Large Matching Markets*

Jacob D. Leshno

Large matching models allow for a simpler description of matching markets. In physics, models of a handful of molecules are complicated, but it is simple to characterize the behavior of a gas consisting of great numbers of molecules moving in all directions. Likewise, by considering continuum models or the limit of increasingly large markets we are able tractably characterize the main features of the market, and abstract away from discreteness constraints or pathological situations.

The analytic tractability of large matching models allows us to ask a richer set of questions. To motivate the discussion, this chapter focuses on one such vein of questions: does the choice of the proposing side in DA matter? If it does, how much does it matter? These questions pose a long-standing puzzle. Gale and Shapley (1962) showed that men receive their optimal matching in the men-proposing DA, but their pessimal matching under the women-proposing DA. However, this theoretical result does little to inform us about the magnitude of the difference, and simple examples show that it is possible that the difference is very large and it is possible that there is no difference at all. Given a real market, should we expect this difference to be large or small? In this chapter, we will show that large matching models can allow us to answer this practical question.

Further, to reap the benefits that large markets offer, we need to develop a language. A lay person person may not find it natural to consider eating an infinitesimally larger apple. Indeed, most students are somewhat perplexed when they first learn about derivatives. On the other hand, a trained economist may be surprised that a grocer will not sell $\sqrt{2}$ apples, being so used to thinking of goods as infinitely divisible. By abstracting away from discreteness constraints, economists gained a great deal of intuition, employing the tools of calculus to analyze problems and gaining important intuition in the form of marginal effects (even if apples only come in discrete units).

The models we review in this chapter develop a language that provides us with a strong set of tools to better understand matching markets. Models of random matching markets allow us to give the more nuanced and useful answer of what is likely to happen, ignoring possible but highly unlikely pathological cases. The analysis will allow us to quantify the magnitude of effects and find which market features are important to consider. Continuum matching markets

*This chapter will appear in: *Online and Matching-Based Market Design*. Federico Echenique, Nicole Immorlica and Vijay V. Vazirani, Editors. Cambridge University Press.

allow us to use the tools of calculus and answer questions such as: how will the market respond to a change? Where should the planner direct investment to improve welfare? In addition, these models also provide a foundation for the econometric tools that allow estimation of the parameters that are necessary inputs to answer these questions.

The remainder of the chapter is organized as follows. Section 1 discusses models of one-to-one matching in which preferences are randomly drawn, focusing on using such models to understand the effect of the choice of the proposing side in DA. Section 2 discusses a many-to-one model in which a continuum of students is matched to a finite number of schools. Either section can be read independently of the other.

1 Random matching markets and the puzzle of the proposing side

When Roth and Peranson (Roth and Peranson, 1999) were tasked with redesigning the NRMP, both students and hospitals lobbied to have their side propose. As we have proved in previous chapters, all men are weakly better off under the men-proposing DA (MPDA) than under any other stable matching. A simple example with two men and and two women illustrates this:

Example 1 (Two-by-Two Economy)**.** Consider a market with two women w_1, w_2 and two men m_1, m_2 with preferences

$$\begin{aligned} m_1 &: w_1 \succ w_2 \qquad\qquad & w_1 &: m_2 \succ m_1 \\ m_2 &: w_2 \succ w_1 \qquad\qquad & w_2 &: m_1 \succ m_2 \end{aligned}$$

Under MPDA both men are matched to their top choice, and under woman-proposing DA (WPDA) both men are matched to their bottom choice.

However, the choice of the proposing side ended up being almost irrelevant in practice. Roth and Peranson simulated both the student and hospital-proposing algorithms on real past data and found that the choice of the proposing side made very little difference. Although some students did benefit from having their side propose, less than one in a thousand medical students was affected. And even those who were affected mostly received a similarly ranked match.

This finding posed a theoretical challenge. While theory correctly predicted that men did weakly benefit from having their side propose, it failed to inform us that the choice of the proposing side is immaterial. In Example 1 the choice of the proposing side affects all agents, and determines whether an agent receives their top choice or least preferred choice. What explains this difference? Does this imply there is something special about the preferences submitted to the NRMP? What should we expect in other markets?

1.1 Saying the market is "large" is not enough

One conjecture (motivated by the title of this chapter) might be that the NRMP matches over 20,000 doctors every year, while Example 1 includes only two men and two women. What if we restrict attention to markets that have many agents?

Example 2 (Large Market with Opposite Preferences)**.** Consider a market with n women $\{w_1, \ldots, w_n\}$ and n men $\{m_1, \ldots, m_n\}$ with preferences

$$\begin{aligned} m_i :& w_i \succ w_{i+1} \succ \cdots \succ w_n \succ w_1 \succ \cdots \succ w_{i-1} \\ w_i :& m_i \prec m_{i+1} \prec \cdots \prec m_n \prec m_1 \prec \cdots \prec m_{i-1} \end{aligned}$$

Under MPDA all men are matched to their top choice, and under WPDA all men are matched to their bottom choice.

Example 2 shows that the choice of the proposing side can have a large effect in a market of any size. The following example goes further, showing that absent of some additional structure the number of agents in the market is, in a sense, irrelevant.

Example 3 (Island Replication)**.** Consider a market with $2n$ women $\{w_1^k, w_2^k\}_{k \leq n}$ and $2n$ men $\{m_1^k, m_2^k\}_{k \leq n}$ with preferences

$$\begin{aligned} m_1^k &: w_1^k \succ w_2^k \succ \emptyset & \qquad w_1^k &: m_2^k \succ m_1^k \succ \emptyset \\ m_2^k &: w_2^k \succ w_1^k \succ \emptyset & \qquad w_2^k &: m_1^k \succ m_2^k \succ \emptyset \end{aligned}$$

This market is equivalent to n copies of the market from Example 1, as men m_1^k, m_2^k can only match to women w_1^k, w_2^k.

1.2 Random matching markets

From these examples and the NRMP simulations we learned that the choice of the proposing side can matter a lot for some economies, but makes little difference for others.

Therefore, our next step is to explore a range of economies. A natural starting point is to explore economies drawn uniformly at random.[1] A *random matching market* is composed of a set of men $\mathcal{M}$ and a set of women $\mathcal{W}$ in which preferences are 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.

