

Matching with Couples: Stability and Incentives in Large Markets

Fuhito Kojima Parag A. Pathak Alvin E. Roth*

First version: May 2010

This version: April 2013

Abstract

Accommodating couples has been a longstanding issue in the design of centralized labor market clearinghouses for doctors and psychologists, because couples view pairs of jobs as complements. A stable matching may not exist when couples are present. This paper's main result is that a stable matching exists when there are relatively few couples and preference lists are sufficiently short relative to market size. We also discuss incentives in markets with couples. We relate these theoretical results to the job market for psychologists, in which stable matchings exist for all years of the data, despite the presence of couples.

*Kojima: Department of Economics, Stanford University, Landau Economics Building, 579 Serra Mall, Stanford, CA 94305-6072. Email: fkojima@stanford.edu. Pathak: Department of Economics (E52), Massachusetts Institute of Technology, 50 Memorial Drive, Cambridge, MA 02142. Email: ppathak@mit.edu. Roth: Department of Economics, Stanford University, Landau Economics Building, 579 Serra Mall, Stanford, CA 94305-6072. Email: alroth@stanford.edu. We thank Jin Chen, Rezwan Haque, Fanqi Shi, and especially Dan Barron, Pete Troyan and Mengxi Wu for superb research assistance. We are also grateful to Delong Meng, Assaf Romm, Ran Shorrer, Joel Sobel and seminar participants at Calgary, Ecole Polytechnique, Washington University in St. Louis, the ERID Matching Conference "Roth and Sotomayor: Twenty Years After" at Duke University, the Coalition Theory Workshop in Marseilles, the EC conference in San Jose, the 2011 AMES "pre-conference" at Hanyang, and the NBER Market Design conference for comments. Elliott Peranson and Greg Keilin provided invaluable assistance in obtaining and answering questions about the data used in this paper on behalf of National Matching Services and the Association of Psychology Postdoctoral and Internship Centers. The National Science Foundation provided research support.

1 Introduction

One of the big 20th century transformations of the American labor market involves the increased labor force participation of married women, and the consequent growth in the number of two-career households.¹ When a couple needs two jobs, they face a hard problem of coordination with each other and with their prospective employers. The search and matching process for spouses can involve very different timing of searches and hiring. A couple may be forced to make a decision on a job offer for one member of the couple before knowing what complementary jobs may become available for the other or what better pairs of jobs might become available elsewhere.

An unusually clear view of this problem can be found in the history of the entry-level labor market for American doctors. Since the early 1900s, new U.S. medical graduates have been first employed as “residents” at hospitals, where they work under the supervision of more senior, licensed doctors. This market experienced serious problems having to do with the timing of offers and acceptances, and this *unraveling* of the market led to the creation of a centralized clearinghouse in the 1950s that drew high rates of participation (see Roth (1984, 2003) and Roth and Xing (1994) for further details). Medical graduates were almost all men throughout this period, but by the 1970s there were enough women graduating from medical school so that it was not unheard of for two new medical graduates to be married to each other.² Many couples felt that the existing clearinghouse did not serve them well, and starting in the 1970s, significant numbers of these couples began seeking jobs outside of the clearinghouse.

Roth (1984) argues that this was because the matching algorithm used until then did not allow couples to appropriately express preferences. That paper shows that, in a market without couples, the 1950s clearinghouse algorithm is equivalent to the deferred acceptance algorithm of Gale and Shapley (1962), and that it produces a *stable* matching for any reported preferences – loosely speaking, this is a matching such that there is no pair of hospital and doctor who want to be matched with each other rather than accepting the prescribed matching.³ It then observes that the algorithm often fails to find a stable matching when there are couples, and argues that a main problem of the mechanism is that (prior to the 1983 match) it did not allow couples to report preferences over *pairs* of positions, one for each member of the couple. Roth and Peranson (1999) describe the current algorithm, which elicits and uses couples’ preferences over pairs of positions, and the Roth-Peranson algorithm has been used by more than 40 centralized clearinghouses including the American labor market for new doctors, the National Resident Matching Program (NRMP).⁴

But the problem is difficult even if couples are allowed to express their preferences over pairs of positions, because there does not necessarily exist a stable matching in markets with couples (Roth,

¹See, for instance, Costa and Kahn (2000) for a description of the trends in the labor market choices for college-educated couples since World War II.

²In the 1967-68 academic year, 8% of the graduates of U.S. medical schools were women. By 1977-78 this fraction had risen to 21%, and by 2008-09 to 49% (Jonas and Etzel (1998), and <http://www.aamc.org/data/facts/charts1982to2010.pdf>).

³Section 3 provides a precise definition of our stability concept.

⁴See Roth (2007) for a recent list of these clearinghouses as well as a survey of the literature. See also Sönmez and Ünver (2009).

1984). However, some matching clearinghouses regularly entertain high rates of participation and produce matchings that are honored by participants. In fact, it has been reported that there have only been a few occasions in which a stable matching was not found over the last decade in several dozen annual markets (Peranson, private communication). Moreover, in the largest of these markets, the NRMP, Roth and Peranson (1999) run a number of matching algorithms using submitted preferences from 1993, 1994 and 1995 and find no instance in which any of these algorithms failed to produce a stable matching. Why do these matching clearinghouses produce stable outcomes from submitted preferences even though existing theory suggests that stable matchings may not exist when couples are present?

This is the puzzle we address, and this paper argues that the answer may have to do with the size of the market. We consider a sequence of markets indexed by the number of hospitals. Doctors have preferences, drawn from a distribution over a number of hospitals, and this number is allowed to grow as the total number of hospitals grows, but more slowly. These are preferences after interviews have taken place.⁵ When the number of couples grows sufficiently slowly relative to market size, under some regularity conditions, our main result demonstrates that the probability that a stable matching exists converges to one as the market size approaches infinity. Moreover, we provide an algorithm that finds a stable matching with a probability that converges to one as the market size approaches infinity. We also discuss incentives for doctors and hospitals to report their preferences truthfully to the clearinghouse in markets with couples.

As our theoretical analysis only provides limit results, we study data on submitted preferences from the centralized market for clinical psychologists. In the late 1990s, the market evolved from a decentralized one (Roth and Xing, 1997) to one employing a centralized clearinghouse, where a key design issue was whether it would be possible to accommodate the presence of couples. Keilin (1998) reports that under the old decentralized system couples had difficulties coordinating their internship choices. In 1999, clinical psychologists adopted a centralized clearinghouse using an algorithm based on Roth and Peranson (1999), in which couples are allowed to express preferences over hospital pairs. We explore a variation of the Roth-Peranson procedure to investigate the existence of a stable matching for years 1999-2007. Using our algorithm, we are able to find a stable matching with respect to the stated preferences of participants in all nine years. To investigate whether our asymptotic arguments provide a guide for realistic market sizes, we also simulate markets under different assumptions on the number of couples. We draw preferences for agents from a uniform distribution and from a distribution calibrated using data from the clinical psychology market. Our simulations show that a stable matching is more likely to be found in large markets. For example, if the market size, i.e. the number of hospital seats, is at least 2,000 and the number of couples is equal to the square root of market size, a stable matching exists at least 96% of the time.

⁵The length of doctors preference lists can be expected to grow more slowly than the number of hospitals because other elements of the environment that limit the number of interviews are not also growing (e.g. the number of days between November and February and the time required for interviews remain constant). We do not model frictions involved in interviewing at different hospitals in this paper. A model with explicit costs of interviewing would be substantially different than the current framework.

Related Literature

This paper is related to several lines of work. First, it is part of research in two-sided matching with couples. Existing studies are mostly negative: Roth (1984) and unpublished work by Sotomayor show that there does not necessarily exist a stable matching when there are couples, and Ronn (1990) shows that it may be computationally hard even to determine if a stable matching exists. Klaus and Klijn (2005) provide a maximal domain of couple preferences that guarantees the existence of stable matchings.⁶ While their preference domain has a natural interpretation, our paper finds that the submitted preferences of almost all couples in our psychology market data violate their condition.⁷ This empirical fact motivates our appeal to large market arguments.

The second line of studies related to this paper is the growing literature on large matching markets. Roth and Peranson (1999) conduct simulations based on data from the NRMP which includes couples and randomly generated data. They find suggestive evidence that in large markets, a stable matching is likely to exist and stable matching mechanisms are difficult to manipulate. We also examine data from the psychologist market and demonstrate that a stable matching exists for all years we have access to data. One of the findings of Roth and Peranson (1999) is that the opportunities for manipulation vanish in large markets if doctor preference lists are bounded as market size grows, but such a result does not hold if each doctor lists all hospitals in her preference list. Moreover, in the market for clinical psychologists and the NRMP, each applicant can physically interview at a small number of potential employers. These two considerations motivate our theoretical analysis of a model where doctor preference lists are allowed to grow, but in a controlled manner.

Several recent papers have studied incentive issues in matching models with a large number of agents, following the setup and techniques in Immorlica and Mahdian (2005) who consider a one-to-one matching model and establish that hospitals are unlikely to be able to manipulate the doctor-optimal stable mechanism in a large market.⁸ Kojima and Pathak (2009) build on and extend this analysis for a many-to-one market. While the use of large market arguments in our paper is similar to these studies, the questions are substantially different. A stable matching always exists in the model without couples, while it is not guaranteed to exist when couples are present. This paper's use of large market arguments to establish the existence of a stable matching is, as far as we know, new in the matching literature. Our analysis also relaxes some commonly used assumptions in the analysis of large matching markets.

Large market arguments have been used in a number of other recent studies of matching mechanisms

⁶Related contributions include Dutta and Massó (1997) and Dean et al. (2006) who investigate the properties of matching models with alternative formulations of couples.

⁷Preference restrictions are also used to study incentives for manipulation in matching markets. The restriction under which incentive compatibility can be established is often very strong, as shown by Alcalde and Barberà (1994), Kesten (2006b), Kojima (2007), and Konishi and Ünver (2006b) for various kinds of manipulations. Similarly, in the context of resource allocation and school choice (Abdulkadiroğlu and Sönmez, 2003), necessary and sufficient conditions for desirable properties such as efficiency and incentive compatibility are strict (Ergin, 2002; Kesten, 2006a; Haeringer and Klijn, 2009; Ehlers and Erdil, 2010). Another approach is based on incomplete information (Roth and Rothblum, 1999; Ehlers, 2004, 2008; Kesten, 2010; Erdil and Ergin, 2008).

⁸See also Knuth et al. (1990) on the probabilistic analysis of large matching markets.

(Bulow and Levin, 2006; Abdulkadirođlu et al., 2008; Budish, 2011; Che and Kojima, 2010; Kojima and Manea, 2010; Manea, 2009; Liu and Pycia, 2011; Azevedo and Leshno, 2012; Lee, 2013). We describe subsequent work by Ashlagi et al. (2011) in Section 6. While the analysis of large markets is relatively new in the matching literature, it has a long tradition in economics. For example, Roberts and Postlewaite (1976) and Jackson and Manelli (1997) show that, under some conditions, the Walrasian mechanism is difficult to manipulate in large exchange economies. Similarly, Gresik and Satterthwaite (1989) and Rustichini et al. (1994) study incentive properties of a large class of double auction mechanisms.

Finally, a couple preference is a particular form of complementarity, and this paper can be put in the context of the larger research program on the role of complementarities in resource allocation. Complementarities have been identified to cause non-existence of desirable solutions in various contexts of resource allocation. There has been a recent flurry of investigations on complementarities and existence problems in auction markets (Milgrom, 2004; Gul and Stacchetti, 2000; Sun and Yang, 2009), general equilibrium with indivisible goods (Bikhchandani and Ostroy, 2002; Gul and Stacchetti, 1999; Sun and Yang, 2006), and matching markets (Hatfield and Kominers, 2009; Ostrovsky, 2008; Sönmez and Ünver, 2010; Pycia, 2012).

The layout of this paper is as follows. The next section describes some features of the market for clinical psychologists and lays out a series of stylized facts on matching with couples based on data from this market. Section 3 defines the model and describes a simple theory of matching with couples in a finite market. Section 4 introduces the large market assumptions. Section 5 states our main results on existence, while Section 6 discusses incentives, robustness, and extensions of our results. Section 7 concludes.

2 The Market for Internships in Professional Psychology

2.1 Background

The story of how design has been influenced by the presence of couples in the National Resident Matching Program (NRMP) has parallels in the evolution of the market for internships in professional psychology.⁹ Roth and Xing (1997) described this market through the early 1990s. From the 1970's through the late 1990's, this market operated in a decentralized fashion (with frequent rule changes), based on a “uniform notification day” system in which offers were given to internship applicants over the telephone within a specific time frame (e.g., a 4 hour period on the second Monday in February). All acceptances and rejections of offers occurred during this period. Keilin (2000) described the system as “problematic, subject to bottlenecks and gridlock, encouraging the violation of guidelines, and resulting in less-than-desirable outcomes for participants.”

⁹To be clear, we will concentrate on the match run by the Association of Psychology Postdoctoral and Internship Centers (APPIC) for predoctoral internships in psychology, which involves clinical, counseling, and school psychologists. (This is distinct from the postdoctoral match in neuropsychology.)

In 1998-1999, the Association of Psychology Postdoctoral and Internship Centers (APPIC) switched to a system in which applicants and internship sites were matched by computer. A major debate in this decision was whether a centralized system could handle the presence of couples. In the old decentralized scheme it was challenging for couples to coordinate their internship choices. Keilin (1998) reports that one partner could be put in the position of having to make an immediate decision about an offer without knowing the status of the other partner. Following the reforms of the National Resident Matching Program, a new scheme which allowed couples to jointly express their preferences was adopted.

