better than a random man. But this result may not hold when preferences are correlated; for example, if all men have the same preferences, the top-ranked woman must be matched with her top choice (regardless of the number of men and women). We therefore conduct numerical experiments to investigate the extent of the men's advantage when there are fewer men and preferences are correlated.

Figure 5 reports the results for different markets generated using the utility model from Section IV.B. Each panel reports the men's average rank of wives under different values of β , γ for markets with 40 women and from 20 to 60 men. Each panel contains as a reference point the graph for $\beta = \gamma = 0$ (which replicates fig. 3). We plot the men's average rank only under the MOSM, as the men's average rank under the WOSM is almost identical (it differs only for balanced markets).

Panel A shows results with $\gamma = 0$ and different values of β as indicated. A higher value of β makes preferences of agents on the same side more correlated. The short side of the market retains an advantage for small values of β . When β is very large, all men have almost the same preferences over women, and therefore the men's average rank is $[\min(|\mathcal{M}|, |\mathcal{W}|) + 1]/2$.

Panel B shows results for varying values of γ , holding $\beta = 0$. If $\gamma > 0$, both agents prefer a close partner, generating alignment between the

²³ See Yariv and Lee (2014) for analysis of fully aligned random preferences.

²⁴ We also measured the size of the core using the difference/ratio between the men's rank of wives under MOSM and under WOSM. These measures produce very similar results and are therefore omitted.




FIG. 4.—The percentage of men with multiple stable partners for correlated preferences, generated as per equation (1): *A*, 40 men; *B*, 41 men; *C*, 45 men; *D*, 60 men. In each plot the *z*-axis (vertical and color) gives the percentage of agents with multiple stable partners, the *x*-axis is the value of the coefficient γ (ranging from -20 to 20), and the *y*-axis is the value of coefficient β (ranging from 0 to 20). The number of men is 40 , 41 , 45 , or 60 , as labeled for each plot. In all markets there are 40 women.





FIG. 4 (Continued)

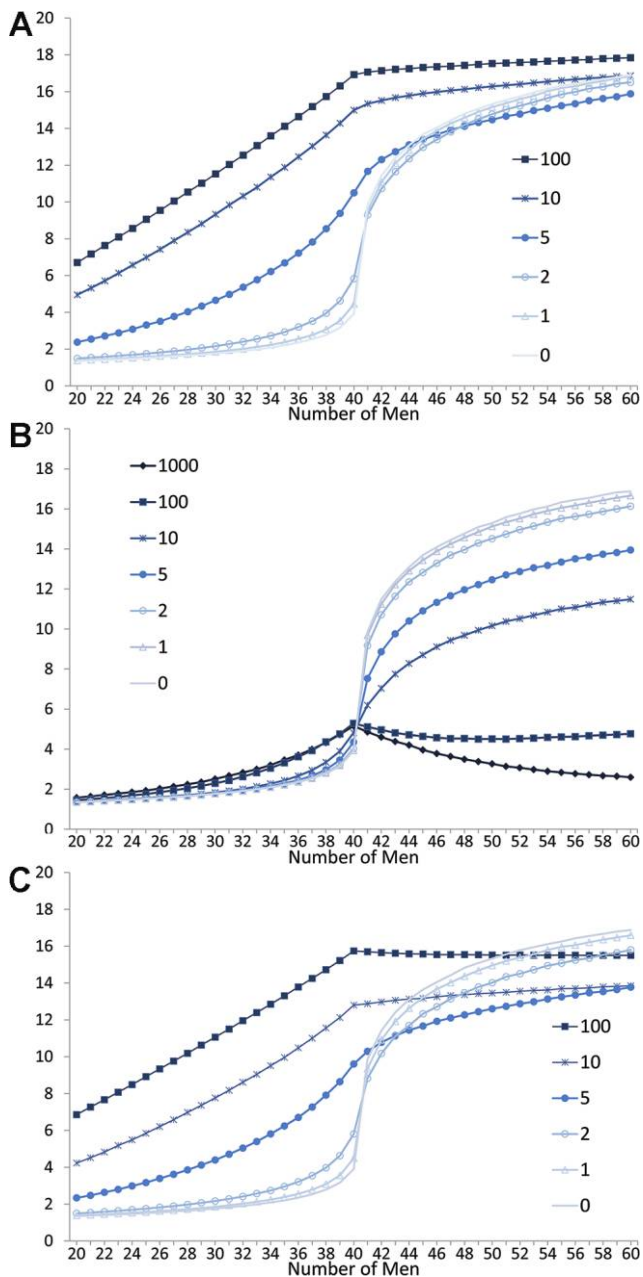


FIG. 5.—Men's average rank of wives under MOSM for correlated preferences, generated as per equation (1), with 40 women and a varying number of men. Panel A plots the average rank when $\gamma = 0$ and β ranges from 0 to 100, with different lines corresponding to different values of β . Similarly, panel B plots the average rank when $\beta = 0$ and γ ranges from 0 to 1,000. Panel C plots the average rank for $\beta = \gamma$.