We introduce two metrics to asses the benefit to the proposing side.

[1] For tractability, the model assumes preferences are uncorrelated. We will allow arbitrary correlations in Section 2

Definition 4. Given a matching μ, the *men's average rank of wives* is given by

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

where $\bar{\mathcal{M}}$ is the set of men who are unmatched under μ.[2]

Similarly, the *women's average rank of husbands* is given by

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

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

The men-proposing DA (MPDA) produces the Men Optimal Stable Matching (MOSM), and the woman-proposing DA (WPDA) produces the Women Optimal Stable Matching (WOSM). Recall that these are the extreme points in the lattice of stable matchings. In particular, each man is matched to his most preferred spouse under the MOSM and to his least preferred spouse under the WOSM (out of all the women he is matched to in some stable matching). If a man is matched to the same woman under both the MOSM and WOSM, then he has a unique stable partner. Pittel (1989) and Knuth et al. (1990) characterize the MOSM and WOSM for a random matching market with $|\mathcal{W}| = |\mathcal{M}| = n$ as n grows large.

Theorem 5 (Pittel (1989))**.** *In a random matching market with n men and n women, the fraction of agents who have multiple stable partners converges to 1 as $n \to \infty$. Furthermore,*

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

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

In words, there is a substantial benefit to the proposing side in a typical random matching market with n men and n women. The benefits are widespread, most men strictly prefer their spouse in the MOSM over their spouse in the WOSM. The benefits are also large, under the MOSM a men is matched to his $\log(n)$ most preferred wife (on average), while under the WOSM a men is matched to his $n/\log(n)$ most preferred wife (on average). For example, if we randomly draw a market with $n = 1,000$ men and women, then each men expects to be matched to his $\log(1000) \approx 7$ most preferred wife under the MOSM, but to his $1000/\log(1000) \approx 145$ most preferred wife under the WOSM. Since the core of the economy is the set of stable matchings, we can equivalently say that the core is large.

[2] $\bar{\mathcal{M}}$ does not depend on μ because an agent who is unmatched in some stable matching is unmatched in any stable matching.

While we may not necessarily be interested in the behavior of large random matching markets, it is useful to consider $n \to \infty$ for two reasons. First, it is much easier to prove and characterize the asymptotic behavior. Second, by taking limits we can gain an intuition for the "first-order effect" in the market, abstracting away from secondary issues. As a result, the proof to can be more elegant and provide more intuition.

1.3 Random matching markets with short preference lists

Theorem 5 poses a puzzle. Roth and Peranson (1999) had data on multiple years and observed a dramatically smaller effect in the NRMP data. What makes the NRMP so different from these random matching markets? Roth and Peranson (1999) conjectured that the difference was due to a different distribution of preference lists. Most students' preference lists include a few dozen hospitals at most, and omit the vast majority of hospitals.

A *random matching market with short preference lists* is composed of a set of men $\mathcal{M}$ and a set of women $\mathcal{W}$ in which women have arbitrary complete preference lists, and the preferences of each men is generated by drawing k women uniformly and independently. The size of the market is $n = |\mathcal{W}| = |\mathcal{M}|$. We consider the case where $k \ll n$, that is, markets where the number of participants is large relative to the length of randomly drawn preference lists.

Theorem 6 (Immorlica and Mahdian (2005))**.** *In a random matching market with short preference lists with n men and n women, for any fixed length of preference list k, the expected fraction of agents who have multiple stable partners converges to 0 as $n \to \infty$.*

In other words, large economies in which one side has constant length randomly drawn preference lists have a small set of stable matchings, in the sense that most agents have the same partner under all stable matchings. This type of result is often referred to in the literature as a *core convergence* result.

Note that this result only holds for *randomly* drawn short preference lists. Example 3 shows a large economy with short preference list that has many stable matchings.

To prove the result, it is sufficient to determine the probability that a given man m_0 is matched to the same woman under the MOSM and the WOSM. While we can easily calculate both using DA with either side proposing, it is not a priory clear how we can tractably obtain the probability that both assign m_0 to the same partner. Therefore, the proof leverages the structure of stable matchings, and it uses *rejection chains* and the McVitie and Wilson (1971) algorithm to determine whether m_0 has multiple stable partners.

Consider man m_0 in a randomly drawn market such that the MOSM μ matches m_0 to $\mu(m_0) = w_0 \in \mathcal{W}$. If there is another stable matching μ' such that $\mu'(m_0) = w' \neq w_0$, then the matching μ' remains stable even if we change the preferences of woman w_0 so that she finds m_0 unacceptable. This change will make woman w_0 reject man m_0 if he proposes to her in MPDA. Because the order in which men propose in men-proposing DA does not affect the resulting match, we can first run MPDA until we reach the MOSM, and only then reject man m_0 and track the following steps of the algorithm.

The sequence of proposals following the rejection of m_0 by $\mu(m_0) = w_0$ is called a *rejection chain.* The algorithm continues with a proposal from man m_0 to his next most preferred wife. If m_0 proposed to a woman who is already matched to a more preferred husband, m_0's proposal will be rejected, and m_0 will keep proposing. If m_0 proposed to a woman who prefers m_0 over her current match m_1, m_0's proposal will be temporarily accepted, and m_1 will start proposing. This generates a chain of proposals, where at any time at most one man is making proposal, and a new man starts proposing when displaced by the previous man.

The chain can terminate in one of three ways: (i) a proposing man exhausted his preference list, (ii) a man proposed to an unmatched woman $w \neq w_0$ and the proposal was accepted, or (iii) a man proposed to the unmatched woman w_0 and the proposal was accepted. If (i) or (ii) happened, the algorithm will output a matching in which the set of agents that are matched is different from the set of agents matched under μ, and therefore the outputted matching cannot be stable under for the original economy (since by the rural hospital theorem any stable matching leaves the same agents unmatched). It is a good exercise to show this implies that there are no stable matchings in which man m_0 is matched to a woman w such that $\mu(m_0) \succ_m w$. If (iii) happened, then the algorithm will output a matching μ' in which the set of agents that are matched is the same as the set of agents matched under μ, and it is a good exercise to show that this is a new stable matching in which m_0 is matched to a different wife such that $\mu'(m_0) \prec_{m_0} \mu(m_0)$.