With the permission of the APPIC, the organization which runs the matching process, National Matching Services, provided us with an anonymized dataset of the stated rank order lists of single doctors, couples and hospitals and hospital capacities for the first nine years of the centralized system. Because of privacy concerns, the APPIC dataset does not include any demographic information on applicants, and includes only limited identifying information on programs.

2.2 Stylized facts

This section identifies some stylized facts from the internship market for professional psychologists. While we do not have as detailed data from the NRMP, we also mention related facts from that market when appropriate using information from the NRMP's annual reports.

The data are the stated preferences of market participants, so their interpretation may require some caution. There are at least two parts to the process by which market participants form their preferences: (1) they determine which applicants or internship programs may be attractive, and participate in interviews, and (2) after interviewing, they determine their rank ordering over the applicants or internship programs they have interviewed. The model in this paper and the data do not allow us to say much about the first stage of the application process. In determining where to interview, applicants likely factor in the costs of traveling to interviews, the program's reputation and a host of other factors. Programs consider, among other things, the applicant's recommendation letters and suitability for their program in deciding whom to interview. Once market participants have learned about each other, they must come up with their rank ordering. For the empirical analysis in this section, we abstract away from the initial phase of mutual decisions of whom to interview, and our interpretation is that the data reflect the preferences formed *after* interviews. This, and the fact that participants only seem to rank those with whom they have had interviews, likely accounts for the relatively short rank order lists.

Even with this interpretation, the reported post-interview preferences may be manipulations of the true post-interview preferences of market participants since truth-telling is not a dominant strategy for all market participants. However, there are at least two reasons why treating submitted preferences as true preferences may not be an unrealistic approximation. First, as noted in Section 6.1, the organizers of the APPIC match emphasize repeatedly that market participants should declare their preferences truthfully. Second, in our working paper (Kojima et al., 2010) we provide assumptions under which truth-telling is an approximate equilibrium in large markets.

Table 1 presents some summary statistics on the market. On average, per year, there are 3,010 single applicants and 19 pairs of applicants who participate as couples. In early years, there were just under 3,000 applicants, but the number of applicants has increased slightly in the most recent years. The number of applicants who participate as couples has remained relatively small, varying between 28 and 44 (i.e. between 14 and 22 couples) which is about 1% of all applicants.¹⁰ On the surface, the small number of couples may appear surprising, but this number represents cases where both couple members look for jobs in professional psychology in the same year. If couple members wish to work in different fields, or even in the same field, but in different years, each couple member simply applies as a single applicant. In the National Resident Matching Program, from 1992-2009, there were on average 4.4% of applicants participating as couples, with slightly more couples participating in the most recent years (NRMP, 2009).

Fact 1: *Applicants who participate as couples constitute a small fraction of all participating applicants.*

Panel A of Table 1 shows the length of the rank order lists for applicants and programs. On average across years, single applicants rank between 7 and 8 programs. Since there are 1,094 programs on average, this means that the typical applicant ranks less than 1% of all possible programs. Even at the extreme, the length of the longest single applicant's rank order list is about 6.7% of all possible programs. In the NRMP, the length of the applicant preference list is about 7-9 programs, which would be roughly 0.3% of all possible programs.¹¹ This may not be surprising because an applicant typically ranks a program only after she interviews at the program, and each applicant receives and can travel to only a limited number of interviews.

For couples, each entry in the rank order list is a pair of programs (or being unmatched). The typical rank order list of couples averages 81 program pairs. However, the rank order list of a couple has entries for both members, so there are many duplicate programs. When we consider the number of distinct programs ranked by a couple, it is similar to the number ranked by single applicants: on average, there are about 10 distinct programs listed by each couple member. At the extreme in our dataset, the maximum number of distinct programs ranked by a couple member is 1.9% of all programs.

Of course, the fact that a doctor has a short preference list does not mean they prefer to leave the profession if they cannot obtain one of their stated choices. Given our interpretation of preferences as those formed after interviews, the short rank order list means doctors only interview at a fraction of possible hospitals and they may have a complete ranking over these options. In the event that they are unassigned, they either participate in the after-market in which they can learn about additional hospitals or postpone their training for a year, as is commonly done by doing a research year to strengthen one's credentials.

Fact 2: *The length of the rank order lists of applicants who are single or couples is small relative to the number of possible programs.*

¹⁰As the example in the next section shows, even one couple in the market may lead to non-existence of a stable matching.

¹¹This information is not available separately for single applicants and those who participate as couples in the NRMP.

The next issue we examine is the distribution of applicant preferences. In Figure 1, we explore the popularity of programs in our data. For each program, we compute the total number of students who rank that program as their top choice. We order programs by this number, with the program with the highest number of top choices on the left and programs that no one ranks as their top choice on the right. Figure 1 shows the distribution of popularity for 2003. In this year, the most popular program was ranked as the top choice by 19 applicants, and there are 189 programs that are not ranked as a top choice by any applicant. The other years of our dataset display a similar pattern. Averaged across all years, the most popular program is ranked as a top choice by 24 applicants, and about 208 programs are not ranked as a top choice by anyone. The fraction of applicants ranking the most popular program as their first choice is only 0.8%. (Recall that these are preferences stated after interviews have been conducted, so it does not preclude the possibility that there are popular programs that receive many applications but only interview a small subset of applicants.)

Fact 3: *The most popular programs are ranked as a top choice by a small number of applicants.*

The only identifying information we have on programs are geographic regions where they are located. The eleven geographic regions in our dataset are ten regions in the US, each of which corresponds to the first digit of the zip code of the program’s location, and one region for all of Canada. Most programs are concentrated on the West Coast and in the Northeast. In Table 1, we report the number of distinct regions ranked by applicants. Half of single applicants rank at most two regions. Couples, on the other hand, tend to rank slightly more regions.

For a given couple rank order list, we also compute the fraction of entries on their submitted list that have both jobs in the same region. On average, 73% of a couple’s rank order list is for programs in the same region.

Fact 4: *A pair of internship programs ranked by doctors who participate as a couple tend to be in the same region.*

In the psychology market, there are about 1,100 internship programs. The average capacity is about 2.5 seats, and more than three quarters of programs have three or fewer spots. The total capacity of internship programs is smaller than the total number of applicants who participate, which implies that each year there will be unmatched applicants. This is also true in the NRMP where the number of positions per applicant ranges from 0.75-0.90 over the period 1995-2009 (NRMP, 2009).

Even though there are more applicants than programs, in the APPIC match, there are a sizable number of programs that are unfilled at the end of the regular match. According to the APPIC’s statistics, during 1999-2007, on average 17% of programs had unfilled positions. In the NRMP, a similar proportion of programs had unfilled seats. In 2009, for instance, 12% of programs had unfilled positions.¹²

¹²In practice, to place these remaining applicants, both of these markets have a decentralized aftermarket where positions are filled. In this market, applicants can communicate and informally interview with places they did not initially consider,

Fact 5: *Even though there are more applicants than positions, many programs still have unfilled positions at the end of the centralized match.*

2.3 Stable Matchings in the Market for Psychologists

We next investigate whether a stable matching exists in the market for psychologists.¹³ Roughly speaking, this is a matching such that there is no pair of hospital and applicant who prefer each other to the prescribed matching.¹⁴ We use a variant of the procedure by Roth and Peranson (1999) to compute a stable matching.¹⁵ For each of the 9 years of data, the first column of Table 2 shows that a stable matching exists in the market with couples with respect to submitted preferences.¹⁶

Fact 6: *A stable matching with respect to submitted preferences exists in all nine years in the market for psychologists.*

We also compare the assignment of single applicants at the stable matching we find in a market with couples to their assignment in the applicant-optimal stable matching in a market without couples in Table A1. While adding couples to the market could in principle affect the assignment received by many single applicants, in practice it has little effect. This can be seen by comparing the overall distribution of choice received for single applicants in a stable matching in markets without couples and with couples. Moreover, Table A2 reports the exact number of single applicants who receive a less preferred assignment in the market with couples. On average, there are 19 couples or 38 applicants who participate as couples

but under a very short time limits. In recent years there have been proposals to eliminate these processes completely (see, e.g., Supplemental Offer and Acceptance Program (SOAP) described on <http://www.nrmp.org/soap.pdf>, accessed on October 22, 2010.) The model in this paper is only about the main match and does not model this decentralized aftermarket, though understanding how it may interact with the main round is an interesting question for future work. Some discussion of after-market scrambles can be found at <http://marketdesigner.blogspot.com/search/label/scramble>; for discussion of the SOAP, see particularly <http://marketdesigner.blogspot.com/2012/06/first-year-of-new-medical-residency.html>.

¹³The model we analyze in this paper allows employers to have preferences over sets of applicants provided that the preferences are responsive. Our data on program rank order lists consist only of preferences over individual applicants. We do not know, for instance, whether a program prefers their first and fourth ranked applicants over their second and third ranked applicants. To compute a stable matching in the market for psychologists, it is necessary to specify how comparisons between individual applicants relate to comparisons between sets of applicants. For the empirical computation, when comparing sets of applicants D_1 and D_2 , we assume that D_1 is more preferred to D_2 if the highest individually ranked applicant in D_1 who is not in D_2 is preferred to the highest individual ranked applicant in D_2 who is not in D_1 . This would imply that the first and fourth ranked applicant are preferred over the second and third ranked applicant. (We take advantage here of the more flexible formulation of preferences over sets that we employ, compared to that used in practice.)

¹⁴Section 3 provides a precise definition of our stability concept.

¹⁵Our variation has a different sequencing of applications from single applicants and couples than that described in Roth and Peranson (1999). That paper gives some evidence that these sequencing decisions have little impact on the success of the procedure.

¹⁶We focus on a particular stable matching in the market with couples, since we are unable to compute the entire set of stable matchings. There may be a reason to suspect that this set is small. In Table A3, we compute stable matchings in the market without couples and find that very few applicants and programs have different assignments across the applicant-optimal and program-optimal stable matchings. Moreover, there is evidence from real-life applications of the college admissions model without couples that the size of the stable set is small: see for example Roth and Peranson (1999) Table 9 which examines the market for thoracic surgeons, and Pathak and Sönmez (2008), p.1645 on Boston school choice.

in the market and because of their presence, only 63, or 2% of single applicants obtain a lower choice. This corresponds to about 3 single applicants obtaining a different assignment per couple.

Fact 7: *Across stable matchings, most single applicants obtain the same position in the market without couples as in the market with couples.*

3 A Simple Theory of Matching with Couples

3.1 Model

A matching market consists of hospitals, doctors, and their preferences. Let H be the set of hospitals and \emptyset be the outside option for doctors. Define $\tilde{H} = H \cup \{\emptyset\}$. S is the set of single doctors and C is the set of couples of doctors. Each couple is denoted by $c = (f, m)$, where f and m denote the first and second members of couple c respectively. When we need to refer to the members of a specific couple c , we sometimes write (f_c, m_c) . Let $F = \{f | (f, m) \in C \text{ for some } m\}$ and $M = \{m | (f, m) \in C \text{ for some } f\}$ be the sets of first and second members that form couples. Let $D = S \cup F \cup M$ be the set of doctors.

Each single doctor $s \in S$ has a preference relation R_s over \tilde{H} . We assume that preferences are strict: if $hR_s h'$ and $h'R_s h$, then $h = h'$. We write $hP_s h'$ if $hR_s h'$ and $h \neq h'$. If $hP_s \emptyset$, we say that hospital h is **acceptable** to single doctor s .

Each couple $c \in C$ has a preference relation R_c over $\tilde{H} \times \tilde{H}$, pairs of hospitals (and being unmatched). We assume that preferences of couples are strict with P_c denoting the asymmetric part of R_c . If $(h, h')P_c(\emptyset, \emptyset)$, then we say that pair (h, h') is acceptable to couple c . We say that hospital h is **listed by** R_c if there exists $h' \in \tilde{H}$ (so h' may be \emptyset) such that either $(h, h')P_c(\emptyset, \emptyset)$ or $(h', h)P_c(\emptyset, \emptyset)$.

Each hospital $h \in H$ has a preference relation over 2^D , all possible subsets of doctors. We assume preferences of hospitals are strict. Let $h \in H$ and κ_h be a positive integer. We say that preference relation \succeq_h is **responsive with capacity** κ_h if it ranks a doctor independently of her colleagues and disprefers any set of doctors exceeding capacity κ_h to being unmatched (see Appendix A.1 for a formal definition). We follow much of the matching literature and assume that hospital preferences are responsive throughout the paper. Let R_h be the corresponding **preference list of hospital** h , which is the preference relation over individual doctors and \emptyset . We write $dP_h d'$ if $dR_h d'$ and $d \neq d'$. We say that doctor d is acceptable to hospital h if $dP_h \emptyset$. We write $\succeq_H = (\succeq_h)_{h \in H}$. We refer to a matching market Γ as a tuple $(H, S, C, (\succeq_h)_{h \in H}, (R_i)_{i \in S \cup C})$.

