MATCHING WITH COUPLES:
A MULTIDISCIPLINARY SURVEY*

PÉTER BIRÓ†

*Institute of Economics, Hungarian Academy of Sciences, Budaörsi út 45
Budapest, H-1112, Hungary
birop@econ.core.hu
http://econ.core.hu/english/inst/biro.html*

FLIP KLIJN‡

*Institute for Economic Analysis (CSIC), Campus UAB
08193 Bellaterra (Barcelona), Spain
flip.klijn@iae.csic.es
http://sites.google.com/site/flipklijn*

This survey deals with two-sided matching markets where one set of agents (workers/residents) has to be matched with another set of agents (firms/hospitals). We first give a short overview of a selection of classical results. Then, we review recent contributions to a complex and representative case of matching with complementarities, namely matching markets with couples. We discuss contributions from computer scientists, economists, and game theorists.

Keywords: matching; couples; stability; computational complexity; incentive compatibility; restricted domains; large markets.

2010 Mathematics Subject Classification: 03D15, 68-02, 68Q17, 68Q25, 68R10, 68U10, 90B10, 90C27, 90C59, 91-02, 91-08, 91A06, 91A80, 91B10, 91B68.

1. Introduction

In a general matching problem, a set of agents has to be matched, either to some other set of agents (two-sided matching), or to itself (one-sided matching). Gale and Shapley [1962] wrote a now classical paper in which they considered the two basic models of one-to-one matching ("marriage markets") and many-to-one matching (labor markets). Their paper generated a vast literature, and especially in the last two decades the number of contributions and researchers in the field has grown very fast: hundreds, if not thousands, of theory papers, empirical studies, laboratory

experiments, and three text books[1] studied matching problems. Apart from the text books, some surveys[2] have reviewed parts of the matching literature. Last but not least, the number of applications (labor markets, school choice, kidney exchange, etc.) has made the field of interest to market designers, policy makers, and the public in general.[3]

In this survey, we focus on a particular but important branch of many-to-one two-sided matching, motivated by a feature of many entry-level labor markets. This feature is the presence of couples whose members each look for a job in the same market. A couple faces a difficult coordination problem since it has to find a "good" pair of jobs. Hence, the agents' preferences contain complementarities. Most matching markets studied in the literature consider agents' preferences to be free of complementarities. It turns out that introducing complementarities, such as the presence of married couples, changes a great deal the properties of the market.

The extent to which the presence of couples causes complications can possibly be most easily illustrated in the case of centralized labor markets. In such centralized matching schemes, the participants submit their preferences to the market organizer and an automated mechanism produces a matching. Ideally, the mechanism is incentive compatible and its outcomes satisfy certain desirable properties related to welfare, optimality, etc. that are studied in the economics and game theory literature. When couples are present, the existence of "desirable" matchings can no longer be guaranteed. The automated mechanism of a centralized scheme is usually realized by a computer program, which should be able to compute an outcome efficiently for any possible input of preference lists. The computational tractability of mechanism design problems is intensively studied in the computer science literature. When couples are present, and even if "desirable" matchings do exist, the problem of finding such a matching is no longer tractable. Therefore, the study of matching markets with couples is profoundly different from matching markets without couples. It also calls for an interdisciplinary approach to tackle the hurdles just described. Obviously, improving our knowledge regarding this complementarity is a first step to dealing with other complementarities as well.

In Section 2 we first give a short overview of the most important results in matching without couples. This makes the survey self-contained and facilitates an easy comparison with Section 3, where we review recent contributions to the study of matching markets with couples. Section 4 concludes. Finally, two remarks are

---

[1]The first book was written by Knuth [1976] with a focus on computational aspects. The second book by Gusfield and Irving [1989] studied stable matching problems from a graph and algorithm theoretical point of view. The third book was written by Roth and Sotomayor [1990] and mainly describes structural and game-theoretical findings.
[2]The survey by Roth [2008] focused on Gale and Shapley's (1962) celebrated deferred acceptance algorithm. The survey by Klijn [2008] focused on recent developments in school choice, which has turned out to be a fruitful application of matching theory. Sönmez and Ünver [2011] and Abdulkadiroğlu and Sönmez [2010] dealt with several models of both one-sided and two-sided matching. Finally, Kojima and Troyan [2011] focused on market design issues in matching markets.
[3]We refer to Al Roth's website and market design blog for further discussion.

in order. First, in view of the main applications, we have opted to use mainly the terminology of students/residents and hospitals. Second, we have purposely omitted most formal definitions: they would require substantial additional space without clear benefits for most non-specialized readers.