It is left to determine the probability that a rejection chain terminates in (iii), and not in (i) or (ii). To do so, we first calculate the expected number of women that are unmatched in any stable matching. Let X be the (random) number of women that do not appear in any man's preference list (i.e., all men find unacceptable). Clearly, X is a lower bound for the number of unmatched women.

By the previous arguments, the probability that man m_0 has an additional stable partner is at most the probability that the rejection chain terminates with a proposal to woman w_0 (his original partner) and not with a proposal to an unmatched woman. By the principle of deferred decision, we can dynamically draw the men's preferences only as the algorithm asks the men to propose. The next proposal by a men will be drawn uniformly at random among all women the man did not previously propose to, and therefore is at least as likely to go to an unmatched woman (who was not previously proposed to) as it is to go to w_0. Thus, the probability that w_0 receives a proposal before all of the unmatched women is at most $1/(1 + X)$.

A probabilistic calculation (using a variant of the occupancy problem) shows that for any fixed length of preferences list k and market size n, we have that

$$E\left[\frac{1}{1+X}\right] \leq \frac{e^{k+1} + k^2}{n} \ .$$

By linearity of expectation, this bounds the expected fraction of men who have multiple stable partners, completing the argument.

1.4 Unbalanced Random Matching Markets

While Theorem 6 provides us with an answer to the puzzle presented by the NRMP data, it leaves some open questions as well. The contrast with Theorem 5 shows that the assumption that preference lists are short is essential, and preference lists need to be so short (and random) that many agents remain unmatched. However, this does not match the economies we see in the NRMP. Preferences lists may be short, but agents choose a preference list that is likely to get them matched.

Another remaining challenge is giving a prediction for general markets that do not have the particular features of the model, namely markets that are not necessarily large or with short preference list. Should we expect a substantial benefit to the proposing side in most matching markets?

As it turns out, in any documented matching market the benefit to the proposing side is negligible, and the choice of the proposing side in DA is immaterial. In other words, the data suggests that the small core documented in the NRMP is the typical case. If a matching market asks you to predict whether the choice of the proposing side will make a difference, the empirical data suggests to answer that the choice will have a negligible effect.

Ashlagi et al. (2017) provide an explanation for why general matching markets will have a small core. A key observation is that Theorem 5 considered markets in which the number of women is exactly equal to the number of men. To gain intuition, consider standard markets with payments. In standard markets an exact balance between buyers and sellers can generate a large core, but it is eliminated by any imbalance. For example, consider a market with 100 buyers with unit demand and willing to pay 1 for an item, and 100 sellers each offering one unit of the item with reservation value of 0. The core of this market is large, and any price between 0 and 1 will generate a core allocation. However, if we add even a single seller there is a unique clearing price of 0, because one seller must be unmatched and willing to sell for any positive price. Therefore, we consider what happens in *unbalanced matching markets* in which there is a different number of men and women.

Theorem 7 (Ashlagi et al. (2017))**.** *In a random matching market with n men and $n+1$ women the fraction of agents who have multiple stable partners converges to 0 as $n \to \infty$. Moreover, for any $\varepsilon > 0$ and sufficiently large n, in every stable matching μ*

$$R_{\text{MEN}}(\mu) \le (1+\varepsilon)\log n\,,$$
$$R_{\text{WOMEN}}(\mu) \ge \frac{n}{(1+\varepsilon)\log n}\,.$$

In other words, even the slightest imbalance eliminates the benefit to the proposing side and, while there may be multiple stable matchings, the choice of the proposing side makes little difference. That is, balanced random matching markets (which have exactly the same number of men and women) are special and atypical (see also Theorem 12). This suggests that in practice we should see a small core in any matching market, as exactly balanced markets are unlikely.[3]

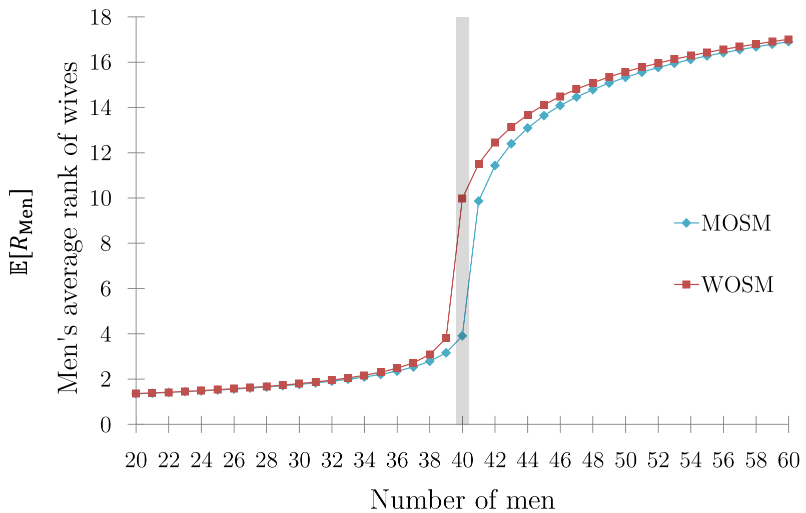

Figure 1: Average rank of men given different number of men and $|\mathcal{W}| = 40$ women.

Theorem 7 also shows the benefit to the short side, regardless of the proposing side. If there are n men and n women, the proposing side receives their $\log(n)$ most preferred partner, and the other side receives their $n/\log(n)$ most preferred partner. If we add additional women to the market, men will be better off and the existing women will be worse off. Therefore, we should expect the men to receive their $\log(n)$ partner if the men propose and we add another woman. More surprising is that in an unbalanced market, men receive their $\log(n)$ most preferred partner in every stable matching. For example, in a market with 1000 men and 1001 women, a man is matched to his 7th most preferred partner (on average), and a woman is matched to her 145th most preferred partner in *any* stable matching.

[3]Balanced random matching markets impose a particular distribution of preferences and assume that agents find every possible partner acceptable. In practice, we might expect agents' preferences to be more highly correlated, and that some agents will generally be found unacceptable. Since both tend to shrink the set of stable matchings, we should expect a small core.

When the imbalance is greater, the benefit to the short side is greater, and the choice of the proposing side is insignificant.