We proceed to define our stability concept in markets with couples. The descriptions are necessarily somewhat more involved than those in the existing literature because we allow for capacities of hospitals larger than one (we will elaborate on this issue in Section 3.1.1). First, it is convenient to introduce the concept of hospital choices over permissible sets of doctors. For any set of doctors and couples $D' \subseteq D \cup C$,

define

$$\begin{aligned} \mathcal{A}(D') = \{D'' \subseteq D \mid & \forall s \in S, \text{ if } s \in D'' \text{ then } s \in D', \\ & \forall c \in C, \text{ if } \{f_c, m_c\} \subseteq D'', \text{ then } (f_c, m_c) \in D', \\ & \text{if } f_c \in D'' \text{ and } m_c \notin D'', \text{ then } f_c \in D', \\ & \text{if } f_c \notin D'' \text{ and } m_c \in D'', \text{ then } m_c \in D'\}. \end{aligned}$$

In words, $\mathcal{A}(D')$ is the collection of sets of doctors available for a hospital to employ when doctors (or couples of doctors) D' are applying to it. Underlying this definition is the distinction between applications by individual couple members and those by couples as a whole. For example, if $(f, m) \in D' \cap C$ but $f, m \notin D'$, then the couple is happy to be matched to the hospital if and only if both members are employed together, while if $(f, m) \notin D'$ but $\{f, m\} \subseteq D'$, then the couple is happy to have one member matched to the hospital but not together.

For any set $D' \subseteq D \cup C$, define the **choice** of hospital h given D' , $Ch_h(D')$, to be the set such that

- $Ch_h(D') \in \mathcal{A}(D')$,
- $Ch_h(D') \succeq_h D''$ for all $D'' \in \mathcal{A}(D')$.

The choice $Ch_h(D')$ is the most preferred subset of doctors among those in D' such that each couple is either chosen or not chosen together if they apply as a couple.¹⁷

A **matching** specifies which doctors are matched to which hospitals (if any). Formally, a matching μ is a function defined on the set $\tilde{H} \cup S \cup C$, such that $\mu(h) \subseteq D$ for every hospital h , $\mu(s) \in \tilde{H}$ for every single doctor s , and $\mu(c) \in \tilde{H} \times \tilde{H}$ for every couple c where

- $\mu(s) = h$ if and only if $s \in \mu(h)$ and
- $\mu(c) = (h, h')$ if and only if $f_c \in \mu(h)$ and $m_c \in \mu(h')$.

When there are only single doctors in D' , the set $\mathcal{A}(D')$ is simply the set of subsets of D' . Hence the choice $Ch_h(D')$ is the subset of D' that is the most preferred by h . This is the standard definition of $Ch_h(\cdot)$ in markets without couples (see Roth and Sotomayor (1990) for example), and hence the current definition is a generalization of the concept to markets with couples.

A matching is **individually rational** if no player can be made better off by unilaterally rejecting some of the existing partners (see Appendix A.1 for a formal definition). We define different cases of how the relevant small coalitions can block a matching as follows:

- (1) A pair of a single doctor s and a hospital $h \in H$ is a **block** of μ if $hP_s\mu(s)$ and $s \in Ch_h(\mu(h) \cup s)$.¹⁸

¹⁷This formulation of hospital preferences involving couples is more general than currently implemented in practice, where hospitals' preferences are elicited only over individual members of a couple.

¹⁸We denote a singleton set $\{x\}$ simply by x whenever there is no confusion.

- (2) (a) A coalition $(c, h, h') \in C \times \tilde{H} \times \tilde{H}$ of a couple and two hospitals¹⁹, where $h \neq h'$, is a **block** of μ if
- $(h, h')P_c\mu(c)$,
 - $f_c \in Ch_h(\mu(h) \cup f_c)$, and
 - $m_c \in Ch_{h'}(\mu(h') \cup m_c)$.²⁰
- (b) A pair $(c, h) \in C \times H$ of a couple and a hospital is a **block** of μ if
- $(h, h)P_c\mu(c)$ and
 - $\{f_c, m_c\} \subseteq Ch_h(\mu(h) \cup c)$.

A matching μ is **stable** if it is individually rational and there is no block of μ .

3.1.1 Discussion of the solution concepts

Models of matching with couples where hospitals have multiple positions are a particular form of many-to-many matching because each couple may seek two positions.²¹ Various definitions of stability have been proposed for many-to-many matching, which differ based on the assumptions on what blocking coalitions are allowed (Sotomayor, 1999, 2004; Konishi and Ünver, 2006a; Echenique and Oviedo, 2006). Consequently, there are multiple possible stability concepts in matching with couples. The present definition of stability allows us to stay as close to the most commonly used pairwise stability as possible, by assuming away deviations involving large groups. Ruling out large coalitions appears to be reasonable because identifying and organizing large groups of agents may be difficult.

It is nevertheless important to understand whether our analysis is sensitive to a particular definition of stability. To address this issue, in Appendix A.2 we present an alternative definition of stability that allows for larger coalitions to block a matching. We show that the results of this paper hold under that definition as well.

Most studies in matching with couples have focused on the case in which every hospital has capacity one.²² Following the standard definition of stability in such models (see Klaus and Klijn (2005) for instance), we say that a matching μ is **unit-capacity stable** if

- (1) μ is individually rational,

¹⁹Note that the definition of a blocking coalition of a couple and two hospitals includes the case in which only one member of the couple changes hospitals, so it might be thought of as being a coalition of a couple and one hospital, with the hospital whose employee doesn't move being a passive participant in the blocking coalition.

²⁰For the purpose of this definition, we adopt a notational convention that Ch_\emptyset is an identity function, so the condition $d \in Ch_\emptyset(\mu(\emptyset) \cup d)$ is satisfied for any μ and $d \in D$.

²¹More precisely, Hatfield and Kojima (2008, 2010) point out that the model is subsumed by a many-to-many generalization of the matching model with contracts as analyzed by Hatfield and Milgrom (2005).

²²Some papers consider multiple positions of hospitals but treat a hospital with capacity larger than one as multiple hospitals with capacity one each. This approach is customary and usually innocuous when there exists no couple because most stability concepts are known to coincide in that setting (Roth, 1985). However, the approach has a consequence if couples exist since it leads to a particular stability concept. A different modeling approach is pursued by McDermid and Manlove (2009).

- (2) there exists no single doctor-hospital pair s, h such that $hP_s\mu(s)$ and $sP_h\mu(h)$, and
- (3) there exists no coalition by a couple $c = (f, m) \in C$ and hospitals (or being unmatched) $h, h' \in \tilde{H}$ with $h \neq h'$ such that $(h, h')P_c\mu(c)$, $fR_h\mu(h)$ and $mR_{h'}\mu(h')$.²³

Our concept of stability is equivalent to the unit-capacity stability as defined above if every hospital has responsive preferences with capacity one. To see this, first observe that condition (3) of unit-capacity stability is equivalent to the nonexistence of a block as defined in condition (2a) of our stability concept. Moreover, condition (2b) of our stability concept is irrelevant when each hospital has capacity one because a hospital with capacity one never prefers to match with two members of a couple. Finally, the remaining conditions for unit-capacity stability have direct counterparts in our definition of stability. Thus the stability concept employed in this paper is a generalization of the standard concept to the case where hospitals have multiple positions.

Also note that our stability concept is equivalent to the standard definition of (pairwise) stability when there exist no couples. More specifically, condition (2) of our stability concept is irrelevant if couples are not present, and condition (1) is equivalent to the nonexistence of a blocking pair which, together with individual rationality, defines stability in markets without couples.

3.2 The Existence Problem with Couples

We illustrate how the existence of couples poses problems in the theory of two sided matching. To understand the role of couples, however, it is useful to start by considering a matching without couples. In that context, the (doctor-proposing) **deferred acceptance algorithm** defined below always produces a stable matching (Gale and Shapley, 1962).

Algorithm 1. DOCTOR-PROPOSING DEFERRED ACCEPTANCE ALGORITHM

Input: a matching market $(H, S, (\succeq_h)_{h \in H}, (R_s)_{s \in S})$ without couples.

- *Step 1:* Each single doctor applies to her first choice hospital. Each hospital rejects its least-preferred doctor in excess of its capacity and all unacceptable doctors among those who applied to it, keeping the rest of the doctors temporarily (so doctors not rejected at this step may be rejected in later steps).

In general,

- *Step t:* Each doctor who was rejected in Step (t-1) applies to her next highest choice (if any). Each hospital considers these doctors and doctors who are temporarily held from the previous step together, and rejects the least-preferred doctors in excess of its capacity and all unacceptable doctors, keeping the rest of the doctors temporarily (so doctors not rejected at this step may be rejected in later steps).

²³We adopt the notational convention that $dR_\emptyset d'$ for any $d, d' \in D \cup \emptyset$.

The algorithm terminates at a step where no doctor is rejected. The algorithm always terminates in a finite number of steps. At that point, all tentative matchings become final. Gale and Shapley (1962) show that for any given market without couples, the matching produced by the deferred acceptance algorithm is stable. Furthermore, they show that it is the doctor-optimal stable matching, the stable matching that is weakly preferred to any other stable matching by all doctors.

By contrast, stable matchings do not necessarily exist even when there is only one couple in the market (shown by Roth (1984) and an unpublished work by Sotomayor). This fact is illustrated in the following example, based on Klaus and Klijn (2005).

Example 1. Let there be a single doctor s and a couple $c = (f, m)$ as well as two hospitals h_1 and h_2 , each with capacity one. Suppose the acceptable matches for each agent, in order of preference, are given by:

$$\begin{array}{ll} R_c : (h_1, h_2) & R_s : h_1, h_2 \\ \succeq_{h_1} : f, s & \succeq_{h_2} : s, m. \end{array}$$

We illustrate that there is no stable matching in this market, by considering each possible matching.

- (1) Suppose $\mu(c) = (h_1, h_2)$. Then single doctor s is unmatched. Thus single doctor s and hospital h_2 block μ because s prefers h_2 to her match $\mu(s) = \emptyset$ and h_2 prefers s to its match $\mu(h_2) = m$.
- (2) Suppose $\mu(c) = (\emptyset, \emptyset)$.
 - (a) If $\mu(s) = h_1$, then (c, h_1, h_2) blocks μ since couple c prefers (h_1, h_2) to their match $\mu(c) = (\emptyset, \emptyset)$, hospital h_1 prefers f to its match $\mu(h_1) = s$ and hospital h_2 prefers m to its match $\mu(h_2) = \emptyset$.
 - (b) If $\mu(s) = h_2$ or $\mu(s) = \emptyset$, then (s, h_1) blocks μ since single doctor s prefers his first choice hospital h_1 to both hospital h_2 and \emptyset while h_1 prefers s to its match $\mu(h_1) = \emptyset$.

Klaus and Klijn (2005) identify a sufficient condition to guarantee the existence of a stable matching called weak responsiveness. A couple's preferences are said to be responsive if an improvement in one couple member's assignment is an improvement for the couple. Preferences are said to be weakly responsive if the requirement applies to all acceptable positions.²⁴ The preferences of couples in Example 1 do not satisfy this condition. If, for instance, the couple's preferences are $(h_1, h_2), (h_1, \emptyset), (\emptyset, h_2), (\emptyset, \emptyset)$, in order of preference, then it satisfies responsiveness and a stable matching exists. Klaus and Klijn (2005) write that "responsiveness essentially excludes complementarities in couples' preferences." They showed that:

- (1) if the preferences of every couple are weakly responsive, then there exists a stable matching.

²⁴See Klaus et al. (2009) for formal definition.

- (2) if there is at least one couple whose preferences violate weak responsiveness while satisfying a condition called “restricted strict unemployment aversion,” then there exists a preference profile of other agents such that preferences of all other couples are weakly responsive but there exists no stable matching.

Their second result says that the class of weakly responsive preferences is the “maximal domain” of preferences. That is, it is the weakest possible condition that can be imposed on individual couples’ preferences that guarantees the existence of stable matchings.²⁵

There seem to be many situations in which couple preferences violate weak responsiveness. One reason may be geographic, as stated as Fact 4 in Section 2.2: both programs ranked as a pair by a couple tend to be in the same geographic region. For example, the first choice of a couple of medical residents may be two residency programs in Boston and the second may be two programs in Los Angeles, while one member working in Boston and the other working in Los Angeles could be unacceptable because these two cities are too far away from each other. The coordinator of the Association of Psychology Postdoctoral and Internship Centers (APPIC) matching program writes in Keilin (1998) that “most couples want to coordinate their internship placements, particularly with regard to geographic location.” This suggests that violation of weak responsiveness due to geographic preferences is one of the representative features of couple preferences.²⁶

To further study this question empirically, we analyze the data on the stated preferences of couples from the APPIC. Stated preferences of couples may not be their true preferences since the truth-telling is not necessarily their dominant strategy. However, there are reasons that couples may not want to misrepresent their preferences. First, it may be complicated for a couple to determine a profitable deviation. Moreover, truth-telling may be focal, especially since clearinghouse organizers explicitly encourage participants to report their preferences truthfully. Finally, Kojima et al. (2010) provides conditions under which truth-telling is an approximate equilibrium when there are large numbers of market participants. During years for which we have data (1999–2007), preferences of only one couple out of 167 satisfy weak responsiveness. Thus the data suggest, in light of the results of Klaus and Klijn (2005), that it is virtually impossible to guarantee the existence of a stable matching in such markets with couples based on a domain restriction of preferences.

However, the fact that the preferences of the overwhelming majority (166 out of 167) of couples violate weak responsiveness does not mean that a stable matching does not exist in the psychologist market. Stable matchings have been found in many labor markets despite the presence of couples, and as we described in Section 2.3, we find a stable match for each of the nine years of the psychology market for which we have data. This motivates our desire to understand what market features enable the existence of stable matchings most of the time, when the known sufficient conditions on couples’ preferences do

²⁵Hatfield and Kominers (2009) show that the substitutes condition is a maximal domain in the absence of restricted strict unemployment aversion.

²⁶For an investigation of decision making among couples in the market for new Ph.D. economists, see Helppie and Murray-Close (2010).

not guarantee existence.