preferences of men and women, as well as correlation in the preferences of men. As γ increases, the men's average rank of wives increases in markets with fewer than 40 men, but it decreases in markets with more than 40 men since matched men are likely to be favorably ranked by their wives and therefore likely to rank their wives favorably.

Panel C shows the results when taking $\beta = \gamma$, and each line is labeled by the common value for both coefficients. The advantage of the short side gradually decreases as correlation increases but is still evident even when there is a sizable correlation.

Put together, we find that for low levels of correlation, men tend to receive a lower average rank when there are fewer men. This advantage is continuously attenuated as agents' preferences become more correlated.

We further simulate markets in which agents find some potential partners to be unacceptable. For each man m and woman w we independently draw $u_m(w) \sim U[0, 1]$, where m 's utility of remaining single is \bar{u}^M . Thus w is unacceptable to m if $u_m(w) < \bar{u}^M$, and if $\bar{u}^M = 0.33$, there is a 33 percent chance that a man will find a potential wife to be unacceptable. Preferences for women were drawn analogously, with \bar{u}^W being the utility of remaining single. Figure 6 reports the men's average rank of wives under the MOSM for markets with 40 women and from 20 to 60 men and $(\bar{u}^M, \bar{u}^W) \in \{0, .33, .66\}$.²⁵ The shape of the marker indicates the value of \bar{u}^M , and the shade indicates the value of \bar{u}^W . The number of matched agents is at least 85 percent of $\min(|\mathcal{M}|, |\mathcal{W}|)$ in all reported markets. We find that only the selectivity of the long side matters: the lines merge according to shade when there are more women and according to shape when there are more men. The benefit of the short side is still apparent, but it is continuously attenuated as the long side becomes more selective.

D. Many-to-One Markets

We run computational experiments to investigate the effect of imbalance in many-to-one matching markets, in which colleges have responsive preferences (Roth 1985) and each college is small relative to the size of the market (see online app. C for further discussion). For each student we independently draw a complete preference list over colleges uniformly at random. For each college we independently draw a complete preference list over individual students uniformly at random.

Denote by \mathcal{S} the set of students, by \mathcal{C} the set of colleges, and by q the number of seats in each college. We say that a market is unbalanced if the number of students differs from the total number of seats, that is,

²⁵ The men's average rank under WOSM is omitted as it is almost identical in all unbalanced markets.

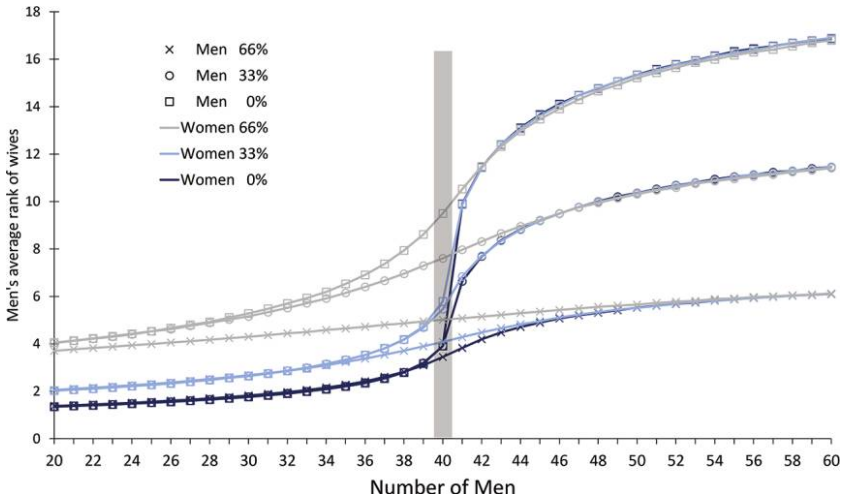


FIG. 6.—Men’s average rank of wives for various levels of imbalance and selectivity in random markets with 40 women and a varying number of men. Each of the 3×3 lines indicates a different level of selectivity for men and women, marked by shade and shape. The shape indicates the selectivity of men, giving the chance a man will find a potential wife to be unacceptable. The shade indicates the selectivity of women, giving the chance that a woman will find a potential husband to be unacceptable.