Theorem 8 (Ashlagi et al. (2017))**.** *Fix $\lambda > 0$. In a random matching market with n men and $(1+\lambda)n$ women the fraction of agents who have multiple stable partners converges to 0 as $n \to \infty$. Moreover, for sufficiently large n, in every stable matching μ the men's average rank of wives is at most κ, and the women's average rank of husbands is at least $n/(1+\kappa)$ where $\kappa = (1+\lambda)\log(1+1/\lambda)$ is a constant that depends only on λ.*

In other words, substantial imbalance leads to an allocation that is very beneficial to the short side. For example, if $\lambda = 0.05$, that is, a market with 5% extra women, then men will be matched, on average, with roughly their 3rd most preferred woman. The women's average rank of husbands (even if women propose) is only a factor of $(1+\kappa)/2 = 2.1$ better than being matched with a random man.

The proof of Theorems 7 and 8 again relies on analysis of rejection chains. As in the proof of Theorem 6, it would have been sufficient to evaluate the result of MPDA and WPDA. MPDA is tractable, since each man proposes to a small number of women (at most $\log(n)$), and the next woman any man will propose to is, approximately, a woman drawn uniformly at random. But WPDA is not as tractable – since a woman make many proposals, the next proposal of a woman substantially depends on the set of men she already proposed to. Therefore, instead of evaluating WPDA we evaluate MPDA and analyze rejection chains to show that MPDA and WPDA will be very close.

The proof of Theorem 8 is similar to that of Theorem 6. Consider a man m_0 that is matched under MPDA to woman w_0. If there are n men and $(1+\lambda)n$ women, λn women will remain unmatched. Since there is a large fraction of unmatched women, the rejection chain that stars with w_0 rejecting m_0 is much more likely to reach one of the λn unmatched woman than w_0.

The proof of Theorem 7 requires a more delicate analysis of rejection chains. If there are n men and $n+1$ women in the market, there will only be a single unmatched woman in the market, and a proposal is almost equally likely to go to w_0 as it is likely to go to an unmatched woman. To prove the result, we need to consider all rejection chains jointly. To do so tractably, the proof follows the calculation of the WOSM through a sequence of rejection chains and keeps tracks of the set of women S that already obtained their WOSM partner. The key observation that simplifies the analysis is that if the run of a rejection chain includes a woman from S accepting a proposal, it must be that the rejection chain terminates without finding a new stable matching (that is, in (i) or (ii)); otherwise the chain also finds a more preferred partner for a woman in S and we have a contradiction. Using this observation, the proof follows from two steps. First, S grows quickly, because once we find a rejection chain that ends in an unmatched woman, all the women that were part of the chain must be matched to their WOSM partner. Second, once S is large, it is likely that any proposal will reach a woman in S and we can terminate the chain knowing that it cannot lead to a new stable matching.

1.5 Small random markets

All the theorems we consider in this section are stated for $n \to \infty$, which greatly simplifies the proofs. Obtaining results for finite random markets is considerably more challenging, but simulations can help us evaluate when we should expected the asymptotic results to be relevant. Simulations of unbalanced markets show that the asymptotic characterization accurately describes markets even with 40 agents (see Figure 1.4).

2 Continuum Matching Markets

The previous section used a sequence of increasingly large matching markets. A natural question is whether we directly capture the limit object, as this limit object may allow us to simplify the analysis. As it turns out, the answer is positive for a different class of matching problems: many-to-one matching markets where one side is large.[4] For example, colleges and schools match to a large number of students. The continuum formulation gives a tractable characterization of stable matchings in the form of cutoffs, an admission threshold for each college. This means that we can describe stable matchings by a low-dimensional object, and find stable matchings by solving simple demand equations.

2.1 Formal Model

Consider a matching market where there is a finite set of colleges and each college can match with a large number of students. We start by describing the limit economy in which each college is matched with a continuum of students (as introduced in Azevedo and Leshno (2016)). The description of such an economy needs to capture (i) the preferences of colleges over students and their capacity limits, and (ii) the preferences of students over colleges. The standard approach to describe (i) is to assume responsive preferences (Roth, 1985) and represent the college's preferences by a preference ordering over all students and a capacity. Because the set of students is infinite, we instead represent the preferences of a college by assigning a *score* to each student, where a student with a higher score is more preferred by the college. This allows us to describe how colleges rank a given student by a vector of scores, one score for each college.

Formally, a continuum matching market is described as follows. There is a finite set of colleges $\mathcal{C} = \{1, \ldots, C\}$ that is to be matched to a continuum of students. A student is described by a type $i = \left(\succ^i, r^i\right)$. The student preferences over colleges are given by $\succ^i$. The vector $r^i \in [0,1]^{\mathcal{C}}$ describes the colleges' rankings of the students, where college c prefers student i over student j if $r^i_c > r^j_c$. We refer to r^i_c as the rank or the score of student i at college c. Given that scores represent ordinal information, we can rescale the scores so that a student i's score at college c is the percentile of i in c's ranking over students.[5] We use $\Theta = \mathcal{L}(\mathcal{C}) \times [0,1]^{\mathcal{C}}$ to denote the set of all possible student types, where $\mathcal{L}(\mathcal{C})$ is the set of all strict orderings over $\mathcal{C}$. Note that different colleges can have different rankings over students. To simplify notation, we assume in this chapter that all students and colleges are acceptable.[6]

A continuum market is given by $E = [\mathcal{C}, \eta, q]$ where η is a measure over the

[4]One can also define a limit economy for large one-to-one economies in which the number of agents on both sides grows large, but in such limit markets each agent has a preference over the infinite agents on the other side of the the market (see Section 3.2).

[5]That is, without loss of generality we can take $r^i_c = \eta\left(\{j \in \Theta : r^i_c > r^j_c\}\right)/\eta(\Theta)$.

[6]This is without loss of generality, as we can introduce an additional college and additional student that represent being unassigned. (Recall that being unassigned is equivalent to matching with a partner that will always be willing to form a blocking pair).

set of students Θ and $q = (q_1, \dots, q_C)$ is the capacity of each college. We can use this representation to describe a standard discrete many-to-one matching market with a finite number of students by taking η to be a measure with finitely many atoms, with each atom corresponding to a single discrete student. While we will later consider discrete economies, this representation is particularly useful in allowing us to think of economies in which each student is infinitesimal.