3.3 Sequential Couples Algorithm

The original deferred acceptance algorithm does not incorporate applications by couples. We consider an extension of the algorithm, which we call the **sequential couples algorithm**. While we defer a formal definition to Appendix A.3 for expositional simplicity, we offer an informal description as follows.

- (1) run a deferred acceptance algorithm for a sub-market composed of all hospitals and single doctors, but without couples,
- (2) one by one, place couples by allowing each couple to apply to pairs of hospitals in order of their preferences (possibly displacing some doctors from their tentative matches), and
- (3) one by one, place singles who were displaced by couples by allowing each of them to apply to a hospital in order of her preferences.

We say that the sequential couples algorithm **succeeds** if there is no instance in the algorithm in which an application is made to a hospital where an application has previously been made by a member (or both members) of a couple except for the couple who is currently applying. Otherwise, we declare a **failure** and terminate the algorithm.

Failure of the sequential couples algorithm does not mean that a stable matching does not exist. Therefore, in practice, a matching clearinghouse would be unlikely to declare failure when the sequential couples algorithm fails, but would instead consider some procedure to try to assign the remaining couples and find a stable matching. This is the main idea behind the Roth-Peranson algorithm (Roth and Peranson, 1999), which is the basis for the mechanism used in the NRMP, APPIC, and other labor markets. If the sequential couples algorithm would succeed, then the Roth-Peranson algorithm produces the matching reached by the sequential couples algorithm. However, the sequential couples algorithm and the Roth-Peranson algorithm are different in two aspects.²⁷

First, where the sequential couples algorithm fails, the Roth-Peranson algorithm proceeds and tries to find a stable matching. The algorithm identifies blocking pairs, eliminating instances of instability one by one, in a manner similar to Roth and Vande Vate (1990). Note that since a stable matching does not necessarily exist in markets with couples, the Roth-Peranson algorithm could cycle without terminating. However, the algorithm forces termination of a cycle and proceeds with processing other applicants. This sometimes ultimately results in a stable matching, and sometimes no stable matching is found. Second, in the Roth-Peranson algorithm, when a couple is added to the market with single doctors, any single doctor who is displaced by the couple is placed before another couple is added. By

²⁷A complete description of the Roth-Peranson algorithm, specifically how the algorithm terminates cycles and proceeds with processing, is not publicly available, but a more detailed description than the one provided here is offered by Roth and Peranson (1999), and a flowchart of the algorithm will appear in Roth (2013).

contrast, the sequential couples algorithm holds any displaced single doctor without letting her apply, until it processes applications by all couples.²⁸

The reason we focus on this simplified procedure is that the success of the sequential couples algorithm turns out to be sufficient to verify the existence of a stable matching (the proof is in Appendix A.3).

Lemma 1. *If the sequential couples algorithm succeeds, then the resulting matching is stable.*

To illustrate the main idea of Lemma 1, we consider how the sequential couples algorithm proceeds for the market in Example 1. In Step 1 of the algorithm, we run the doctor-proposing deferred acceptance algorithm in the sub-market without couples. Single doctor s proposes to hospital h_1 and is assigned there. Then in Step 2, we let couple c apply to their top choice (h_1, h_2) . Couple member f is preferred to s by h_1 and couple member m is preferred to a vacant position by h_2 . Thus f and m are tentatively assigned to h_1 and h_2 respectively while s is rejected. Then in Step 3, we let s apply to her next highest choice. In this case, she applies to hospital h_2 , where a couple member m has applied and been assigned before. At this point we terminate the algorithm and declare that it has failed.

To see why declaring a failure of the sequential couples algorithm is useful, suppose that we hypothetically continue the algorithm by allowing h_2 reject m as h_2 prefers s to m . Then the couple prefers being unassigned rather than having only f be matched to h_1 , so doctor f would like to withdraw his assignment from hospital h_1 . Suppose we terminate the algorithm at this point once f becomes unmatched. Then the resulting matching assigns no doctor to h_1 and s to h_2 . This matching is unstable because doctor s can block with hospital h_1 . On the other hand, if we continue the algorithm further by allowing s to match with h_1 , then the resulting matching is identical to the one obtained at the end of Step 1 of the sequential couples algorithm. This suggests that reasonable algorithms would cycle without terminating in this market.

The idea of declaring failure of the sequential couples algorithm is to avoid a situation like the above example, and turns out to be a useful criterion for judging whether the algorithm produces a stable matching. Of course, the algorithm sometimes fails even if there exists a stable matching, so the success of the algorithm is only a sufficient condition for the existence of a stable matching. What is remarkable is that looking at this particular sufficient condition turns out to be enough for establishing that a stable matching exists with a high probability in the environment we study in this paper. Moreover, there is a sense in which it is *necessary* to use an algorithm that finds a stable matching only in some instances, rather than one that always finds a stable matching whenever it exists. Ronn (1990) shows that the problem of determining whether or not a market with couples has a stable matching is computationally hard (NP-complete). The result suggests that it may be inevitable to employ an approach that does not always find a stable matching like our sequential couples algorithm.

Example 1 illustrates that the sequential couples algorithm does not necessarily succeed, and suggests that markets of any finite size would allow such a failure. We instead consider a large market environ-

²⁸As we will point out subsequently, our result also holds if we follow the sequencing of doctors as in the Roth-Peranson algorithm. We chose the current definition of the sequential couples algorithm for expositional simplicity.

ment with a random component in the preferences of the market participants. Our contribution is to demonstrate that, with high probability, the sequential couples algorithm succeeds, and hence a stable matching exists in this environment.

4 Large Markets

4.1 Random Markets

We have seen that a stable matching does not necessarily exist in a finite matching market with couples. To investigate how often a stable matching exists in large markets, we introduce the following random environment. A **random market** is a tuple $\tilde{\Gamma} = (H, S, C, \succeq_H, k, \mathcal{P}, \rho)$, where k is a positive integer, $\mathcal{P} = \{p_d(\cdot)\}_{d \in D}$ is a collection of probability distributions for each doctor d on H , and ρ is a function which maps two preferences over \tilde{H} to a preference list for couples (explained below). Each random market induces a market by randomly generating preferences of doctors as follows:

Preferences for Single Doctors: For each single doctor $s \in S$,

- Step 1: Select a hospital h with probability $p_s(h)$. List this hospital as the top ranked hospital of single doctor s .

In general,

- Step $t \leq k$: Select a hospital h with probability $p_s(h)$. Repeat until a hospital is drawn that has not been previously drawn in steps 1 through $t - 1$ or every hospital h such that $p_s(h) > 0$ has already been drawn. List this hospital as the t^{th} most preferred hospital of single doctor s .

In other words, to construct the preference list, we draw hospitals repeatedly without replacement (at most) k times according to distribution $p_s(\cdot)$. Single doctor s finds these k hospitals acceptable, and all other hospitals unacceptable. For example, if $p_s(\cdot)$ is the uniform distribution on H , then the preference list is drawn from the uniform distribution over the set of all preference lists of length k .

Preferences for Doctors who are Couples: Couples' preferences are formed by drawing preferences, R_f and R_m , for each doctor in the couple $c = (f, m)$. R_f and R_m are constructed from distributions $p_f(\cdot)$ and $p_m(\cdot)$ following the process used to generate preferences for a single doctor described above.

To construct the preference list for the couple $c = (f, m)$, define $\rho(R_f, R_m)$ to be a preference of the couple with the following restriction: if (h_1, h_2) is acceptable according to $\rho(R_f, R_m)$, then $h_1 R_f \emptyset$ and $h_2 R_m \emptyset$. This is the only restriction we place on ρ .

Preferences for Hospitals: Each hospital h has a responsive preference relation defined over sets of doctors \succeq_h such that all doctors are acceptable.

Discussion of modeling choices

The model assumes that doctors' preferences are drawn independently from one another in a statistical sense. However, doctors' preferences are not necessarily drawn from identical distributions and hence can be substantially different from one another. This modeling assumption can capture various possibilities for doctor preferences that may be important empirically. For example, the assumption allows a situation in which some popular hospitals are listed with higher probability than others by many doctors. It also allows preference distributions where doctors from different regions prefer programs from their own region to those outside of their region. Moreover, the assumption also allows for couples to have systematically different views on desirability of hospitals than single doctors. Allowing the doctors' preferences to differ in this way makes the assumption more general than the identical preference distribution assumption in Immorlica and Mahdian (2005), Kojima and Pathak (2009), and Manea (2009). In Section 6.3, we discuss the assumption in more detail and examine ways it can be relaxed.

The function ρ is a mapping that outputs a preference relation for each couple (f, m) given the pair of preferences R_f and R_m over \tilde{H} . One could interpret $\rho(R_f, R_m)$ as describing the outcome of household bargaining when preferences of the members are R_f and R_m , respectively. For example, the function ρ can represent a process in which any pair of hospitals that are too far away from each other is declared unacceptable, which seems to be consistent with the observed rank order lists of couples described earlier. We remain agnostic about ρ except that a hospital pair (h, h') is weakly acceptable for the couple under $\rho(R_f, R_m)$ only when h and h' are listed under R_f and R_m , respectively. In other words, no hospital appears in the preference list of a couple unless it is considered by the relevant member of the couple. Note that this, of course, does not impose that the couples preferences are weakly responsive. All our results are unchanged if we allow the function ρ to vary across different couples, but we model a common function ρ for all couples for expositional simplicity. Moreover, our results also hold when couples draw their preferences jointly from some distribution over pairs of hospitals.

Some NRMP participants who participate as couples are advised to form preferences by first forming individual rank order lists after interviewing with programs. Then, these individuals' lists serve as an input into the joint ranking of the couple. For instance, medical students who are couples at the University of Kansas Medical School are suggested to make a list of all possible program pair combinations from both individual rank order lists by computing the difference between the ranking number of the program on each individual's rank order list and trying to minimize this difference in their joint rank order list. This would be one example of a ρ function.²⁹

²⁹The details on this advice are available at <http://www.kumc.edu/som/medsos/cm.html>, accessed on March 20, 2010. The clearinghouse for new doctors in Scotland only allows couple members to submit individual rank order lists, in contrast to the model we analyze here. In that context, their mechanism combines these lists into a preference over pairs for the couple

The probabilistic structure we place on doctor preferences is unneeded for hospital preferences. Rather, hospital preferences can be arbitrary except for two important restrictions. First, hospital preferences are assumed to be responsive as in much of the literature on two-sided matching. The labor market clearinghouses which motivate our study impose this restriction by eliciting preferences over individual doctors.

The second important assumption on hospital preferences is that hospitals find all doctors acceptable. We make this assumption so that there are enough hospitals that can actually hire doctors in large markets. At first glance, this assumption seems violated in the data from the market for clinical psychologists as no program submits a rank order list of all doctors (for instance, as seen in Table 1, the average number of doctors listed in a hospital's preference list is 16.7 in our APPIC data). However, the programs rank most doctors who have ranked them, and an equivalent assumption is that each hospital finds acceptable only doctors who list that hospital, e.g. because hospitals (and doctors) will only rank doctors (and hospitals) they have interviewed. Clearly the existence result follows under this alternative assumption, because any stable matching at the original preference profile is also stable under the modified preference profile. The results also follow, at additional notational complexity, in a model where at least a constant fraction of hospitals find all doctors acceptable.

4.2 Regular Sequence of Random Markets

To analyze limit behavior of the matching market as the market becomes large, we consider a sequence of markets of different sizes. A **sequence of random markets** is denoted by $(\tilde{\Gamma}^1, \tilde{\Gamma}^2, \dots)$, where $\tilde{\Gamma}^n = (H^n, S^n, C^n, \succeq_{H^n}, k^n, \mathcal{P}^n, \rho^n)$ is a random market in which $|H^n| = n$ is the number of hospitals. Consider the following regularity conditions.

Definition 1. *A sequence of random markets $(\tilde{\Gamma}^1, \tilde{\Gamma}^2, \dots)$ is **regular** if there exist $\lambda > 0$, $a < \frac{1}{2}$, $b > 0$, $r \geq 1$, and $\gamma < \frac{1-2a}{r\lambda}$ such that for all sufficiently large n ,*

$$(1) |S^n| \leq \lambda n, |C^n| \leq bn^a,$$

$$(2) k^n \leq \gamma \log(n),$$

$$(3) \frac{p_d(h)}{p_d(h')} \leq r \text{ for all doctors } d \text{ in } D^n \text{ and hospitals } h, h' \text{ in } H^n.$$

Condition (1) requires that the number of single doctors does not grow much faster than the number of hospitals. Moreover, couples do not grow at the same rate as the number of hospitals and instead grow at the slower rate of $O(n^a)$ where $a < \frac{1}{2}$. This condition is motivated by Fact 1 that the number of couples is small compared with the number of hospitals or single doctors. Note that the assumption also implies that the total number of applicants $|S^n| + |C^n|$ is of order at most n and is consistent with either

using their individual lists and a table of positions that are determined to be geographically compatible by the mechanism. See the discussion of the Scottish Foundation Allocation Scheme at <http://www.nes.scot.nhs.uk/sfas/About/default.asp>, accessed on March 29, 2010.

more doctors than hospitals or fewer. Condition (2) assumes that the length of doctors' preference lists does not grow too fast when the number of market participants grows.³⁰ This assumption is motivated by Fact 2 in Section 2.2 that the length of rank order lists of doctors is small relative to the number of programs. Allowing the length of doctor's preference lists to vary does not change any of our results as long as each doctor's rank order list is no longer than $\gamma \log(n)$. Condition (3) requires that the popularity of different hospitals (as measured by the probability of being listed by doctors as acceptable) does not vary too much, as suggested by Fact 3.