2. Matching Without Couples

There are two finite and disjoint sets of agents: a set S of single students (or residents) and a set H of hospitals. For each hospital h, there is a fixed quota q_h that represents the number of positions it offers.[4]

Each student has a complete, transitive, and strict preference relation over the hospitals and the prospect of being unmatched. Hence, student s's preferences can be represented by a strict ordering P_s of the elements in $H \cup \{s\}$, where the element s represents being unmatched. Hospital h is acceptable to student s if s prefers h to s.

Each hospital has a complete, transitive, and strict preference relation over the individual students and the prospect of getting no residents. Hence, hospital h's preferences over individual students can be represented by a strict ordering P_h of the elements in $S \cup \{\emptyset\}$. Student s is acceptable to hospital h if h prefers s to $\emptyset$. A set of students $S' \subseteq S$ is feasible for hospital h if $|S'| \leq q_h$. Each hospital h has a complete and transitive preference relation over feasible sets of students, which can be represented by a weak ordering P_h^* of the elements in $\mathcal{P}(S, q_h) := \{S' \subseteq S : |S'| \leq q_h\}$. A widely used assumption in the literature is that P_h^* is a responsive extension[5] of P_h. This connects the preferences over sets of students with the preferences over individual students in the following way. Given a set of students, an improvement is obtained by (i) replacing a student by a better student, or (ii) filling an empty position with an acceptable student. (In other words, each hospital ranks a student independently of his colleagues.) By responsiveness, P_h^* coincides with P_h on the set of individual students. Note also that P_h allows for multiple responsive extensions.

A (matching) market is a triple (S, H, P), where $P = (P_S, P_H^*)$. A matching for market (S, H, P) is a function μ on the set $S \cup H$ that specifies which students are matched to which hospitals. Formally,

- for all $s \in S$, either $\mu(s) \in H$ or $\mu(s) = s$;
- for all $h \in H$, $\mu(h) \in \mathcal{P}(S, q_h)$;
- for all $s \in S$ and $h \in H$, $\mu(s) = h$ if and only if $s \in \mu(h)$.

A matching μ is individually rational if any student and any hospital that are matched to one another are mutually acceptable. If a student s and a hospital h are not matched to one another in a matching μ, but the student would prefer to be matched to the hospital and the hospital would prefer to either add the student

---

[4]The marriage model is the special case of one-to-one matching, where for all $h \in H$, $q_h = 1$.
[5]See Roth [1985] and Roth and Sotomayor [1990] for a discussion of this assumption.

or replace another student by student s, then we would expect the pair (s, h) to block the matching, i.e., carry out the mutually beneficial adjustment. A matching is stable if it is individually rational and there are no blocking pairs.[6] Stability is a key property of matchings. Apart from its theoretical appeal, empirical studies have found that in many centralized labor markets, clearinghouses are most often successful if they produce stable matchings.[7] Since stability does not depend on the particular responsive extensions of the hospitals' preferences over individual students, a matching market is henceforth a triple (S, H, P), or simply P, where $P = (P_S, P_H)$. In view of stability, we can (and often do) omit, without loss of generality, the unacceptable hospitals and sets of students from the agents' preference lists.

2.1. *Existence of Stable Matchings and Optimality*

The famous deferred acceptance (DA) algorithm (Gale and Shapley [1962]) yields a stable matching.[8] Let $P = (P_S, P_H)$ be a preference profile (of preferences over individual agents). The student-proposing DA algorithm applied to P, denoted by $\mathrm{DA}(P)$ for short, finds a matching through the following steps.[9]

STEP 1: Each student s proposes to the hospital that is ranked first in P_s (if there is no such hospital then s remains single). Each hospital h considers its proposers and tentatively assigns its q_h positions to these students one at a time following the preferences P_h. All other proposers are rejected.

STEP k, $k \geq 2$: Each student s that is rejected in Step $k - 1$ proposes to the next hospital in his list P_s (if there is no such hospital then s remains single). Each hospital h considers the students it has been holding together with its new proposers and tentatively assigns its q_h positions to these students one at a time following the preferences P_h. All other proposers are rejected.

The algorithm stops when no student is rejected. Then, all tentative matches become final. With some abuse of notation, let $\mu(P)$ denote the matching. Gale and Shapley [1962] proved that matching $\mu(P)$ is stable with respect to preference profile P. Moreover, they also showed (one-side) optimality: $\mu(P)$ is the best (worst) stable matching for the students (hospitals) with respect to preference profile P.

Taking the requirement of stability as granted, an important question is whether the choice of a particular stable mechanism affects the numerical distribution of students; and if not, whether it affects the composition of the hospitals that do not fill all their vacant positions. For instance, Sudarshan and Zisook [1981] observed that

---

[6]It is not necessary to consider any other possible blocking coalition since the existence of such a blocking coalition implies the existence of a blocking pair (see for instance, Roth and Sotomayor [1990], Lemma 5.5). Also, the set of stable matchings coincides with the core defined by weak domination (see for instance, Roth and Sotomayor [1990], Proposition 5.36).