Assumption 9. Students and colleges have strict preferences. That is, every college's indifference curves are of η-measure 0. That is, for any $x \in [0,1]$ we have that $\eta\left(\{j \in \Theta : r_c^j = x\}\right) = 0$.

Assumption 9 ensures that $\eta(i) = 0$ for any $i \in \Theta$. The following example illustrates a continuum economy for school choice.

Example 10 (The Line City)**.** Consider a matching market two schools $\mathcal{C} = \{1,2\}$ with capacities q_1, q_2. Students live on the segment $[0,1]$ connecting school 1 and school 2, and are uniformly distributed on this line. A student in location $x \in [0,1]$ prefers school 1 over school 2 with probability $pr(x) = 1-x$ and prefers school 2 with probability $1 - pr(x) = x$. Both schools give higher priority to students who live closer to the school. That is, the score of a student in location x is $(1-x, x)$.

This economy is captured by $\mathcal{C} = \{1,2\}$ and the measure η given by

$$\begin{aligned}\eta\left(\{i : 2 \succ^i 1, r_2^i \in [a,b], r_1^i = 1 - r_2^i\}\right) &= \int_a^b (1 - pr(x))dx \\ &= \frac{b^2 - a^2}{2}\end{aligned}$$

for $a < b \in [0,1]$, and $\eta\left(\{i : 1 - r_1^i \neq r_2^i\}\right) = 0$.

As the example shows, a continuum market allows arbitrary correlations between the preferences of students over colleges and the preferences of colleges over students. This model can also describe markets in which college priorities are distributed independently of student preferences (e.g., in school choice Abdulkadiroglu et al. (2011)) as a particularly tractable case.

A matching for a continuum economy is given by a mapping $\mu : \Theta \to \mathcal{C} \cup \{\emptyset\}$. For each student i is assigned to a college $\mu(i)$, with $\mu(i) = \emptyset$ denoting the student is unassigned. With slight abuse of notation we write $\mu(c)$ for the set $\mu^{-1}(c)$ of students assigned to c. To avoid measure theoretic issues, we also require that for each $c \in \mathcal{C} \cup \{\emptyset\}$ we have that the set $\mu(c)$ is measurable and that the set $\{i : \mu(i) \prec^i c\}$ is open.

As in the discrete model, a student-college pair (i,c) blocks a matching μ if they can both benefit from matching to each other. A formal definition for the continuum setting is (i,c) blocks if $c \succ^i \mu(i)$ and either (i) c did not fill its capacity, that is $\eta(\mu(c)) < q_c$, or (ii) there exists $j \in \mu(c)$ such that $r_c^i > r_c^j$.[7] A matching μ is stable if it is not blocked by any student-college pair and $\eta(\mu(c)) \le q_c$ for all $c \in \mathcal{C}$.

[7] Because we are working with a continuous mass of students, a college can technically always match to an additional infinitesimal student. Such a definition will lead to a paradox,

2.2 Cutoffs and demand

To obtain a tractable representation of stable matchings, we introduce some additional notation. A cutoff for college c is a minimal score $p_c \in [0,1]$ required for admission to c. Given a vector of cutoffs $p = (p_c)_{c\in\mathcal{C}}$, we define the budget set of student i to be the set of colleges student i can be admitted to, denoted by $B^i(p) = \{c \in \mathcal{C} : r_c^i \geq p_c\}$. The demand of student i given cutoffs p is i's most preferred college from his budget set

$$D^i(p) = \max_{\succ^i} B^i(p)$$

where $D^i(p) = \emptyset$ if $B^i(p) = \emptyset$. Finally, aggregate demand for college c given cutoffs p is given by

$$D_c(p|\eta) = \eta\left(\{i : D^i(p) = c\}\right) \ .$$

We write $D_c(p)$ when η is clear from context, and denote overall demand by $D(p) = (D_c(p))_{c\in\mathcal{C}}$.

Consider a decentralized admission mechanism which posts cutoffs p and assigns each student to their most preferred college out of their budget set, that is $\mu(i) = D^i(p)$. The resulting matching μ will be envy free; if student i prefers college c over her assigned college $\mu(i)$ it must be that $c \notin B^i(p)$, and therefore for any student $j \in \mu(c)$ we have that $r_c^j > p_c > r_c^i$. The resulting matching may not be stable because a college c can be assigned more students than its capacity, or because a college is assigned strictly less students than its capacity and some student would like to be assigned to that college.

We say that cutoffs p are market clearing cutoffs if for all $c \in \mathcal{C}$ we have that $D(p) \leq q$, and if $D_c(p) < q_c$ then $p_c = 0$. That is, no college is assigned more students than its capacity, and a college that did not fill its capacity has a cutoff equal to 0.[8]

By the previous arguments, market clearing cutoffs p induce a stable matching given by $\mu(i) = D^i(p)$. The following lemma from Azevedo and Leshno (2016) shows the reverse is also true, and stable matchings are equivalent to market clearing cutoffs.

Lemma 11. *Stable matchings are equivalent to market clearing cutoffs. That is, every stable matching μ corresponds to market clearing cutoffs p defined by $p_c = \inf\{r_c^i : i \in \mu(c)\}$. Every market clearing cutoff p corresponds to a stable matching μ defined by $\mu(i) = D^i(p)$.*

The correspondence also preserves the lattice structure. If p, p' are both market clearing cutoffs, we can define cutoffs $p^+ = \max(p, p')$ and $p^- = \min(p, p')$

as the college cannot sequentially add all blocking students without violating its capacity constraint. The definition above gives a formulation of the continuum model that provides that desired behavior.

[8] If $\sum_c q_c < \eta(\Theta)$, that is, there are more students than college seats, then it is sufficient to require that $D(p) = q$. This is because in such a case it is impossible to have market clearing cutoffs such that $D_c(p) < q_c$ for some c.

(the min and max are taken coordinate by coordinate). The cutoffs p^+, p^- are also market clearing cutoffs. Moreover, if μ, μ' are the stable matchings corresponding to p, p' then p^+ corresponds to $\mu \wedge \mu'$ which matches each student i to $\min_{\succ^i}\{\mu(i), \mu'(i)\}$, and p^- corresponds to $\mu \vee \mu'$ which matches each student i to $\max_{\succ^i}\{\mu(i), \mu'(i)\}$.