5 Existence of Stable Matchings

As seen in Example 1, a stable matching does not necessarily exist when some doctors are couples. However, there is a sense in which a stable matching is likely to exist if the market is large as formalized by our main result:

Theorem 1. *Suppose that $(\tilde{\Gamma}^1, \tilde{\Gamma}^2, \dots)$ is a regular sequence of random markets. Then the probability that there exists a stable matching in the market induced by $\tilde{\Gamma}^n$ converges to one as the number of hospitals n approaches infinity.*

We defer the formal proof to Appendix A.3 and describe the argument here. Our proof involves analysis of the sequential couples algorithm in a regular sequence of random markets. By Lemma 1, we know that a stable matching exists whenever the algorithm succeeds. Our proof strategy is to show that the probability that the sequential couples algorithm succeeds converges to one as the market size approaches infinity.

Suppose that there are a large number of hospitals in the market. Given our assumptions on the distribution of couples' preferences, different couples are likely to prefer different pairs of hospitals. Hence, in Step 2 of the algorithm, members of two distinct couples are unlikely to apply to the same hospital. In such an instance, this step of the algorithm tentatively places couples without failure. Given that, it suffices to show that the single doctors displaced in Steps 2 and 3 (if any) are likely to be placed without applying to a hospital where a couple has applied. To show this, first we demonstrate that if the market is large, then it is a high probability event that there are a large number of hospitals with vacant positions at the end of Step 2 (even though there could be more applicants than positions: see Proposition 2 in the Appendix).³¹ Then, any single doctor is much more likely to apply to a hospital with a vacant position than to one of the hospitals that has already received an application by a couple member. Since every doctor is acceptable to any hospital by assumption, a doctor is accepted whenever an application is made to a vacant position. With high probability the algorithm places all the single doctors in Step 3, resulting in a success. Together with Lemma 1, we conclude that if the market is large

³⁰In Kojima et al. (2010), we considered the case where the doctor rank order list is bounded. We thank the referees for suggesting that we relax this assumption to allow for a growing rank order list length.

³¹Note that the feature that there are many hospitals with vacant positions is consistent with Fact 5 in Section 2.2, which states that there are many resident programs with vacant positions in practical matching markets.

enough, then the probability that there exists no stable matching can be made arbitrarily small. This completes the argument.

As explained in Section 3.3, the sequential couples algorithm is similar to but slightly different from the Roth-Peranson algorithm in the order of which doctors apply to hospitals. However, it is clear from the proof that the argument can be modified for the Roth-Peranson algorithm. Therefore, we have the following result as a corollary.

Corollary 1. *Suppose that $(\tilde{\Gamma}^1, \tilde{\Gamma}^2, \dots)$ is a regular sequence of random markets. Then the probability that the Roth-Peranson algorithm produces a stable matching in the market induced by $\tilde{\Gamma}^n$ converges to one as the number of hospitals n approaches infinity.*

6 Incentives, Robustness, and Extensions

6.1 Incentives

The previous section establishes our main result on the existence of a stable matching with respect to reported preferences that follow certain distributional assumptions. In practice, however, preferences are private information of market participants, and the matching clearinghouse needs to elicit this information. Thus a natural question is whether there is a mechanism that induces participants to report true preferences and produces a stable matching with respect to the true preferences.

One motivation for studying this question comes from the market for psychologists. The following advice is given to participants by clearinghouse organizers:³²

IMPORTANT: There is only one correct “strategy” for developing your Rank Order List: simply list your sites based on your true preferences, without consideration for where you believe you might be ranked by them. List the site that you want most as your #1 choice, followed by your next most-preferred site, and so on.

The previous paragraph is so important that we are going to repeat it: simply list your sites based on your true preferences.

Similar recommendations are made in other labor markets with couples. Below is the advice for participants offered by the National Resident Matching Program (NRMP).³³

Programs should be ranked in sequence, according to the applicant’s true preferences. . . . It is highly unlikely that either applicants or programs will be able to influence the outcome of the Match in their favor by submitting a list that differs from their true preferences.

³²“FAQ for Internship Applicants” in the APPIC website, http://www.appic.org/match/5_2.1.2.6.html, accessed on November 11, 2009.

³³“Rank Order Lists” in the NRMP website, http://www.nrmp.org/fellow/rank_order.html, accessed on November 11, 2009.

In these quotes, market participants are advised to report their true preferences to the matching authority. This advice may have been based on the incentive properties of stable mechanisms in markets without couples. For instance, without couples, in the doctor proposing deferred acceptance algorithm, truth-telling is a dominant strategy for doctors. However, in markets with couples, this is no longer the case.

In the working paper version (Kojima et al., 2010), we consider the issue of incentives after participants have conducted their interviews.³⁴ The question we study is whether given a particular mechanism which finds a stable matching with high probability, do applicants have an incentive to rank the programs that interviewed them truthfully? At a first glance, a positive result seems elusive: there exists no mechanism that is stable and strategy-proof even without couples. However, we provide conditions under which truth-telling is an approximate Bayes-Nash equilibrium in a large regular market.

6.2 Robustness

Theorem 1 is an asymptotic result, and the probability that a stable matching exists is not guaranteed to reach one in any finite market. On the other hand, any particular market has only finite numbers of market participants. In this section, we examine the relevance of the asymptotic result to the applications which motivate our analysis.

6.2.1 Speed of Convergence

The first issue is the order (speed) of convergence. Our model is quite general in terms of a number of parameters, making it elusive to establish an appealing result about the order of convergence.

However, the proof of the theorem allows us to evaluate the convergence probability for each given sequence of random markets, and this enables us to obtain a sharp result for special cases. In Appendix A.3.1, we show that if the number of couples is bounded and the length of each doctor’s preference list is bounded along the sequence of random markets (that is, $a = 0$ and $k^n \leq k$ for some constant k in Definition 1), then the probability that there is no stable matching approaches zero at least with the rate of convergence $O(1/n)$.

The key to this observation is inequality (19) in Appendix A.3, which provides a lower bound of the probability that a sequential couples algorithm successfully finds a stable matching. As explained there, the convergence order result above is obtained through a bounding exercise of this inequality for the special case. Note however that inequality (19) itself holds generally for any choice of parameters as long as our regularity and thickness conditions hold. Therefore, one can evaluate inequality (19) to obtain an order of convergence more generally. In that sense our analysis may shed light on the speed of convergence given any relevant information about the sequence of random markets.

There is an important sense in which our bounds do not exactly match the bounds required for finite markets like APPIC with 2,000 participants largely because of constants in expression (19). To

³⁴Before their interviews, an applicant may find many hospitals acceptable. However, doctors typically interview with only a small subset of hospitals, which naturally restricts the set of hospitals they consider.

see whether the qualitative insights from the theory are relevant for actual market sizes, we turn to simulation evidence in Section 6.2.5.

6.2.2 Proportion of Market Participants who are Couples

Our model made a number of assumptions, some of which could be relaxed. Perhaps the strongest assumption in our analysis relates to the growth rate of couples. Empirically, there are few couples in actual markets, but this fact does not directly imply the appropriate growth rate for a sequence of markets. Subsequent work by Ashlagi et al. (2011) studies a model with similar features. They state that a stable matching exists (with a slightly different notion of stability) when couples are allowed to grow at rate $n^{1-\epsilon}$ for $\epsilon > 0$. Their paper differs from ours in that they examine the implications of considering a particular sequence of proposals by couples and find the order which is least likely to generate existence problems.³⁵

6.2.3 Length of Doctor Rank Order Lists

As stated in Definition 1, our result holds even when the number k^n of hospitals that a doctor finds acceptable grows without a bound as the market size n grows, as long as its growth speed is slow enough. More specifically, our proof goes through if $k^n \leq \gamma \log(n)$ where $\gamma < \frac{1-2a}{16\lambda}$.

Our result is a limit one based on probability bounds. As such, our asymptotic prediction is not directly applicable to any finite economy. That being said, we can use the above expression to make some qualitative predictions. To do so, note that the upper bound of γ is decreasing in a and λ : Recall that a is the parameter representing the growth speed of couples, and λ is the upper bound of the ratio of the number of doctors to the number of hospitals. Thus the comparative statics here suggests that asymptotic existence holds for a faster growth speed of doctor preference list length when the number of couples grows more slowly and the number of doctors are smaller. The caveat of this argument is, of course, that our result provides only a lower bound of the existence probability, so our exercise here cannot be conclusive.

6.2.4 Performance of Sequential Couples Algorithm in APPIC

We found a stable matching in the APPIC dataset for all nine years using the algorithm described in Section B.2 in the Appendix. This algorithm is a variant of the sequential couples algorithm (SCA) with particular cycling and termination rules. Here we investigate whether the intuition expressed from the formal analysis of the SCA provides insight into existence for the clinical psychology market. Recall that we define the SCA as succeeding if there is no instance in the algorithm in which an application is

³⁵See also Biró and Irving (2010), for simulations and analysis related to a special case of the couples problem that arises in a medical labor market in Scotland, in which hospitals rank all applicants (including the individual members of couples) according to a common exam score. Biro and Irving show that the problem of determining if a stable matching exists remains computationally hard even in this special case, but simulations show that the probability of the set of stable matchings being empty is low when the proportion of couples is low.

made to a hospital where an application has previously been made by a member (or both members) of a couple except for the couple who is currently applying. Under this definition, failure of the SCA does not imply that a stable matching does not exist. Our formal results focused on this definition because it was sufficient to derive formal properties in the market with couples.

To examine how the sequential couples algorithm performs for the APPIC market, in column (2) of Table 2 we report on whether there is an ordering of couples for which a stable matching is found using the sequential couples algorithm. The definition of failure used in our proof provides only a lower bound, and for Table 2 we investigate a stronger definition of failure which still works for our purposes. We declare failure when a single doctor or couple applies to a program where a couple member is assigned and this applicant is more preferred than the couple member. This definition is preferred because declaring failure when an applicant applies to a program where a couple is assigned is relatively common in the APPIC dataset, even though it suffices for our formal argument and a stable matching continues to be found even when it occurs. If the SCA fails under this slightly stronger definition of failure, it also fails under the definition used in our proof. Even though the SCA by itself is a relatively naive procedure, remarkably, a simple application of the algorithm together with the appropriate sequencing of couples finds a stable matching in six out of nine years.³⁶

The fact that SCA finds a stable matching two thirds of the time even in a finite market suggests that analyzing its properties may provide intuition for why stable matchings are found in practice. The key phenomena are doctor’s short rank order lists, the small number of couples, and the limited overlap between preferences in doctor’s rank order lists. Indeed, in two of the six successful years, 2003 and 2005, it is never the case that an applicant even applies to a program where a couple is assigned as shown in column (3).³⁷ In another two of those successful years, there is limited overlap in couple’s rank order lists. In two years (1999 and 2001) a couple does not apply to a program in the SCA where another couple is assigned as shown in column (4).

6.2.5 Simulations varying Market Size and Number of Couples

Data from the APPIC market does not provide guidance on how the likelihood of existence changes with market size and the number of couples in the market. We next turn to simulations of markets under different assumptions on size, the number of couples, and the distribution of preferences. Figure 2 reports the fraction of markets where we find a stable matching for a one-to-one matching market with preferences drawn from a uniform distribution. We set single doctors’ rank order lists to have length 10 and couples’ joint rank order lists to have length of about 40 on average, and focus on varying the fraction of couples in the market.³⁸ (Section B.1 precisely describes how we construct market primitives for each market size.)

³⁶Since the ordering of couples is arbitrary and we do not want our conclusions to be influenced by a particular ordering in these columns we investigate what happens for 1000 random permutations of the ordering of couples and declare success if we find a stable matching for at least one of those orderings.

³⁷This is consistent with the definition of failure which is sufficient for our statement of Lemma 1.

³⁸Roth and Peranson (1999) report simulations varying the length of doctor rank order lists in markets without couples.

Each line in Figure 2 corresponds a particular number of couples in the market as we vary market size. For a given number of couples, as market size increases, the fraction of simulations for which we find a stable matching increases. For instance, with 20 couples, we always find a stable matching under each simulation once the market size is greater than 1,000. For an economy like the NRMP, where over the last decade, about 600 applicants are couples (corresponding to 1,200 couple members) and the market size averages 26,000, the probability a stable matching is found is roughly 95%. Finally, the red line in Figure 2 shows the fraction of markets with a stable matching when the number of couples is equal to \sqrt{n} , corresponding to an upper bound of the growth rate in Theorem 1. Once the market size is at least 2,000, a stable matching exists in at least 96% of the markets.

Even though Figure 1 suggests that doctor preferences are widely distributed across programs, it is possible that the uniform distribution does not adequately capture the distribution of doctor preferences in real markets. In Figure 3, we use the APPIC dataset to calibrate the distribution for applicant and program preferences. We fit models of doctor and program demand for one another and simulate an economy based on these estimates. (The precise details on how we generate preferences are in Section B.1.) The pattern of how existence changes with market size and the number of couples is similar between markets with the APPIC-calibrated distribution and the uniform distribution. For a given number of couples, the likelihood of finding a stable matching weakly increases with market size for the 100 markets we simulate.

Overall, Figures 2 and 3 suggests that the existence probability increases with market size for a given number of couples. Market sizes and couples fractions like those observed in APPIC or the NRMP display a high likelihood of stable matching existence in our simulation.

6.3 Extensions

This paper focuses on regular sequences of random markets and makes use of each condition in our arguments. A notable implication of the model is that, if the market is large, then it is a high probability event that there are a large number of hospitals with vacant positions, even if there are more applicants than positions (for formal statements, see Proposition 2 in the Appendix). Note that the feature that there are many hospitals with vacant positions is consistent with Fact 5 in Section 2.2.