$|\mathcal{S}| \neq |\mathcal{C}| \times q$. For each market size we draw 1,000 realizations and compute the extreme stable matchings: the student-optimal stable matching (SOSM) and the college-optimal stable matching (COSM). The students’ average rank of their colleges (under both SOSM and COSM) is defined as before. The colleges’ average rank of students is computed by averaging the rank of students in all filled seats.

Tables 4, 5, and 6 report simulation results for markets with an overall number of seats of $q \times |\mathcal{C}| = 100, 200, 500, 1,000$, with a varying number of seats per college ($q = 2, 5$) and a varying number of students $|\mathcal{S}|$ (varying from $q|\mathcal{C}| - 10$ to $q|\mathcal{C}| + 10$). As in the one-to-one markets, we find that there is a big difference between the SOSM and COSM when the market is balanced, but in unbalanced markets the SOSM and COSM give almost the same average rank to both colleges and students. The percentage of students with multiple stable matches is large when the market is balanced, but it is small under imbalance.

If we increase q while holding $q \times |\mathcal{C}|$ and $|\mathcal{S}|$ fixed, we find that the core shrinks, and the college’s average rank increases. Note that an increase in q effectively makes preferences over seats more correlated, and our findings here resemble the effects of increased correlations reported in Sections IV.B and IV.C. The average rank of students is not

TABLE 4
STUDENTS' AVERAGE RANK OF COLLEGES IN THE SOSM AND COSM
FOR DIFFERENT MARKETS

$ \mathcal{C} \times q$	$ \mathcal{S} - \mathcal{C} \times q$									
	-10	-5	-2	-1	0	1	2	5	10	
100 (50×2):										
SOSM	1.8	2.0	2.4	2.7	3.1	10.4	12.0	14.4	16.3	
COSM	1.8	2.1	2.6	3.1	10.5	12.0	13.0	14.9	16.6	
200 (100×2):										
SOSM	2.1	2.4	2.8	3.1	3.6	17.8	20.4	24.7	28.1	
COSM	2.2	2.5	3.1	3.5	17.9	20.6	22.3	25.6	28.7	
500 (250×2):										
SOSM	2.6	2.9	3.4	3.6	4.1	38.0	43.3	51.6	59.1	
COSM	2.6	3.0	3.6	4.1	38.2	43.5	46.9	53.4	60.2	
1,000 (500×2):										
SOSM	3.0	3.3	3.8	4.0	4.5	69.1	78.0	91.7	103.7	
COSM	3.0	3.4	4.0	4.5	68.5	78.1	83.8	94.6	105.2	
100 (20×5):										
SOSM	1.3	1.5	1.6	1.7	1.9	4.5	5.1	6.0	6.8	
COSM	1.3	1.5	1.7	1.9	4.5	5.1	5.5	6.2	6.9	
200 (40×5):										
SOSM	1.5	1.7	1.8	2.0	2.1	7.5	8.5	10.2	11.6	
COSM	1.5	1.7	1.9	2.1	7.6	8.5	9.3	10.6	11.8	
500 (100×5):										
SOSM	1.8	1.9	1.1	2.3	2.4	15.6	17.7	21.0	23.9	
COSM	1.8	2.0	2.2	2.4	15.6	17.8	19.1	21.7	24.3	
1,000 (200×5):										
SOSM	1.9	2.1	2.3	2.5	2.7	27.7	31.4	36.8	41.6	
COSM	2.0	2.2	2.5	2.6	28.0	31.4	33.8	38.0	42.3	

NOTE.—A student's most preferred college has rank 1, and a larger rank indicates a less preferred college.

directly comparable across different values of q , since the length of the students' preference list changes.

V. Discussion

Competition has been at the heart of economic theory since Edgeworth (1881) and may have a stark effect as in the famous left glove, right glove example by Shapley and Shubik (1969). We show that there is a similar stark effect of competition in random matching markets, despite the heterogeneity of preferences and the lack of monetary transfers. This allows us to characterize stable matchings and to provide an explanation as to why small cores are empirically ubiquitous.

Under more general preference distributions the structure of stable matchings may be more involved. As an example, consider a tiered one-to-one market with 30 workers and 40 firms in two tiers: 20 firms are "top" and 20 firms are "middle." Preferences are drawn uniformly at random, except that every worker prefers any top firm over any middle firm.