The equivalence of stable matchings and market clearing cutoffs is more general, and there are many related characterizations in the literature (e.g., Balinski and Sönmez (1999), Adachi (2003), Abdulkadiroglu et al. (2011), Fleiner and Jankó (2014)). It is particularly useful in continuum economies where the cutoff characterization offers a tractable way to solve for stable matchings.

2.3 Calculating the stable matching

Let us revisit the market from Example 10. Consider cutoffs p_1, p_2. College 1 is in the budget set of a student if $r_1^i \geq p_1$, who are the students whose location is $x \in [0, 1 - p_1]$. College 2 is in the budget set of students whose location is $x \in [p_2, 1]$. Students remain unassigned if the student's location is in $[1 - p_1, p_2]$ and the student's budget set is empty. If $q_1 + q_2 < 1$ some students must remain unassigned, and therefore $1 - p_1 < p_2$, all students are assigned to the single college in their budget set, and $1 - p_1 = q_1, 1 - p_2 = q_2$. If $q_1 + q_2 \geq 1$, all students must be assigned and we have that $1 - p_1 \geq p_2$. In which case, demand for college 1 is all the students whose location is in $[0, p_2]$ (which have the budget set $\{1\}$) plus the students whose location is in $[p_2, 1 - p_1]$ (which have the budget set $\{1\}$) who prefer college 1:

$$D_1(p) = p_2 + \int_{p_2}^{1-p_1} pr(x)dx = \frac{1 - p_1^2 + p_2^2}{2}$$

Likewise, the demand for college 2 is $D_2(p) = (1 + p_1^2 - p_2^2)/2$.

Suppose $q_1 + q_2 > 1$, and that, without loss, $q_1 > 1/2$. Because it is impossible for both colleges to fill their capacity, one of the cutoffs must be equal to 0, and it is easy to verify that it must be that $p_1 = 0$. If $q_2 < 1/2$ there is a unique solution to $D_2(p) = q_2$ which is $p_2 = \sqrt{1 - 2q_2}$. If $q_2 \geq 1/2$ we have a unique solution $p_2 = p_1 = 0$.

If $q_1 + q_2 = 1$ there are multiple market clearing cutoffs and multiple stable matchings. For example, if $q_1 = q_2 = 1/2$ then any $p_1 = p_2 \in [0, 1]$ are market clearing cutoffs.

2.4 Generic uniqueness of stable matching

The calculation of market clearing cutoffs shows that it is possible for a continuum economy to have multiple market clearing cutoffs, and correspondingly, multiple stable matchings. However, it suggests that multiplicity of stable matching may not be typical, as it only happens in a knife edge event.[9] This

[9]Namely, a very particular set of parameters in which the amount of college seats is exactly equal to the amount of students.

aligns with the intuition from random matching markets – multiplicity of stable matchings arises only when the market is balanced. The continuum model allows us to formalize that this is a knife edge situation, and generically a matching market has a unique stable matching.

The following theorem requires a technical condition, namely that η is regular. The distribution of student types η is regular if the image under $D(\cdot|\eta)$ of the closure of the set

$$\{P \in [0,1]^C : D(\cdot|\eta) \text{ is not continuously differentiable at } P\}$$

has Lebesgue measure 0. This technical condition is satisfied if, for example, $D(\cdot|\eta)$ is continuously differentiable or if η admits a continuous density.

Theorem 12 (Azevedo and Leshno (2016))**.** *Suppose η is regular. Then for almost any vector of capacities q the market (η, q) has a unique stable matching.*

In other words, we can ensure a unique stable matching by slightly perturbing college capacities q.

2.5 Calculating and optimizing for welfare

Let us revisit the market from Example 10 again. The market clearing conditions allowed us to solve for the stable matching for any q_1, q_2. This allows us to treat the capacities q_1, q_2 as decision variables. For example, it is natural to ask what are the optimal capacities q_1, q_2 that maximize the sum of student utilities subject to some constraints or costs.

To define welfare, we need to specify a cardinal utility. A simple specification that is consistent with the distribution of ordinal preferences in Example 10 is that a student in location $x \in [0,1]$ has utility $1 - x/2$ if assigned to school 1 and utility $1 - (1-x)/2 + \varepsilon$ if assigned to school 2, where $\varepsilon \sim U[-1/2, 1/2]$ is an independent random taste shock. Unassigned students receive utility of 0.

We can obtain tractable expressions for welfare by conditioning on a student's location and budget set. The expected utility of a student in location x that has the budget set $\{1\}, \{2\}$ is $u_{\{1\}}(x) = 1 - x/2$ and $u_{\{2\}}(x) = x/2$, respectively. The expected utility of student in location x that has the budget set $\{1, 2\}$ is given by

$$\begin{aligned}
u_{\{1,2\}}(x) &= \int_{-\frac{1}{2}}^{\frac{1}{2}} \max\left[1 - x/2, x/2 + \varepsilon\right] d\varepsilon \\
&= \int_{-\frac{1}{2}}^{x} \max\left[1 - x/2\right] d\varepsilon + \int_{x}^{\frac{1}{2}} \max\left[x/2 + \varepsilon\right] d\varepsilon \\
&= \frac{1}{2}\left(x^2 - x + 2\right)
\end{aligned}$$

Combining these expressions with the previously derived expressions for the cutoffs, we obtain the following expressions for welfare. If $q_1 + q_2 < 1$ then

$p_1 = 1 - q_1, p_2 = 1 - q_2$ and welfare is

$$\int_0^{q_1} u_{\{1\}} dx + \int_{1-q2}^{1} u_{\{2\}} dx = q_1 - q_1^2/4 + q_2 - q_2^2/4 \; .$$

If $q_1 + q_2 > 1$ with $q_1 > 1/2 > q_2$ we have that $p_1 = 0$ and $p_2 = \sqrt{1 - 2q_2}$ and welfare is

$$\int_0^{p_2} u_{\{1\}} dx + \int_{p_2}^{1} u_{\{1,2\}}\, dx = \frac{11}{12} - \frac{1}{6}(1 - 2q_2)^{3/2} \; .$$

2.6 Random sampling and relation to discrete economies

One important interpretation of the continuum matching model is that the measure η represents the distribution from which finite economies are sampled. Understanding properties of finitely drawn economies is crucial for empirical work (e.g., Abdulkadiroğlu et al. (2017)). This interpretation is also useful for understanding the continuum model and its relation to finite matching model. The results below show that indeed we can think of continuum models as an approximation for finite models in which each student is small relative to the capacity of schools.