6.3.1 Distribution of Preferences

Our model assumed certain regularities in the way that random markets grow. In particular, there are some nontrivial restrictions on doctors' preference distributions. Of course, some distributional assumptions are needed in large market analysis: For instance, Immorlica and Mahdian (2005) offer an example where preference distributions violate a regularity assumption and their result fails even without couples.

One of the modeling assumptions in the preceding text is that doctors' preferences are drawn independently from one another in a statistical sense. However, doctors' preferences are not necessarily drawn

from identical distributions and hence can be substantially different from one another. As discussed in Section 4.1, this modeling assumption can capture various possibilities for doctor preferences that may be important empirically. These include doctors from a particular region ranking own-region programs higher than programs outside of their region. Also, the model allows single doctors and couple members to draw their rankings from different distributions.

In fact, statistical independence was introduced for expositional simplicity, but we do not even need this assumption. On the contrary, the result can be generalized to environments with preference correlation and even aggregate shocks in other variables such as supply of doctors and hospitals' job openings. To see this point, consider a model in which there is a state variable σ that is drawn randomly from a certain distribution. For each realization of the state variable σ , there is a sequence of random markets $(\tilde{\Gamma}^1(\sigma), \tilde{\Gamma}^2(\sigma), \dots)$. We assume that $(\tilde{\Gamma}^1(\sigma), \tilde{\Gamma}^2(\sigma), \dots)$ satisfies the regularity condition for each σ , but allow $\tilde{\Gamma}^n(\sigma)$ to be different for different realizations of σ . Doctor preferences are conditionally independent given σ , but this framework allows for doctor preferences to be correlated through their dependence on the common shock represented by the state variable. Clearly, our result generalizes to this model: The probability that a stable matching exists conditional on state variable σ converge to one as the market size approaches infinity for each σ by our previous analysis, and the unconditional probability of existence is merely a weighted average of these conditional probabilities.

This model allows for correlations in doctor preferences as well as other variables that may be useful in applications. For instance, σ could represent the state of the economy in different regions of the country, which then determines how many positions are open in each hospital and which hospitals are popular. In that case, the realization of σ determines the preference distributions from which doctors draw their preference list in such a way that hospitals in regions with positive wage shocks are popular and those in regions with negative wage shocks are unpopular. Similarly, σ could represent changes in funding of medical residency training by Medicare. Medicare is a major source of funding for residency training, which is subject to heated debate and many reform proposals (Rich et al., 2002). Thus the Medicare funding level may be seen as a relevant state variable, which affects parameters such as the number of advertised positions in hospitals and the content of, and hence the popularity of hospital residency programs.

6.3.2 Other Complementarities

As mentioned in Introduction, a couple preference is a particular form of complementarity, and this paper can be put in the context of the research program on the role of complementarities in resource allocation. Examples could include groups of workers who care about externalities and firms that need a team of workers with complementary skills, among others. Given our analysis, a natural question is whether our asymptotic existence result of stable matchings continues to hold in the presence of these and other forms of complementarities.

To study this issue, consider a model of firm-worker matching, where each firm can hire up to κ workers

while each worker can work for at most one firm, where κ is a constant. Suppose that a small fraction of firms draw non-substitutable preferences over workers, while others draw substitutable preferences which satisfy the law of aggregate demand. Moreover, assume regularity of the sequence of random markets similar to Definition 1. The couple matching is the case with $\kappa = 2$, where we relabel firms and workers as doctors and hospitals, respectively (i.e. think of each couple as a firm needing to hire two positions).

Our asymptotic existence result extends to the model described above. To see why, consider a variant of the sequential couples algorithm, which places firms with substitutable preferences first and then places firms with non-substitutable preferences one by one. This algorithm succeeds if no firms apply to workers where other firms with non-substitutable preferences are already tentatively placed. As in the case of couples, as long as the fraction of firms with non-substitutable preferences shrinks sufficiently fast, each firm has a capacity bounded by a constant κ , and each firm finds up to a constant number of workers to be acceptable, then the probability of a success can be shown to converge to one by the same argument as the one for Theorem 1. In fact, Ashlagi et al. (2011) consider a similar model with complementarities and present a convergence result (see their Theorem 5).

7 Conclusion

This paper contributes towards understanding the consequences of the complementarities caused by couples in matching markets, a phenomenon that has grown in importance as dual-career households have become an important part of the labor force. We investigate this issue by studying couples in centralized labor market clearinghouses. Even though a stable matching does not necessarily exist when couples are present, as long as the complementarities caused by couples are small in an appropriate way, our main result is that the market has a stable matching with a high probability. More broadly, our study suggests that not only does large market analysis help understand economies with indivisibilities, but it also generates new kinds of results in markets with complementarities, which have been challenging for existing approaches.

We have complemented our theoretical results with analysis of data from the market for psychologists. The stylized facts from the data motivate some of the modeling assumptions. In every year of the data we are able to find a stable matching with respect to the stated preferences. Our simulations show that a stable matching is more likely to be found in large markets. Since the mechanism we analyze is similar to the actual procedures used in markets such as the NRMP for American medical residents, our results help explain why some mechanisms in practice provide a stable matching with high probability despite the presence of couples.

There are a number of additional questions motivated by this paper. One question is whether stability itself is the reason for the enduring success of the NRMP and post-doctoral psychology market. Field and laboratory evidence suggests that stability is responsible for the persistence of certain centralized clearinghouses (Roth (1991) and Kagel and Roth (2000)). Within the context of couples, however, an alternative may simply be to consider a weaker approximate notion of stability such as a requirement that

the number of blocking coalitions be small. Under our assumptions it is obvious that there always exists a matching that is approximately stable in the above sense.³⁹ What is more remarkable and interesting is that the markets we study have exactly stable matchings, and our analysis provides conditions for this fact.

Another question involves the interpretation of preferences. As we have emphasized, the analysis we undertake is after applicants interview for positions. This is perhaps the major reason why applicants' rank order lists are short in a large market. A richer, but substantially different, model could consider a two-stage game where participants have imperfect information about their preferences and first decide where to interview. This type of analysis could provide a way to endogenize the short rank order lists of applicants. While interesting, we expect this sort of analysis to first focus on the decision problem of where to apply, before adding the complications of how to participate in a matching market with couples.⁴⁰ Loosely speaking, we would expect that in such a model, most applicants would apply to many positions, and many hospitals would have more applicants than they can interview, while some applicants might receive more interview invitations than they can or feel a need to accept, so that a good deal of sorting among doctors and hospitals would take place even before interviews were conducted.

A further topic for research is how *decentralized* markets might be organized to handle couples better. For instance, Niederle and Roth (2009a,b) study how the rules regarding exploding offers influence market outcomes. The issues here would involve the formal and informal rules by which couples search for two positions, and by which offers and responses are made, so as to increase the efficiency of the market in finding matches when some applicants are looking for pairs of positions.

In summary, labor markets in which the pool of applicants includes two-career households have proved challenging to study even as they have become more common, and have demanded adaptation in labor market rules and institutions. While many open questions remain, the results of the present paper suggest that some of the potential problems that couples and market designers face may become more tractable in large markets.

³⁹To see this point, consider a matching that is stable in the sub-market composed of hospitals and single doctors only while keeping all couples unmatched. Clearly the number of coalitions that may block this matching is at most the sum of the lengths of the rank order lists over all couples, which is small in large markets under our maintained assumptions.

⁴⁰Related to this, another issue is to examine whether couples may have an incentive to manipulate by pretending to be singles or vice versa (as in Klaus et al. (2007)), or even whether a dual-career joint location problem encourages or discourages doctors from marrying other doctors (see Hurder (2013)).

Table 1. Summary Statistics for Market for Clinical Psychologists

	Total	Mean	Min	25th	Median	75th	Max
<i>A. Length of Rank-Order List (ROL)</i>							
Single Applicants	3,010	7.6	1.0	4.0	7.1	10.4	73.1
Couples	19	81.2	7.3	29.4	52.3	115.0	249.9
Distinct Programs Ranked		10.2	2.0	6.4	9.9	13.0	20.9
Programs	1,094	16.7	1.0	7.6	14.3	23.9	80.9
<i>B. Program Capacities</i>							
Capacity	2,721	2.5	1.0	1.0	2.0	3.0	21.4
<i>C. Geographic Similarity of Preferences</i>							
Single Applicants							
# Regions Ranked		2.5	1.0	1.0	2.0	3.3	9.4
Couples							
# Regions Ranked		4.0	1.1	2.6	4.0	4.9	6.9
Fraction of ROL where both Members Rank Same Region		72.7%	29.2%	46.6%	77.3%	97.9%	100.0%

Notes: This table reports descriptive information from the Association of Psychology Postdoctoral and Internship Centers match, averaged over 1999-2007. Single applicants' rank order lists consist of a ranking over hospitals, while couples indicate rankings over program pairs. Distinct programs ranked are the set of distinct programs ranked by each couple member. Programs include only those which have positive capacity. There are 11 regions, corresponding to the first digit of US Zipcodes and Canada.

Table 2. Stable Matchings in APPIC Market

Year	Stable	Stable	<i>Sequential Couples Algorithm: Reasons for Failure</i>			
	Matching	Matching	Applicant	Couple Applies	Applicant	Couple Displaces
	Found	Found by SCA	Applies to	to Program	Displaces a	a Couple
	(1)	(2)	Program where	where Couple	Couple Member	Member
	(1)	(2)	Couple Assigned	Assigned	(5)	(6)
1999	YES	YES	100%	0%	0%	0%
2000	YES	NO	100%	100%	100%	74%
2001	YES	YES	100%	0%	0%	0%
2002	YES	YES	100%	95%	0%	0%
2003	YES	YES	0%	0%	0%	0%
2004	YES	YES	100%	100%	52%	52%
2005	YES	YES	0%	0%	0%	0%
2006	YES	NO	100%	52%	100%	0%
2007	YES	NO	100%	52%	100%	0%

Notes: This table reports the source of failure of the sequential couples algorithm. The sequential couples algorithm fails when a doctor applies to a program where a couple member is assigned and is more preferred than the couple member. The table reports the outcomes of 1000 permutations of the ordering of couples.

Figure 1. Popularity of Programs as Top Choice (2003)