TABLE 5
COLLEGES' AVERAGE RANK OF STUDENTS IN THE SOSM AND COSM FOR DIFFERENT MARKETS

$ C \times q$	$ \mathcal{S} - C \times q$								
	-10	-5	-2	-1	0	1	2	5	10
100 (50×2):									
SOSM	33.4	31.5	28.9	27.1	24.5	7.1	5.9	4.7	3.9
COSM	32.9	30.6	26.8	23.9	7.2	6.0	5.4	4.5	3.8
200 (100×2):									
SOSM	61.3	56.8	50.9	48.3	42.8	8.5	7.2	5.7	4.8
COSM	60.4	55.1	47.4	42.9	8.5	7.2	6.5	5.4	4.6
500 (250×2):									
SOSM	134.7	123.1	110.3	103.3	92.4	10.0	8.6	7.0	6.0
COSM	132.6	119.1	102.9	92.7	10.1	8.6	7.9	6.8	5.8
1,000 (500×2):									
SOSM	242.3	220.4	199.1	188.3	168.1	11.1	9.6	8.0	7.0
COSM	238.7	213.9	187	169.4	11.2	9.6	8.9	7.8	6.9
100 (20×5):									
SOSM	38.1	36.9	34.8	33.4	31.5	13.4	11.3	9.0	7.5
COSM	37.8	36.3	33.2	31.2	13.4	11.3	10.3	8.6	7.3
200 (40×5):									
SOSM	72.2	68.1	63.3	60.5	56.4	16.1	13.8	11.0	9.3
COSM	71.4	66.5	60.2	56.4	16.1	13.8	12.5	10.6	9.1
500 (100×5):									
SOSM	162.6	152.1	140.8	132.7	123.7	19.4	16.8	13.8	11.9
COSM	160.7	148.4	133.9	123.7	19.5	16.7	15.5	13.3	11.6
1,000 (200×5):									
SOSM	300.0	278.6	256.8	244.5	229.3	22.1	19.1	16.0	13.9
COSM	296.7	272.3	244.7	228.4	21.9	19.2	17.7	15.5	13.7

NOTE.—A college's most preferred student has rank 1, and a larger rank indicates a less preferred student.

Stable matchings can be decomposed to a stable matching between the top 20 firms and the 30 workers and a stable matching between the remaining 10 unmatched workers and the 20 middle firms. Applying our results to each part implies that the core is small, and stable matchings can roughly be described by the 20 top firms choosing their worker and

TABLE 6
PERCENTAGE OF STUDENTS WHO HAVE MULTIPLE STABLE MATCHES

$ C \times q$	$ \mathcal{S} - C \times q$								
	-10	-5	-2	-1	0	1	2	5	10
100 (50×2)	1.8	3.4	7.8	12.4	70.7	14.1	8.4	3.9	2.1
200 (100×2)	1.8	3.3	7.3	11.4	79.8	14.1	8.6	4.0	2.2
500 (250×2)	1.7	3.4	6.9	10.4	89.0	12.7	7.9	3.7	2.0
1,000 (500×2)	1.6	3.0	6.2	10.2	93.3	11.7	7.0	3.2	2.0
100 (20×5)	1.2	1.9	5.0	6.8	57.5	13.8	7.9	3.9	1.7
200 (40×5)	1.3	2.6	5.2	7.1	71.4	13.1	8.5	3.9	2.1
500 (100×5)	1.2	2.6	5.0	6.8	84.1	12.5	7.2	3.7	2.1
1,000 (200×5)	1.1	2.3	4.8	6.7	90.3	12.0	7.1	3.4	1.8

the remaining 10 workers choosing their firm from among the middle firms. A similar argument can be used to characterize stable matchings in any tiered market.

In online appendix C, we discuss how to extend our analysis to many-to-one markets between colleges and students, assuming that colleges have responsive preferences and that each college is small relative to the size of the market. The sharp effect of competition is also present in these markets: whenever there is even a slight difference between the number of seats and the number of students in the market, the core is small, and the short side will “choose.”