As an illustration, consider Israeli college admission. As in many countries (the U.S.A. being an exception), college admission is determined by the student's grade in standardised national exams. Each program in each university ranks students according to a score that is calculated based on the subjects the student took and the grades the student received. Different programs use different weights and can have different rankings over students.

Each year, there will be a different cohort of students applying. But assuming that exams are consistent from one year to the next, we expect that the overall distribution of students is similar from one year to the next. A standard statistical approach is to consider each year as a draw from the same population.

Formally, we model college admission in a given year as a finite sample $F^k = [\eta^k, q^k]$ of size k from a continuum economy $E = [\eta, q]$. We normalize the total mass of students to be $\eta(\Theta) = 1$. The continuum economy represents the potential population of students, and η captures the joint distribution over student preferences and student scores (generated by the distribution of student grades). A finite sample $F^k = [\eta^k, q^k]$ is generated by randomly drawing k students from η and scaling the capacity vector $q^k = \lfloor kq \rfloor$.

Azevedo and Leshno (2016) show that if E has a unique stable matching (which generically holds) then that stable matching of F^k will converge to the stable matching of E as $k \to \infty$. Moreover, they give a characterization of the distribution of market clearing cutoffs P^k for the randomly drawn economy F^k.

Theorem 13 (Azevedo and Leshno (2016))**.** *Let P^* be a market clearing cutoff for the continuum economy $E = [\eta, q]$. Assume that that $\sum_c q_c < 1$, $D(\cdot|\eta)$ is differentiably continuous, and that $\partial D(P^*)$ is nonsingular. Then the asymptotic distribution of the difference between P^k and P^* satisfies*

$$\sqrt{k} \cdot (P^k - P^*) \xrightarrow{d} \mathcal{N}(0, \partial D(P^*)^{-1} \cdot \Sigma^D \cdot (\partial D(P^*)^{-1})'),$$

where $\mathcal{N}(\cdot,\cdot)$ denotes a C-dimensional normal distribution with given mean and covariance matrix. The matrix Σ^D is given by $\left(\Sigma^D\right)_{cd} = \begin{cases} q_c(1-q_d) & \text{if } c = d, \\ -q_c q_d & \text{if } c \neq d\,. \end{cases}$

3 End Notes

3.1 Other applications of random matching markets and rejection chains

Random matching markets and the analysis of rejection chains can be applied to answer other questions as well. Kojima and Pathak (2009) show that a college can benefit from misreporting its preferences if there is a rejection chain that brings a more preferred student to apply to the school. In a random many-to-one economy with short lists, in which a school can be matched to a constant number of students and the number of both schools and students is sufficiently large, schools will not be able to benefit by misreporting their preferences. Biró et al. (2020) analyze a large matching market in which colleges offer two kinds of positions, with and without a scholarship, as considered in Chapter **??**. In their model preferences are random, but a student that is rejected from a position with a scholarship is likely to apply to the same college without a scholarship. As a result, rejection chains are likely to return to the college, and colleges can benefit from misreporting their preference.

Kojima et al. (2013) and Ashlagi et al. (2014) use rejection chains to analyze a matching market with couples. While it is possible that a market with couples will have no stable matching, the NRMP succeeded to find a stable matching with couples since its redesign. These papers provide a theoretical explanation, showing that the market is likely to have a stable matching as long as there are not too many couples in the market.

Menzel (2015) estimates a random matching market. He formulates stable matching as a discrete choice problem, in which each agent chooses his match from a budget of agent from the other side who are willing to form a block pair. Analysis of rejection chains is used to show that this choice problem is well formulated because the agent's budget set is almost independent of the agent's stated preferences. As the market grows large, each agent's budget set includes more options. It is well known that agent utility in discrete choice models is inflated with the number of options, and Menzel (2015) provides proper scaling that ensures the market maintains its key properties as it grows. Lee (2017) analyzes markets with idiosyncratic and common taste shocks with bounded support and shows that as the market grows large all stable matchings become assortative on the common taste shocks and match agents to spouses that give almost the maximal possible idiosyncratic taste shock.

Che and Tercieux (2019) and Yariv and Lee (2014) use markets with randomly drawn preferences to evaluate welfare of different mechanisms and preference distributions. Liu and Pycia (2016) show that many mechanisms coincide as the market grows large.

3.2 Additional applications of continuum models

Abdulkadiroğlu et al. (2015) introduced a matching model for school choice with random tie breaking. Azevedo (2014) uses a continuum model to tractably study the market power of large firms in a matching markets, and shows that a large firm can benefit from reducing its capacity. The intuition is that the firm rejects marginal matches of low value and receives instead infra-marginal matches of high value. Che and Koh (2016) use a continuum model to study decentralized college admission when college face a risk of being over or under capacity.

Ashlagi and Shi (2016), Shi (2019) use a continuum model to optimize over different school choice mechanisms. Leshno and Lo (2020) provide a cutoff characterization for the top trading cycles mechanism for school choice. A tractable continuum model allows for optimization of investment in school quality. Chapter **??** discusses the use of continuum models to study assignment mechanisms without monetary transfers.

Continuum models are also used to analyze markets with complementarities (Azevedo and Hatfield (2012), Che et al. (2019)) and show existence of stable matchings under more general preferences. Leshno (2020) uses the continuum model to give a simple cutoff characterization of stable matchings when students have peer-dependent preferences.

Jagadeesan (2017) develops a more general matching model that allows for multi-sided matchings. Greinecker and Kah (2018) generalize the notion of stability to a one-to-one matching between arbitrary sets of men and women. Gonczarowski et al. (2019) use logical compactness to characterize matching markets with infinitely many agents.