[7]See, for instance, Roth [1984b], Roth [1990], Roth [1991], and Roth and Xing [1994].

[8]We refer to Roth [2008] for a history on the deferred acceptance algorithm.

[9]Note that the DA algorithm does not depend on the particular responsive extensions.

in the US market of medical graduates the National Resident Matching Program (NRMP) fails to fill the posts of many rural hospitals. Roth [1984a] and Roth [1986] showed that the problem of the rural hospitals cannot be attributed to the NRMP's particular stable mechanism, since the number of assignees is the same for every hospital in each stable matching. Moreover, a hospital that does not fill its quota in all stable matchings gets the same assignees in each stable matching.

Many matching markets, and in particular job markets, work in a decentralized way. This means that the agents contact each other randomly and they form new pairs (and possibly break old ones) whenever this is in their interest (it is assumed that the market is flexible; the costs of (re-)contracting can be neglected). Knuth [1976] posed the question whether the above stochastic process converges to a stable matching. Roth and Vande Vate [1990] showed that there always exists a path to stability in one-to-one matching markets, implying that the stochastic process converges to a stable matching with probability one. This result has been extended for many-to-one markets by Kojima and Ünver [2006]. However, real markets are not completely flexible, and even if they were, the time[10] and efforts needed to reach a stable solution in a medium-size market are significant. Therefore, in many large markets, especially when the matching is due to be created in a short time interval (e.g., for resident allocation and student admissions), centralized matching schemes are used.

2.2. *Efficiency and Incentive Compatibility in Centralized Markets*

Two important desirable properties for centralized matching schemes are computational tractability and incentive compatibility. They are compatible with stability if the students are the only strategic agents: the DA algorithm induces a mechanism that satisfies the first and the restricted second additional requirements.

The DA algorithm can be implemented in linear time (see, e.g., Knuth [1976]). This means that the running time of the algorithm is proportional to the number of applications in the scheme. For example, the algorithm used in the Hungarian higher education matching scheme calculates a student-optimal stable matching[11] in 10 seconds for about 150,000 applicants. The linear running time suggests that the same method could be efficiently used in any other country (including higher education admissions in China, where a scoring system similar to that of Hungary is used). Note that even a quadratic time algorithm would very likely be too slow for such large scale applications.

The other important property of the DA-algorithm is incentive compatibility. Roth [1985] proved that under the direct-revelation mechanism induced by μ it is a weakly dominant strategy for the students to reveal their true preferences. So, no

---

[10] Ackermann *et al.* [2008] showed that the expected time of convergence to a stable matching may be exponential in two-sided matching markets.
[11] This application has some special features (see Biró [2008] and Biró, Fleiner, Irving and Manlove [2010]), but the automated mechanism of the scheme is based on the DA algorithm.

student has incentives to misrepresent his preferences.[12] However, it is not possible to ensure incentive compatibility for both sides of (one-to-one) markets as Roth [1982] showed that any stable mechanism is manipulable via preference lists: there is always some participant for whom reporting a true preference list is not a best response to the other participants' reported preferences. Furthermore, in many-to-one markets we cannot even obtain incentive compatibility for the side of the hospitals. Even the hospital-proposing variant of the DA algorithm (which induces a hospital-optimal stable matching) is not immune to manipulation from hospitals (see Roth [1985]).

Sönmez [1997] and Sönmez [1999] showed that if hospitals have more than one slot then there are additional strategic concerns. First, any stable mechanism is manipulable via capacities, i.e., hospitals may sometimes benefit by underreporting their quotas. Second, any stable mechanism is manipulable via prearranged matches, i.e., a hospital and a student may benefit by agreeing to match before receiving their allocations from the centralized matching mechanism.

2.3. *Large Centralized Markets*

Roth and Peranson [1999] conducted simulations based on data from the NRMP. They found suggestive evidence that in large markets the set of stable matchings is likely to be very small and that stable matching mechanisms are difficult to manipulate. A key assumption in their analysis is that the length of students' preference lists is bounded as market size grows. Motivated by the findings of Roth and Peranson [1999], Immorlica and Mahdian [2005] and Kojima and Pathak [2009] analyzed the scope for manipulation in one-to-one and many-to-one matching markets, respectively, under the student-optimal stable mechanism when the number of participants is large. Under some regularity conditions (including the bounded preference lists of the students), they showed that the fraction of participants with incentives to misrepresent their preferences when the others are truthful approaches zero as the market becomes large. Moreover, if the market is sufficiently "thick," truthful reporting by every participant is an approximate equilibrium.