References

- Abdulkadiroglu, A., P. A. Pathak, and A. E. Roth. 2005. “The New York City High School Match.” *A.E.R.* 105:364–67.
- Abdulkadiroglu, A., P. A. Pathak, A. E. Roth, and T. Sönmez. 2006. “Changing the Boston School Choice Mechanism.” Working Paper no. 11965, NBER, Cambridge, MA.
- Abdulkadiroglu, A., and T. Sönmez. 1998. “Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems.” *Econometrica* 66 (3): 689–701.
- Abramitzky, R., A. Delavande, and L. Vasconcelos. 2011. “Marrying Up: The Role of Sex Ratio in Assortative Matching.” *American Econ. J.: Appl. Econ.* 3 (3): 124–57.
- Azevedo, E. D., and J. D. Leshno. 2016. “A Supply and Demand Framework for Two-Sided Matching Markets.” *J.P.E.* 124 (5): 1235–68.
- Banerjee, A., E. Duflo, M. Ghatak, and J. Lafortune. 2013. “Marry for What: Caste and Mate Selection in Modern India.” *American Econ. J.: Microeconomics* 5 (2): 33–72.
- Becker, G. S. 1973. “A Theory of Marriage: Part I.” *J.P.E.* 81:813–46.
- Che, Y.-K., and O. Tercieux. 2014. “Efficiency and Stability in Large Matching Markets.” Working paper, Soc. Sci. Res. Network. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2638009.
- Coles, P., and R. Shorrer. 2014. “Optimal Truncation in Matching Markets.” *Games and Econ. Behavior* 87:591–615.
- Crawford, V. P. 1991. “Comparative Statics in Matching Markets.” *J. Econ. Theory* 54 (1): 389–400.
- Demange, G., D. Gale, and M. Sotomayor. 1987. “A Further Note on the Stable Matching Problem.” *Discrete Appl. Math.* 16 (3): 217–22.
- Dubins, L. E., and D. A. Freedman. 1981. “Machiavelli and the Gale-Shapley Algorithm.” *American Math. Monthly* 88:485–94.
- Edgeworth, F. Y. 1881. *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*. London: C. K. Paul.
- Gale, D., and L. S. Shapley. 1962. “College Admissions and the Stability of Marriage.” *American Math. Monthly* 69:9–15.
- Hitsch, G. J., A. Hortaçsu, and D. Ariely. 2010. “Matching and Sorting in Online Dating.” *A.E.R.* 100:130–63.
- Holzman, R., and D. Samet. 2014. “Matching of Like Rank and the Size of the Core in the Marriage Problem.” *Games and Econ. Behavior* 88:277–85.
- Immorlica, N., and M. Mahdian. 2005. “Marriage, Honesty, and Stability.” In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 53–62. Philadelphia: Soc. Indus. and Appl. Math.

- Knuth, D. E. 1976. *Marriages stables et leurs relations avec d'autres problèmes combinatoires*. Montreal: Presses de l'Université de Montreal.
- Knuth, D. E., R. Motwani, and B. Pittel. 1990. "Stable Husbands." In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, 1–14. Philadelphia: Soc. Indus. and Appl. Math.
- Kojima, F., and P. A. Pathak. 2009. "Incentives and Stability in Large Two-Sided Matching Markets." *A.E.R.* 99:608–27.
- Lee, S. Forthcoming. "Incentive Compatibility of Large Centralized Matching Markets." *Rev. Econ. Studies*.
- McVitie, D. G., and L. B. Wilson. 1971. "The Stable Marriage Problem." *Communications ACM* 14 (7): 486–90.
- Pathak, P. A., and T. Sönmez. 2008. "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism." *A.E.R.* 98 (4): 1636–52.
- Pittel, B. 1989a. "The Average Number of Stable Matchings." *SIAM J. Discrete Math.* 2 (4): 530–49.
- . 1989b. "On Likely Solutions of a Stable Matching Problem." In *Proceedings of the 3rd ACM Symposium on Discrete Algorithms*, 10–15. Philadelphia: Soc. Indus. and Appl. Math.
- Rao, V. 1993. "The Rising Price of Husbands: A Hedonic Analysis of Dowry Increases in Rural India." *J.P.E.* 101:666–77.
- Roth, A. E. 1982. "The Economics of Matching: Stability and Incentives." *Math. Operations Res.* 7:617–28.
- . 1985. "The College Admissions Problem Is Not Equivalent to the Marriage Problem." *J. Econ. Theory* 36 (2): 277–88.
- . 1986. "On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets." *Econometrica* 54:425–27.
- . 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica* 70:1341–78.
- Roth, A. E., and E. Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *A.E.R.* 89:748–80.
- Roth, A. E., and X. Xing. 1994. "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *A.E.R.* 84:992–1044.
- Schwarz, M., and M. B. Yenmez. 2011. "Median Stable Matching for Markets with Wages." *J. Econ. Theory* 146 (2): 619–37.
- Shapley, L. S., and M. Shubik. 1969. "Pure Competition, Coalitional Power, and Fair Division." *Internat. Econ. Rev.* 10 (3): 337–62.
- . 1971. "The Assignment Game I: The Core." *Internat. J. Game Theory* 1 (1): 111–30.
- Yariv, L., and S. Lee. 2014. "On the Efficiency of Stable Matchings in Large Markets." Working paper, Soc. Sci. Res. Network. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2464401.