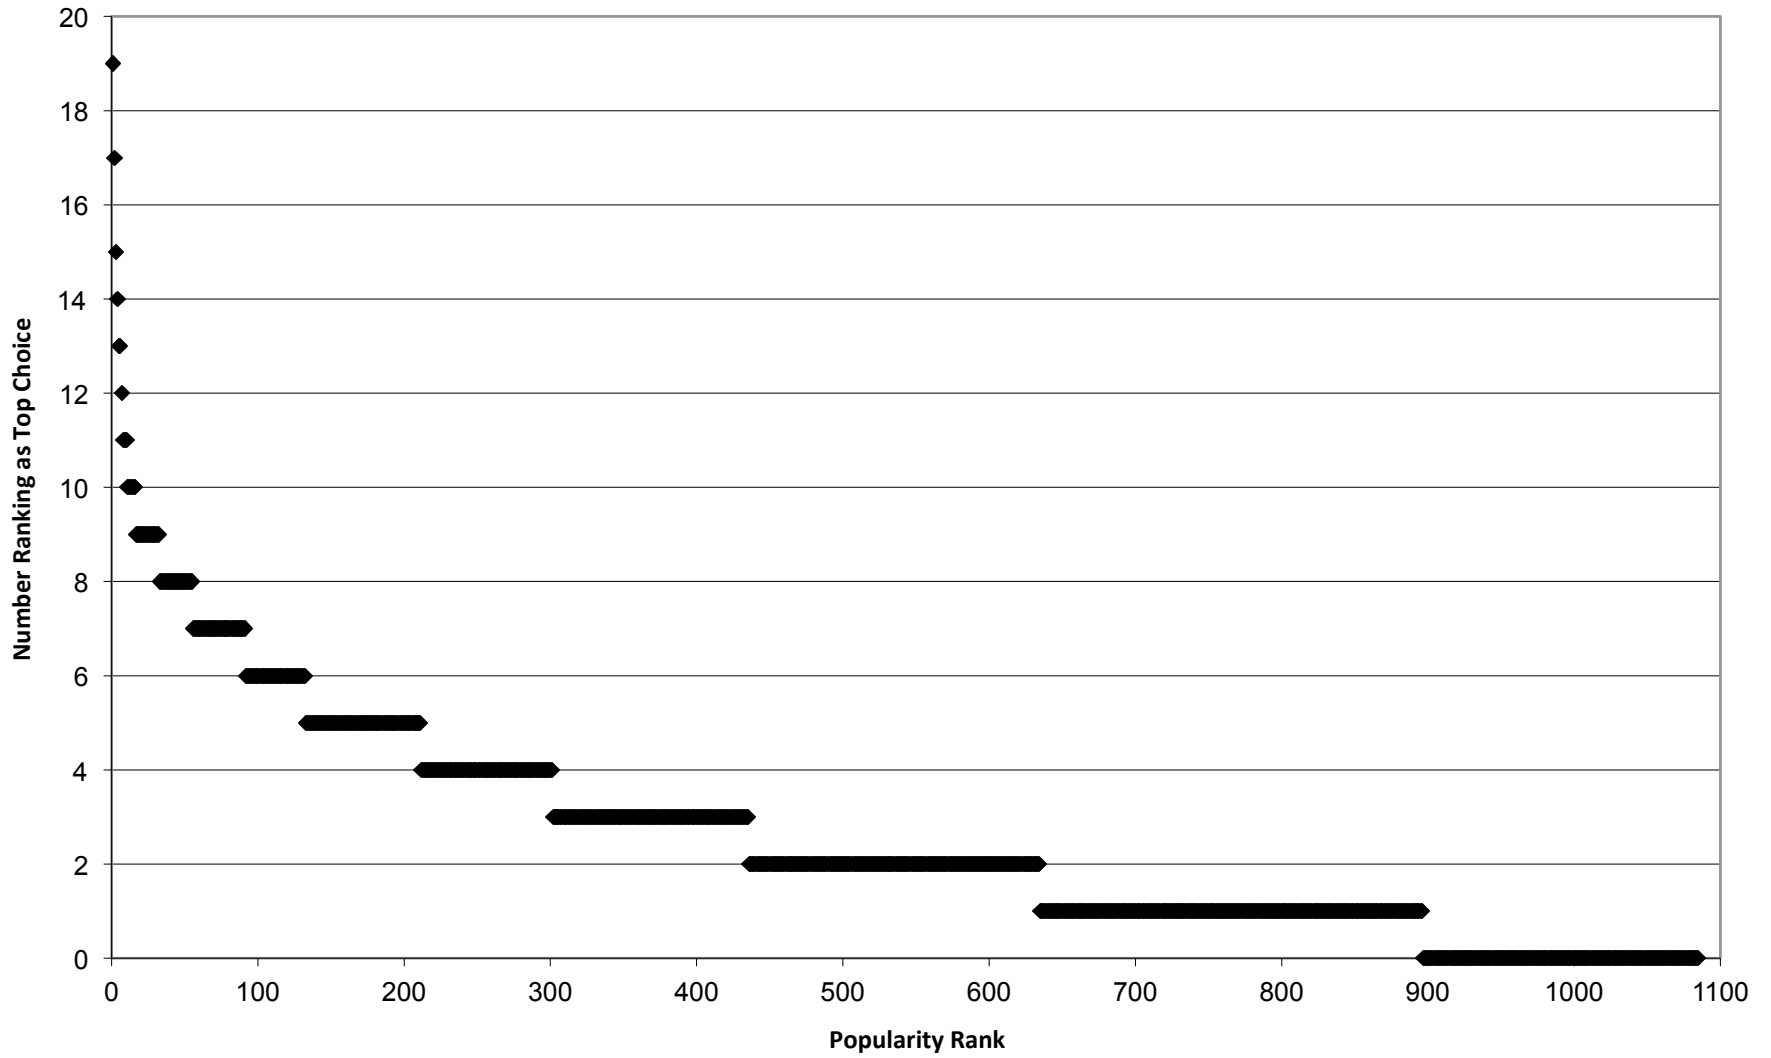


Figure 2. Existence of Stable Matching with Preferences from Uniform Distribution

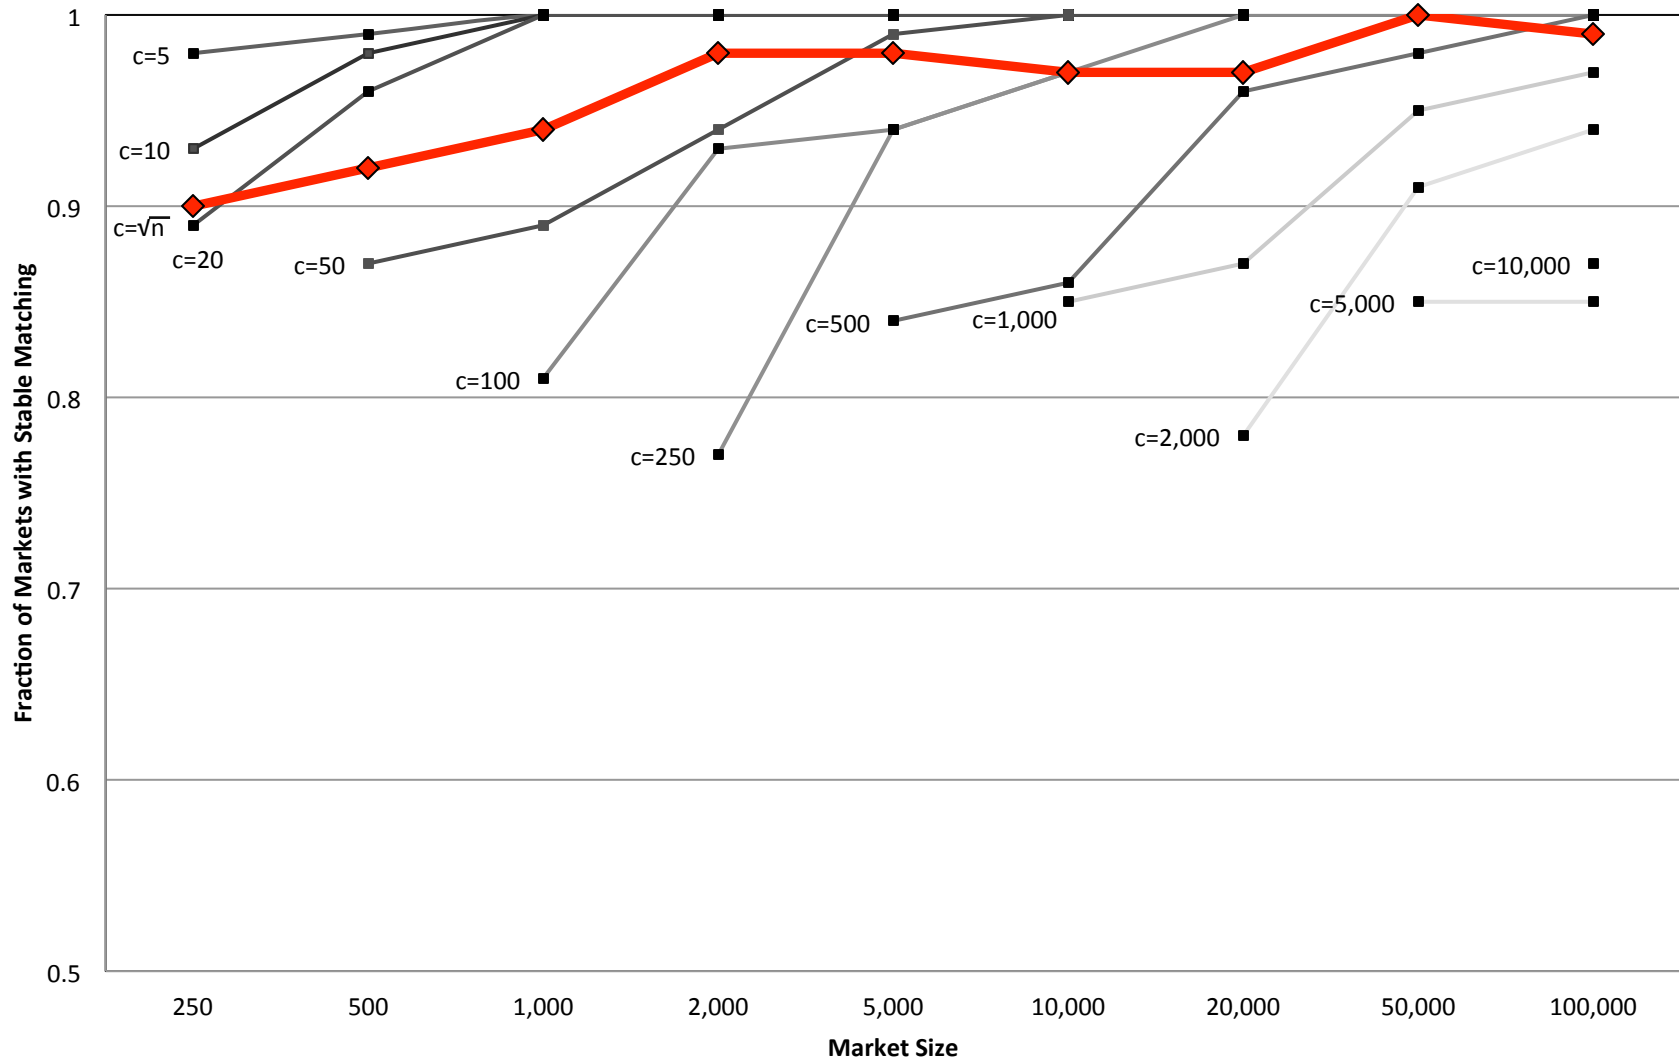
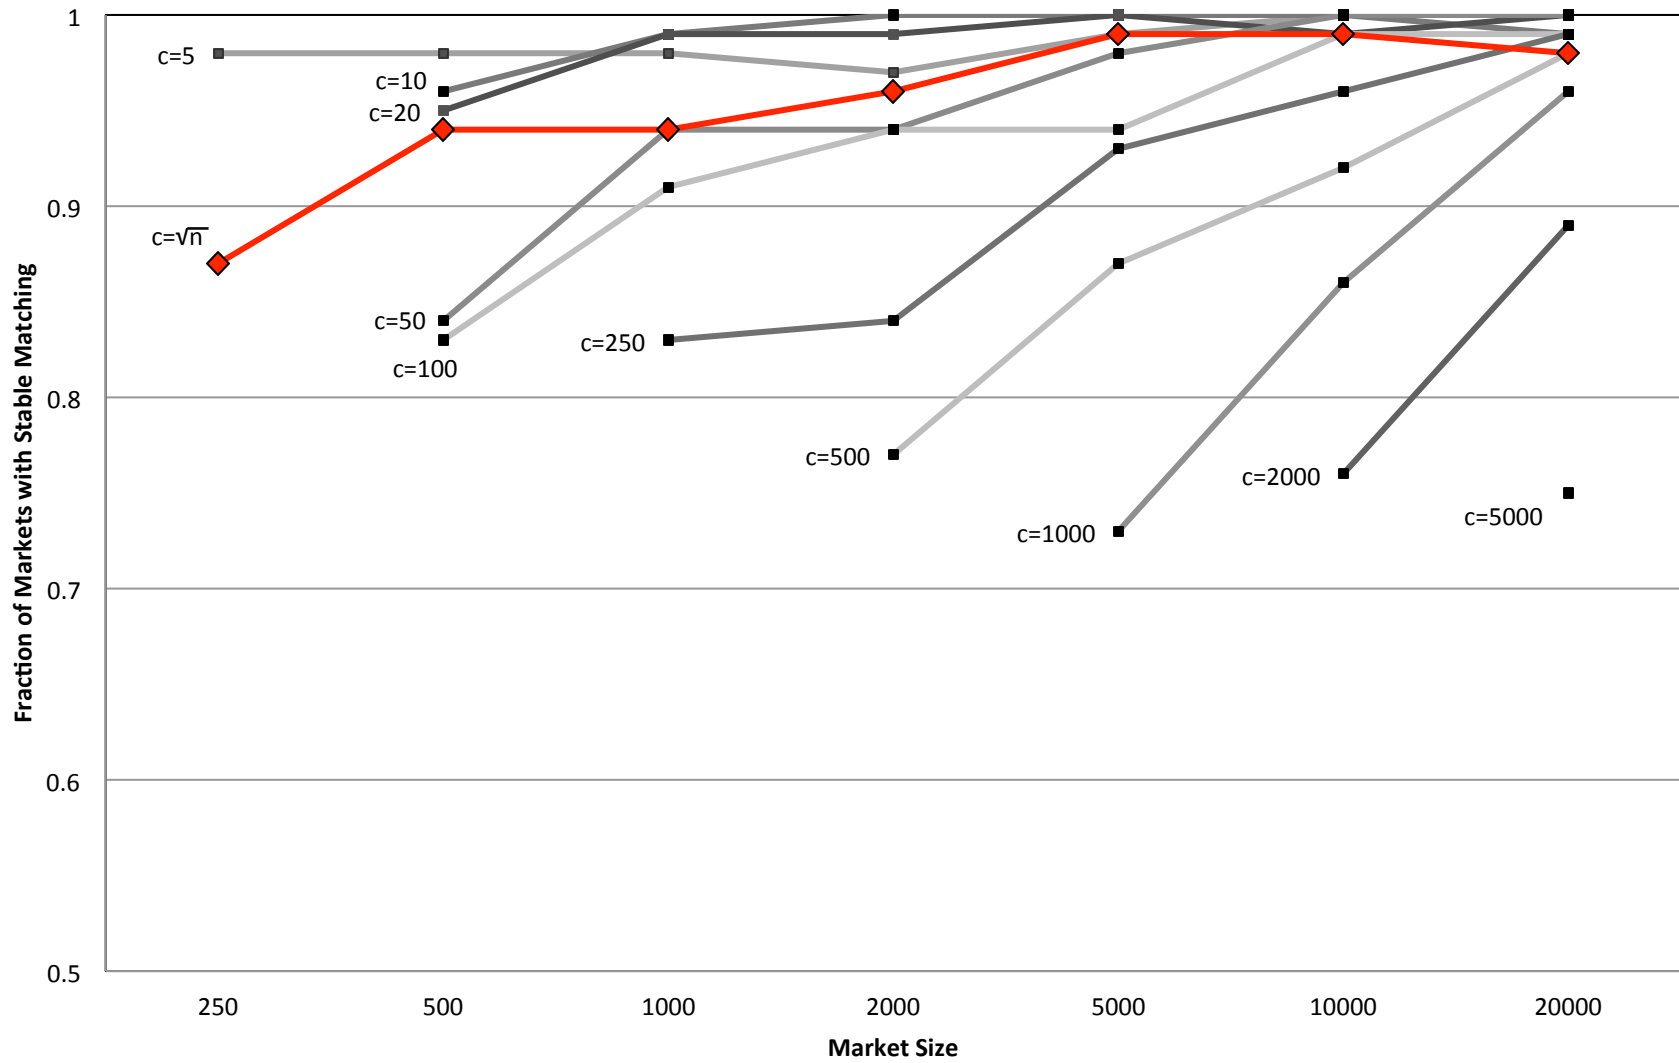