3. Matching With Couples

In this section we consider matching markets where couples are present. We first describe some of the most important real-life matching schemes involving couples. Then, we explain through several simple situations why and how the stability concept can be generalized in different ways to the extended model with couples. However, no matter the precise definition, the set of stable matchings can be empty and no longer satisfies several desirable structural properties. Also, the problem of finding a stable matching is computationally hard. Next, we compare some heuristics

---

[12]In other words, the mechanism is strategy-proof for the students.

that are used in practice. We also discuss the papers that have considered preference domain restrictions (but nonetheless have obtained several negative results regarding structural properties of the set of stable matchings). Finally, we describe some very recent results for large markets with couples.

3.1. *Applications*

In the National Resident Matching Program (NRMP) couples can submit joint preference lists over pairs of positions since 1983. Before that, they could only submit individual preferences for each spouse and assign priority to one spouse. But even after this change, the participation of married couples in the program was very low (see Roth [1984b]). However, in 1995–1998 the program was successfully redesigned by Roth and Peranson [1999]. One of the important issues that triggered the redesign was the presence of couples. Since the redesign of the NRMP, a number of similar applications in the US have adopted the method (see Roth [2008] for a summary).

In Scotland, the Scottish Foundation Allocation Scheme (SFAS) has been responsible for the allocation of residents to hospitals since 2000. Since 2010 a pair of residents can indicate that they wish to be treated as a couple. A couple submits individual preference lists and then a joint preference list is created automatically from them in a very specific way, keeping only those pairs of positions that are located in hospitals geographically close to each other. At the same time, all residents are ranked at each hospital according to a common scoring method. Hence, there is a so-called "master list" of resident scores (see Biró, Irving and Schlotter [2011] for more details).

The US Navy assignment process involves hundreds of thousands of sailors to be assigned, commands seeking sailors and detailers who advise both sailors and commands. According to recent studies (see Short [2000], Robards [2001], and Yang, Giampapa and Sycaratech [2003]), the three most important requirements for an optimal assignment of sailors to billets are the following. The matching should be stable; it should accommodate the wishes of married couples; and it should ensure that all sailors are assigned to a position and that certain critical positions are filled. Although there is no publicly available description of the algorithm used in this (or any similar) program, it is clear that the presence of couples plays an important role.[13]

In the Hungarian higher education matching scheme, students apply for studies via a centralized program since 1985. The matching mechanism is based on a sta-

---

[13]Note that if we consider the third requirement primordial and take only one of the other two requirements, then we already end up with an NP-hard problem. If we omit the stability requirement and aim to find a maximum size matching with couples then the problem is NP-hard (see Biró and McDermid [2010] and Marx and Schlotter [2011]). If we neglect the wishes of couples and try to find a maximum size matching with a minimum number of blocking pairs, then the problem becomes again NP-hard (see Biró, Manlove and Mittal [2010]).

ble score-limit algorithm, which is a generalization of the DA algorithm (see Biró [2008]). Apart from the issue of ties (i.e., when there are students with identical scores), there are at least two other special features in the scheme. Universities and colleges not only have different kinds of common (upper) quotas, but they can also set lower quotas for each study (see Biró, Fleiner, Irving and Manlove [2010]). Finally, students can apply for pairs of MSc studies, which are special studies for future teachers. For instance, in 2010, 5,578 students applied for teachers' studies and 2,091 of them applied for pairs of studies. The matching problem is essentially the same as the one in matching markets with couples, since students applying to pairs of studies act like couples in resident allocation programs. The proportion of paired applications is very high since they are crucial in pursuing a successful teaching career.

3.2. *Stability: Definitions*

If each hospital has only one position then the stable matching problem in markets with couples can be formulated in terms of a hedonic coalition formation game[14] as follows. The single applicants, the couples, and the hospitals are the players and each coalition corresponds to a possible contract between a single applicant and a hospital or between a couple and a pair of hospitals. The preferences of the hospitals over the possible contracts are derived from their preferences over the individual candidates who would be assigned to them under the contracts. A stable matching in a market with couples is equivalent to a core element in the above defined NTU-game. More precisely, an individually rational matching is not stable if there is a blocking pair (consisting of a single applicant and a hospital) or a blocking coalition (consisting of a couple and two hospitals) whose members would all be better of by signing a new contract and possibly breaking existing ones.

If hospitals have multiple positions but no couple applies for a pair of positions at the same hospital, then the extension of the stability definition is straightforward.[15] However, if hospitals have multiple positions and couples are allowed to apply to a pair of positions at the same hospital, then different definitions of stability are possible.[16] Assuming hospitals have responsive preferences we know how they would decide between two single candidates, but it is not obvious how they would rank a single applicant with respect to a couple, or how they would rank two couples.