Exercises

1. Show that if (i) or (ii) happen in the algorithm given in 1.3, then there are no stable matchings in which man m_0 is matched to a woman w such that $\mu(m_0) \succ_m w$.

2. Show that if (iii) happen in the algorithm given in 1.3, the resulting matching is a new stable matching in which m_0 is matched to a different wife such that $\mu'(m_0) \prec_{m_0} \mu(m_0)$.

3. Write the economy that captures the following and calculate its stable matching and market clearing cutoffs. There are n colleges with identical capacities $q_1 = q_2 = \cdots = q_n = q$. Student references are independent of any college's rankings, and a college's ranking over students is independent of other college's ranking. Student preferences drawn uniformly at random.

4. Write the economy that captures the following and calculate its stable matching and market clearing cutoffs. There are n colleges with identical capacities $q_1 = q_2 = \cdots = q_n = q$. Student references are independent of

any college's rankings, and a college's ranking over students is independent of other college's ranking. All student have the preference ordering $1 \succ 2 \succ 3 \cdots \succ n$.

5. Give an example with two colleges and a continuum of students where there are multiple stable matchings (hint: think of the opposite preferences economy). Find a slight perturbation to this economy that leads to an economy with a unique stable matching.

References

Abdulkadiroglu, Atila, Che, Yeon-Koo, and Yasuda, Yosuke. 2011. Resolving conflicting preferences in school choice: The" boston mechanism" reconsidered. *American Economic Review*, **101**(1), 399–410.

Abdulkadiroğlu, Atila, Che, Yeon-Koo, and Yasuda, Yosuke. 2015. Expanding "Choice" in School Choice. *American Economic Journal: Microeconomics*, **7**(1), 1–42.

Abdulkadiroğlu, Atila, Angrist, Joshua D, Narita, Yusuke, and Pathak, Parag A. 2017. Research design meets market design: Using centralized assignment for impact evaluation. *Econometrica*, **85**(5), 1373–1432.

Adachi, Hiroyuki. 2003. A search model of two-sided matching under nontransferable utility. *Journal of Economic Theory*, **113**(2), 182–198.

Ashlagi, Itai, and Shi, Peng. 2016. Optimal allocation without money: An engineering approach. *Management Science*, **62**(4), 1078–1097.

Ashlagi, Itai, Braverman, Mark, and Hassidim, Avinatan. 2014. Stability in large matching markets with complementarities. *Operations Research*, **62**(4), 713–732.

Ashlagi, Itai, Kanoria, Yash, and Leshno, Jacob D. 2017. Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, **125**(1), 69–98.

Azevedo, Eduardo M. 2014. Imperfect competition in two-sided matching markets. *Games and Economic Behavior*, **83**, 207–223.

Azevedo, Eduardo M, and Hatfield, John William. 2012. Complementarity and multidimensional heterogeneity in matching markets. *Unpublished mimeo.*

Azevedo, Eduardo M, and Leshno, Jacob D. 2016. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, **124**(5), 1235–1268.

Balinski, Michel, and Sönmez, Tayfun. 1999. A tale of two mechanisms: student placement. *Journal of Economic theory*, **84**(1), 73–94.

Biró, Péter, Hassidim, Avinatan, Romm, Assaf, Shorrer, Ran I, and Sóvágó, Sándor. 2020. Need versus Merit: The Large Core of College Admissions Markets. *arXiv preprint arXiv:2010.08631*.

Che, Yeon-Koo, and Koh, Youngwoo. 2016. Decentralized College Admissions. *J. Polit. Econ.*, **124**(5), 1295–1338.

Che, Yeon-Koo, and Tercieux, Olivier. 2019. Efficiency and stability in large matching markets. *Journal of Political Economy*, **127**(5), 2301–2342.

Che, Yeon-Koo, Kim, Jinwoo, and Kojima, Fuhito. 2019. Stable matching in large economies. *Econometrica*, **87**(1), 65–110.

Fleiner, Tamás, and Jankó, Zsuzsanna. 2014. Choice function-based two-sided markets: stability, lattice property, path independence and algorithms. *Algorithms*, **7**(1), 32–59.

Gale, D., and Shapley, L.L. 1962. College admissions and the stability of marriage. *American Mathematical Monthly*, **69**, 9–15.

Gonczarowski, Yannai A, Kominers, Scott Duke, and Shorrer, Ran I. 2019. To Infinity and Beyond: Scaling Economic Theories via Logical Compactness.

Greinecker, Michael, and Kah, Christopher. 2018. *Pairwise stable matching in large economies*. Tech. rept. Working Papers in Economics and Statistics.

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

Jagadeesan, Ravi. 2017. Complementary Inputs and the Existence of Stable Outcomes in Large Trading Networks. Page 265 of: *EC*.

Knuth, D. E., Motwani, R., and Pittel, B. 1990. Stable husbands. *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, **1**, 1–14.

Kojima, F., and Pathak, P. A. 2009. Incentives and Stability in Large Two-Sided Matching Markets. *American Economic Review*, **99**, 608–627.

Kojima, Fuhito, Pathak, Parag A, and Roth, Alvin E. 2013. Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics*, **128**(4), 1585–1632.

Lee, SangMok. 2017. Incentive compatibility of large centralized matching markets. *The Review of Economic Studies*, **84**(1), 444–463.

Leshno, Jacob, and Lo, Irene. 2020. The cutoff structure of top trading cycles in school choice. *The Review of Economic Studies*.

Leshno, Jacob D. 2020. Stable Matching with Peer Effects in Large Markets – Existence and Cutoff Characterization. *Unpublished mimeo*.

Liu, Qingmin, and Pycia, Marek. 2016. Ordinal efficiency, fairness, and incentives in large markets. *working paper*.

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

Menzel, Konrad. 2015. Large matching markets as two-sided demand systems. *Econometrica*, **83**(3), 897–941.

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

Roth, A.E. 1985. The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory*, **36**(2), 277–288.

Roth, Alvin E, and Peranson, Elliott. 1999. The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American economic review*, **89**(4), 748–780.

Shi, Peng. 2019. Optimal Priority-Based Allocation Mechanisms. *Available at SSRN 3425348*.

Yariv, L., and Lee, S. 2014. *On the Efficiency of Stable Matchings in Large Markets*. Working paper.