Figure 3. Existence of Stable Matching with Preferences Calibrated from Psychology Market



References

- Abdulkadiroğlu, Atila and Tayfun Sönmez**, “School Choice: A Mechanism Design Approach,” *American Economic Review*, 2003, *93*, 729–747.
- , **Yeon-Koo Che, and Yosuke Yasuda**, “Expanding Choice in School Choice,” 2008. Economic Research Initiatives at Duke Research Paper No. 20.
- Alcalde, José and Salvador Barberà**, “Top Dominance and the Possibility of Strategy-Proof Stable Solutions to Matching Problems,” *Economic Theory*, 1994, *4*, 417–435.
- Ashlagi, Itai, Mark Braverman, and Avinatan Hassidim**, “Stability in large matching markets with complementarities,” 2011. Unpublished mimeo, MIT Sloan School.
- Azevedo, Eduardo M. and Jacob D. Leshno**, “A Supply and Demand Framework for Two-Sided Matching Markets,” 2012. Unpublished manuscript, Wharton.
- Bikhchandani, S. and J.M. Ostroy**, “The package assignment model,” *Journal of Economic Theory*, 2002, *107* (2), 377–406.
- Biró, P. and R.W. Irving**, “Stable Matching with Couples—an Empirical Study,” 2010. Technical Report TR-2010-319, University of Glasgow.
- Budish, Eric**, “The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes,” *Journal of Political Economy*, 2011, *119*(6), 1061–1103.
- Bulow, Jeremy and Jonathan Levin**, “Matching and Price Competition,” *American Economic Review*, 2006, *96*, 652–668.
- Che, Yeon-Koo and Fuhito Kojima**, “Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms,” *Econometrica*, 2010, pp. 1625–1672.
- Costa, Dora and Matthew Kahn**, “Power Couples: Changes in the Locational Choice of the College Educated, 1940-1990,” *Quarterly Journal of Economics*, 2000, *115*, 1287–1315.
- Dean, Brian C., Michael X. Goemans, and N. Immorlica**, “The Unsplittable Stable Marriage Problem,” *Proceedings of the 4th IFIP International Conference on Theoretical Computer Science*, 2006, pp. 65–76.
- Dutta, B. and J. Massó**, “Stability of Matchings when Individuals have Preferences over Colleagues,” *Journal of Economic Theory*, 1997, pp. 464–475.
- Echenique, Federico and Jorge Oviedo**, “A Theory of Stability in Many-to-Many Matching,” *Theoretical Economics*, 2006, *1*, 233–273.

- Ehlers, Lars**, “In Search of Advice for Participants in Matching Markets which use the Deferred Acceptance Algorithm,” *Games and Economic Behavior*, 2004, *48*, 249–270.
- , “Truncation Strategies in Matching Markets,” *Mathematics of Operations Research*, 2008, *33*, 327–335.
- **and Aytek Erdil**, “Efficient assignment respecting priorities,” *Journal of Economic Theory*, 2010, *145* (3), 1269–1282.
- Erdil, Aytek and Haluk Ergin**, “What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice,” *American Economic Review*, 2008, *98*, 669–689.
- Ergin, Haluk**, “Efficient Resource Allocation on the Basis of Priorities,” *Econometrica*, 2002, *70*, 2489–2498.
- Gale, David and Lloyd S. Shapley**, “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 1962, *69*, 9–15.
- Gresik, Thomas and Mark Satterthwaite**, “The Rate at Which a Simple Market Converges to Efficiency as the Number of Traders Increases,” *Journal of Economic Theory*, 1989, *48*, 304–332.
- Gul, Faruk and Ennio Stacchetti**, “Walrasian Equilibrium with Gross Substitutes,” *Journal of Economic Theory*, 1999, *87*, 95–124.
- **and** – , “The English auction with differentiated commodities,” *Journal of Economic Theory*, 2000, *92* (1), 66–95.
- Haeringer, Guillaume and Flip Klijn**, “Constrained School Choice,” *Journal of Economic Theory*, 2009, *144*, 1921–1947.
- Hatfield, John William and Fuhito Kojima**, “Matching with Contracts: Comment,” *American Economic Review*, 2008, *98*, 1189–1194.
- **and** – , “Substitutes and Stability for Matching with Contracts,” *Journal of Economic Theory*, 2010, *145*, 1703–1723.
- **and Paul Milgrom**, “Matching with Contracts,” *American Economic Review*, 2005, *95*, 913–935.
- **and Scott Duke Kominers**, “Many-to-Many Matching with Contracts,” 2009. mimeo.
- Helppie, Brooke and Marta Murray-Close**, “Moving Out or Moving Up? New Economists Sacrifice Job Opportunities for Proximity to Significant Others – and Vice Versa,” 2010. working paper, University of Michigan.

- Hurder, Stephanie**, “An Integrated Model of Occupation Choice, Spouse Choice, and Family Labor Supply,” 2013. Harvard University, Department of Economics, Available at: <http://www.people.fas.harvard.edu/~hurder/papers/SHurderJMP.pdf>.
- Immorlica, Nicole and Mohammad Mahdian**, “Marriage, Honesty, and Stability,” *SODA*, 2005, pp. 53–62.
- Jackson, Matthew O. and Alejandro M. Manelli**, “Approximately competitive equilibria in large finite economies,” *Journal of Economic Theory*, 1997, pp. 354–376.
- Jonas, Harry S. and Sylvia I. Etzel**, “Undergraduate Medical Education,” *JAMA*, 1998, *260*, 1063–1071.
- Kagel, John and Alvin E. Roth**, “The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Field Experiment,” *Quarterly Journal of Economics*, 2000, *115*(1), 201–235.
- Keilin, W. Gregory**, “Internship selection 30 years later: An overview of the APPIC matching program,” *Professional Psychology: Research and Practice*, 1998, *29*, 599–603.
- , “Internship Selection in 1999: Was the Association of Psychology Postdoctoral and Internship Centers’ Match a Success?,” *Professional Psychology: Research and Practice*, 2000, *31*, 281–287.
- Kesten, Onur**, “On two competing mechanisms for priority-based allocation problems,” *Journal of Economic Theory*, 2006, *127*, 155–171.
- , “On Two Kinds of Manipulation for School Choice Problems,” 2006. Unpublished mimeo.
- , “School Choice with Consent,” *Quarterly Journal of Economics*, 2010, *125*(3), 1297–1348.
- Klaus, Bettina and Flip Klijn**, “Stable Matchings and Preferences of Couples,” *Journal of Economic Theory*, 2005, *121*, 75–106.
- , – , and **Jordi Masso**, “Some Things Couples Always wanted to know about Stable Matchings (but were afraid to ask),” *Review of Economic Design*, 2007, *11*, 175–184.
- , – , and **Toshifumi Nakamura**, “Corrigendum: Stable Matchings and Preferences of Couples,” *Journal of Economic Theory*, 2009, *144*, 2227–2233.
- Knuth, Donald E., Rajeev Motwani, and Boris Pittel**, “Stable Husbands,” *Random Structures and Algorithms*, 1990, *1*, 1–14.
- Kojima, Fuhito**, “When can Manipulations be Avoided in Two-Sided Matching Markets? Maximal Domain Results,” *Contributions to Theoretical Economics*, 2007, *7*, Article 32.

- **and Mihai Manea**, “Incentives in the Probabilistic Serial Mechanism,” *Journal of Economic Theory*, 2010, *145*, 106–123.
- **and Parag A. Pathak**, “Incentives and Stability in Large Two-Sided Matching Markets,” *American Economic Review*, 2009, *99*, 608–627.
- , – , **and Alvin E. Roth**, “Matching with Couples: Stability and Incentives in Large Markets,” 2010. NBER Working Paper 16028.
- Konishi, Hideo and M. Utku Ünver**, “Credible Group Stability in Many-to-Many Matching Problems,” *Journal of Economic Theory*, 2006, *129*, 57–80.
- **and** – , “Games of Capacity Manipulation in the Hospital-Intern Market,” *Social Choice and Welfare*, 2006, *27*, 3–24.
- Lee, SangMok**, “Incentive Compatibility of Large Centralized Matching Markets,” 2013. Unpublished manuscript, University of Pennsylvania.
- Liu, Qingmin and Marek Pycia**, “Ordinal Efficiency, Fairness, and Incentives in Large Markets,” 2011. UCLA Working paper.
- Manea, M.**, “Asymptotic ordinal inefficiency of random serial dictatorship,” *Theoretical Economics*, 2009, *4* (2), 165–197.
- McDermid, Eric J. and David Manlove**, “Keeping partners together: algorithmic results for the hospitals/resident problem with couples,” *Journal of Combinatorial Optimization*, 2009, *19*, 279–303.
- Milgrom, Paul R.**, *Putting Auction Theory to Work*, Cambridge: Cambridge University Press, 2004.
- Motwani, R. and P. Raghavan**, “Randomized algorithms,” *ACM Computing Surveys (CSUR)*, 1996, *28* (1), 37.
- Niederle, Muriel and Alvin E. Roth**, “The Effects of a Central Clearinghouse on Job placement, Wages, and Hiring Practices,” in David Autor, ed., *Labor Market Intermediation*, Chicago: The University of Chicago Press, 2009a.
- **and** – , “Market Culture: How Rules Governing Exploding Offers Affect Market Performance,” *American Economic Journal: Microeconomics*, 2009b, *1*, 199–219.
- NRMP**, “Results and Data: 2009 Main Residency Match,” 2009. Accessed at: <http://www.nrmp.org/data/resultsanddata2009.pdf> on March 15, 2010.
- Ostrovsky, Michael**, “Stability in supply chain networks,” *American Economic Review*, 2008, *98* (3), 897–923.

- Pathak, Parag A. and Tayfun Sönmez**, “Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism,” *American Economic Review*, 2008, *98*(4), 1636–1652.
- Pycia, Marek**, “Stability and Preference Alignment in Matching and Coalition Formation,” *Econometrica*, 2012, *80*(1), 323–362.
- Rich, E.C., M. Liebow, M. Srinivasan, D. Parish, J.O. Wolliscroft, O. Fein, and R. Blaser**, “Medicare financing of graduate medical education,” *Journal of General Internal Medicine*, 2002, *17*(4), 283–292.
- Roberts, Donald John and Andrew Postlewaite**, “The Incentives for Price-Taking Behavior in Large Exchange Economies,” *Econometrica*, 1976, *44*, 115–127.
- Ronn, E.**, “NP-complete stable matching problems,” *Journal of Algorithms*, 1990, *11*(2), 285–304.
- Roth, Alvin E.**, “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory,” *Journal of Political Economy*, 1984, *92*, 991–1016.
- , “The College Admission Problem is not Equivalent to the Marriage Problem,” *Journal of Economic Theory*, 1985, *36*, 277–288.
- , “A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the U.K.,” *American Economic Review*, 1991, *81*, 415–440.
- , “The Origins, History, and Design of the Resident Match,” *Journal of the American Medical Association*, 2003, *289*, 909–912.
- , “Deferred Acceptance Algorithms: History, Theory, Practice and Open Questions,” 2007. forthcoming, *International Journal of Game Theory*.
- , “What Have We Learned from Market Design?,” in Nir Vulkan, Alvin E. Roth, and Zvika Neeman, eds., *The Handbook of Market Design*, Oxford University Press, 2013.
- **and Elliott Peranson**, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 1999, *89*, 748–780.
- **and John H. Vande Vate**, “Random Paths to Stability in Two-Sided Matching,” *Econometrica*, 1990, *58*, 1475–1480.
- **and Marilda A. O. Sotomayor**, *Two-sided Matching: a Study in Game-theoretic Modeling and Analysis*, Cambridge: Econometric Society monographs, 1990.
- **and Uriel Rothblum**, “Truncation Strategies in Matching Markets: In Search of Advice for Participants,” *Econometrica*, 1999, *67*, 21–43.

- **and Xiaolin Xing**, “Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions,” *American Economic Review*, 1994, *84*, 992–1044.
- **and –**, “Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists,” *Journal of Political Economy*, 1997, *105*, 284–329.
- Rustichini, Aldo, Mark Satterthwaite, and Steven Williams**, “Convergence to Efficiency in a Simple Market with Incomplete Information,” *Econometrica*, 1994, *62*, 1041–1064.
- Sönmez, Tayfun and M. Utku Ünver**, “Matching, Allocation, and Exchange of Discrete Resources,” 2009. forthcoming, Handbook of Social Economics, eds. Jess Benhabib, Alberto Bisin, and Matthew Jackson, Elsevier.
- **and –**, “Course bidding at business schools,” *International Economic Review*, 2010, *51(1)*, 425–445.
- Sotomayor, Marilda A. O.**, “Three Remarks on the Many-to-Many Stable Matching Problem,” *Mathematical Social Sciences*, 1999, *38*, 55–70.
- , “Implementation in the many-to-many matching market,” *Games and Economic Behavior*, 2004, *46* (1), 199–212.
- Sun, Ning and Zaifu Yang**, “Equilibria and Indivisibilities: Gross Substitutes and Complements,” *Econometrica*, 2006, *74*, 1385–1402.
- **and –**, “Double-Track Adjustment Process for Discrete Markets With Substitutes and Complements,” *Econometrica*, 2009, *77*, 933–952.