---

[14]See Bogomolnaia and Jackson [2002] for a definition of hedonic coalition formation games. To be precise, we can model a market with couples as a coalition formation game only if there is no couple that applies to a pair of hospitals in two ways (i.e., (h, h') and (h', h)). In the general case, the matching problem can be formulated in terms of an NTU-game, where the core corresponds to the set of stable matchings.

[15]Here again, the set of stable matchings coincides with the core defined by weak domination.

[16]The complexities with multiple positions have not been considered until lately. The only papers that mention this issue are Kojima, Pathak and Roth [2010], McDermid and Manlove [2010], Marx and Schlotter [2011], and Biró, Irving and Schlotter [2011]. Both the NRMP and the SFAS allow couples to apply to a pair of positions at the same hospital.

Below, we describe three critical situations that illustrate the basic differences in the possible definitions of stability.[17]

In the first situation we have a hospital h with two positions, a single applicant, s_2, and a couple (s_1, s_3). Hospital h prefers s_1 to s_2 and s_2 to s_3. The hospital can decide to admit either the couple or the single applicant. Both decisions can be rational, depending on the matching scheme. In Scotland, the single applicant would be selected. The reason behind this decision is the following. In Scotland, the residents are ranked in each hospital according to a common scoring system and a matching mechanism provides score limits for each hospital such that admitting a new single applicant or a couple by rejecting some other assignees cannot decrease this score limit (more about this in Biró, Irving and Schlotter [2011]). Based on the stability definition[18] given in the paper by Kojima, Pathak and Roth [2010], the single applicant would be chosen, therefore supposedly this solution would be returned in the NRMP as well.[19] This is because the preference of each hospital is assumed to be responsive, and hence the choice of h from the set $\{s_1, s_2, s_3\}$ would be $\{s_1, s_2\}$. So, the single applicant can block with the hospital if it is not matched, but the couple cannot. However, in the two other papers where this option (i.e., having applications from a couple for pairs of positions at the same hospital) has been considered (McDermid and Manlove [2010] and Marx and Schlotter [2011]), the opposite assumption has been made, namely that the couple would be selected. Note that this assumption can be equally reasonable, since admitting (s_1, s_3) instead of s_2 can be seen as an improvement for the hospital, as s_1 is better than s_2 and it is also better to fill the empty slot with s_3 than leaving it unfilled.

The second critical situation is when two couples, (s_1, s_4) and (s_2, s_3), apply for two positions at the same hospital h, where there are exactly two positions. The candidates are ranked by the hospital in the following order: $s_1 P_h s_2 P_h s_3 P_h s_4$ (i.e., s_1 is the best and s_4 is the worst candidate). Either of the two possible matchings may be considered as stable, and actually both are stable according to the definitions presented in Kojima, Pathak and Roth [2010], McDermid and Manlove [2010], and Marx and Schlotter [2011]. However, in Biró, Irving and Schlotter [2011] the unique stable matching matches (s_2, s_3). Essentially, in the SFAS every couple that applies

---

[17]For two reasons we do not provide the full formal definitions of stability for markets with couples. First, the formal definitions would require a considerable amount of space without clear benefits for most non-specialized readers. Second, most results in the literature are negative and even hold when each hospital has only one position (i.e., when all definitions of stability coincide).

[18]In Kojima, Pathak and Roth [2010], a necessary condition for a hospital h to form a blocking pair with a single candidate s is $s \in Ch_h(\mu(h) \cup \{s\})$, and to form a blocking coalition with a couple (s_1, s_2) is $\{s_1, s_2\} \subseteq Ch_h(\mu(h) \cup \{s_1, s_2\})$. Here, $Ch_h(S)$ denotes the choice of hospital h from the set of residents S.

[19]The algorithm used in the NRMP is described in detail in Roth and Peranson [1999], we give a short summary in Subsection 3.4. At this point we only note that in each step of this algorithm either a single resident (or a couple) makes a proposal to a hospital (or to a pair of hospitals). The proposal is accepted if and only if the resident forms a blocking pair with the hospital (or the couple forms a blocking coalition with the hospital or pairs of hospitals) with respect to the actual matching.

for a pair of positions at the same hospital is ranked according to its weakest member.

Finally, the third critical situation is when two couples, (s_1, s_3) and (s_2, s_4), apply for two positions at the same hospital h, and the applicants are again ranked according to their indices. Both possible matchings are stable when stability is defined as in Kojima, Pathak and Roth [2010]: the choice of h from the set $\{s_1, s_2, s_3, s_4\}$ would be $\{s_1, s_2\}$, so neither of the couples can form a blocking coalition with the hospital when the other couple is matched there. However, one could argue that having s_1 instead of s_2 and s_3 instead of s_4 are both strict improvements for the hospital, so only $\{\langle s_1, h_1 \rangle, \langle s_3, h_2 \rangle\}$ should be considered stable. Actually, this is the case in the other three papers mentioned above. Therefore, the stability condition defined in Kojima, Pathak and Roth [2010] (and presumably used in the NRMP) is the least restrictive among the four variants that are present in the literature.

3.3. *Stability: (Loss of) Structure and Computational Complexity*

In this subsection we provide a number of "simple" matching markets that illustrate the difficulties that may arise when extending the many-to-one model to account for the presence of couples. In all our examples each hospital has only one position. Therefore, our findings are valid no matter which definition of stability we use from the previously described literature.

The first example, presented by Roth [2008] and accredited by him to Klaus and Klijn [2005], describes an instance that does not admit a stable matching.[20]

Example 1. (No stable matching.) There are three residents, comprising one single applicant s_2, one couple (s_1, s_3) and two hospitals, each with one position. The applicants are ordered by both hospitals according to their indices (s_1 highest, s_3 lowest), and the preference lists of the residents are as shown in Figure 1. There are three non-empty match-

$$(s_1, s_3) : (h_1, h_2)$$
$$s_2 : h_1 \ h_2$$

Fig. 1. A couples market without a stable matching.

ings for this instance: $M_1 = \{\langle s_1, h_1 \rangle, \langle s_3, h_2 \rangle\}$, $M_2 = \{\langle s_2, h_2 \rangle\}$ and $M_3 = \{\langle s_2, h_1 \rangle\}$. It may readily be verified that M_1 is blocked by $\langle s_2, h_2 \rangle$, M_2 by $\langle s_2, h_1 \rangle$, and M_3 by $\langle (s_1, s_3), (h_1, h_2) \rangle$. ◇

The following basic example illustrates that in the presence of couples, the set of stable matchings does not necessarily have extreme elements and nice structural

---

[20]Roth [1984b] and Marilda Sotomayor [unpublished note] were the first to exhibit couples markets without stable matchings.

properties.[21]

Example 2. (No one-side optimality.) There are four residents, comprising two couples (s_1, s_4) and (s_2, s_3), and two hospitals, each with one position. The applicants are ordered by both hospitals according to their indices (s_1 highest, s_4 lowest), and the preference lists of the residents are as shown in Figure 2. There are two stable matchings: $\mu_1 =$

$$(s_1, s_4) : (h_1, h_2)$$
$$(s_2, s_3) : (h_1, h_2)$$

Fig. 2. A couples market without one-side optimality.

$\{\langle s_1, h_1 \rangle, \langle s_4, h_2 \rangle\}$ and $\mu_2 = \{\langle s_2, h_1 \rangle, \langle s_3, h_2 \rangle\}$, but neither of them is optimal for either side of the market. (This follows since h_1 and (s_1, s_4) prefer μ_1, but h_2 and (s_2, s_3) prefer μ_2.) ◇

We can extend the previous example by adding a pair of positions to the end of one joint preference list to show that the rural hospital theorem (Roth [1984a] and Roth [1986]) does not carry over to markets with couples.[22]

Example 3. (No extension of the rural hospital theorem.) There are four residents, comprising two couples (s_1, s_4) and (s_2, s_3), and four hospitals, each with one position. The applicants are ordered by all hospitals according to their indices (s_1 highest, s_4 lowest), and the preference lists of the residents are as shown in Figure 3. There are only two

$$(s_1, s_4) : (h_1, h_2)$$
$$(s_2, s_3) : (h_1, h_2) \ (h_3, h_4)$$

Fig. 3. A couples market with stable matchings of different sizes.

stable matchings $\{\langle s_1, h_1 \rangle, \langle s_4, h_2 \rangle, \langle s_2, h_3 \rangle, \langle s_3, h_4 \rangle\}$ and $\{\langle s_2, h_1 \rangle, \langle s_3, h_2 \rangle\}$, and they have different sizes. ◇

Combining the markets in Examples 1 and 2, we can easily construct another market (with a stable matching) such that if we start from some unstable matching, there is no path to a stable matching by consecutively satisfying blocking pairs.[23]

---

[21] Aldershof and Carducci [1996] were the first to exhibit a couples markets with stable matchings but without optimal stable matchings.

[22] Aldershof and Carducci [1996] were the first to exhibit a couples markets with stable matchings of different sizes.

[23] Klaus and Klijn [2007] were the first to exhibit a couples markets (with a stable matching) where there is an unstable matching from which there exists no path to a stable matching.

Example 4. (No path to stability.) There are five residents, comprising one single applicant s_3 and two couples (s_1, s_5) and (s_2, s_4), and four hospitals, each with one position. The applicants are ordered by all hospitals according to their indices (s_1 highest, s_5 lowest), and the preference lists of the residents are as shown in Figure 4. There is a unique

$$(s_1, s_5) : (h_1, h_2)$$
$$(s_2, s_4) : (h_1, h_2) \ (h_3, h_4)$$
$$s_3 : h_3 \ h_4$$

Fig. 4. A couples market without a path to stability.

stable matching $\{\langle s_2, h_1 \rangle, \langle s_4, h_2 \rangle\}$. If we start with matching $\{\langle s_1, h_1 \rangle, \langle s_5, h_2 \rangle\}$, then no path of matchings obtained by satisfying blocking pairs leads to the stable matching. $\diamond$

Finally, we illustrate with a similar example that there is no incentive compatible mechanism that yields a stable matching whenever one exists.

Example 5. (No incentive compatibility + "stability whenever possible.") There are five residents, comprising one single applicant s_3 and two couples (s_1, s_5) and (s_4, s_2), and four hospitals, each with one position. The applicants are ordered by all hospitals according to their indices (s_1 highest, s_5 lowest), and the preference lists of the residents are as shown in Figure 5. There is a unique stable matching given by

$$(s_1, s_5) : (h_1, h_2)$$
$$(s_4, s_2) : (h_1, h_2) \ (h_3, h_4)$$
$$s_3 : h_1 \ h_4$$

Fig. 5. A couples market with a unique stable matching.

$\{\langle s_1, h_1 \rangle, \langle s_5, h_2 \rangle, \langle s_4, h_3 \rangle, \langle s_2, h_4 \rangle\}$, where s_3 is unmatched. However, if s_3 were to replace h_1 with h_2 in his preference list, then the unique stable matching for the manipulated instance would be $\{\langle s_4, h_1 \rangle, \langle s_2, h_2 \rangle, \langle s_3, h_4 \rangle\}$. So, any mechanism that always yields a stable matching whenever one exists violates incentive compatibility. $\diamond$

Another piece of bad news for policy makers and mechanism designers is that it is very unlikely that there is an efficient algorithm that always finds a stable matching whenever one exists. Indeed, the problem of finding a stable matching in a market with couples is NP-hard. This was first proved by Ronn [1990] for markets where each hospital has only one position. The same negative result has been proved for even more restricted settings, namely McDermid and Manlove [2010] showed that the problem of finding a stable matching is NP-hard even for "consistent" couples, i.e., couples that apply only for pairs of positions at the same hospital. Furthermore,

Biró, Irving and Schlotter [2011] showed that the problem is still NP-hard if the residents are ranked according to a common order (or "master list") and the joint list of each couple is constructed in a systematic way (which is how for instance the SFAS proceeds).

The problem is also hard from a parameterized point of view when the number of couples is the parameter considered. Marx and Schlotter [2011] proved that the problem is so-called $W[1]$-hard, which means that there exists no FPT-algorithm for this problem unless P=NP.[24] Note that for the restricted problem setting that is present in the SFAS, and which is still NP-hard, Biró, Irving and Schlotter [2011] provide an FPT-algorithm.

Finally, there do exist some positive results for the special case where the set of possible matchings is restricted to those which are stable with respect to the individual preferences as well. Sethuraman, Teo and Qian [2006] and McDermid and Manlove [2010] gave polynomial time algorithms for such problems in different settings.

As most variants of the problems, and certainly all the variants that are present in the current applications, are NP-hard, the mechanism designers can do nothing but try different heuristics to solve the real instances. As we will see, computational experiments show that the problem is relatively easy to solve in practice for resident allocation schemes, where the proportion of couples is small (about 5%). However, in other applications, where the proportion of couples is high, the problem seems very challenging.

3.4. *Heuristics*

Since the National Resident Matching Program has been redesigned (see Roth and Peranson [1999]), the matching algorithm that is used in the program is a heuristic based on the incremental algorithm by Roth and Vande Vate [1990]. The idea of the algorithm is that the applicants are added to the market one-by-one and the matching regains stability after satisfying a bounded number of blocking pairs. The algorithm that is used in the NRMP is very similar; the difference is that couples are also accommodated. The algorithm starts by adding all the single applicants first, and then the couples are added one-by-one with the hope that a stable solution can be reached by the same proposal-rejection sequence for each sub-market. This simple algorithm seems to work well for markets with a small proportion of couples (see Roth and Peranson [1999]). However, in Klaus, Klijn and Massó [2007] it was shown that (i) this method may not find a stable solution, even if the instance

---

[24]An FPT-algorithm would have a running time $f(c) \cdot g(n)$, where c is the number of couples, $f(c)$ is an arbitrary function of c and $g(n)$ is a polynomial function of n, n being the size of the instance. The existence of an FPT-algorithm might suggest that the problem can be tractable for instances where the number of couples is small, even if the problem is NP-hard. However, $W[1]$-hardness implies that the problem is practically intractable, even if the number of couples is small in a matching scheme.

satisfies some special properties that guarantee the existence of a stable solution (see Subsection 3.5), and (ii) the induced mechanism is also vulnerable to strategic manipulation.

Motivated by the practical challenge of constructing an algorithm for the SFAS, Biró, Irving and Schlotter [2011] designed and implemented some "more complex" algorithms and compared their performance with the NRMP-algorithm described in Roth and Peranson [1999] by testing them on random samples. Their finding was that most variants were able to return a stable matching for almost every instance when the proportion of couples was low (around 5%). However, as the proportion of couples increased the success rate of the new heuristics was clearly higher than that of the Roth-Peranson algorithm.

There is a third approach described in a very recent paper by Biró, Fleiner and Irving [2011], which uses the Scarf-algorithm (Scarf [1967]) to find a stable matching. Biró and Fleiner [2010] showed how the Scarf-algorithm can be employed to find fractional core solutions for NTU-games where agents can have capacities. As a possible application, they illustrated how their theory can be applied to matching markets with couples. The Scarf-algorithm always produces a stable fractional solution (which can be interpreted as part-time contracts), and when the solution is integral it corresponds to a stable matching. Biró, Fleiner and Irving [2011] implemented and tested this method and compared its performance with the two above mentioned algorithms for the same random samples that Biró, Irving and Schlotter [2011] studied. The results of a particular experiment are shown in Table 1.

Algorithm	Number of couples										
	12	25	50	75	100	125	150	175	200	225	250
R–P (1999)	952	897	701	547	395	277	170	83	41	9	3
B–I–S (2011)	976	958	911	870	811	752	682	546	281	71	10
Scarf by B–F–I (2011)	895	813	649	532	426	356	316	261	202	174	158

Table 1. Randomly created matching markets with couples for 500 residents.

In this computational experiment, there were 500 residents and the proportion of couples varied from 5% to 100%. For each proportion of couples, the authors generated 1000 random instances and counted how many instances each variant could solve. They considered the Roth-Peranson method (used in the NRMP), some more sophisticated heuristics due to Biró, Irving and Schlotter [2011] (used in the SFAS), and the Scarf-algorithm. The table shows that the heuristics due to Biró, Irving and Schlotter [2011] obtained a much better success ratio than the Roth-Peranson heuristics, especially for high proportions of couples. But surprisingly, the Scarf-algorithm was much better than the others when (almost) all the applicants form couples. Note that this situation occurs in some applications, such as the Hungarian higher education matching scheme (where many applicants apply for a pair of studies).

Another possibility that has not been explored thoroughly is the application of genetic algorithms to generate stable matchings. Aldershof and Carducci [1999] provided such an algorithm. They applied their algorithm to some couples markets of size 50. They noticed that their examples are much smaller than most real-world problems, but the latter involve preference lists that are very short with respect to the number of available choices (e.g., a resident ranks only a very small number of hospital positions from the thousands of available ones). Considering some larger instances with smaller number of applications, they found that genetic enumeration of stable matchings in a market of size 2,000 can be accomplished in a relatively short amount of time. This experiment led them to suggest that it could be possible to use the genetic algorithm for markets of the size of the NRMP.

3.5. *Restricted Preference Domains*

One possible way to get around the non-existence of stable matchings is to consider restricted preference domains for couples. Dutta and Massó [1997] considered a restriction that they called togetherness and proved the existence of a stable matching using an adapted deferred acceptance algorithm. Unfortunately, togetherness is a very strong restriction since it basically makes the couples behave like individual students.

Klaus and Klijn [2005] introduced a natural preference domain for couples, namely the domain of "(weakly) responsive" preferences, and showed that on this domain stable matchings exist.[25] A couple's preferences are responsive if the unilateral improvement of one partner's job is considered beneficial to the couple as well. If responsiveness only applies to acceptable positions, then preferences are weakly responsive. Hence, (weakly) responsive preferences may reflect situations where couples search for jobs in the same metropolitan area. They proved that under a restricted unemployment aversion condition, the domain of weakly responsive preferences is maximal for the existence of stable matchings.[26]

Hatfield and Kojima [2010] introduced a condition called "bilateral substitutes" in the matching with contracts framework of Hatfield and Milgrom [2005]. Matching with couples (where each hospital has only one position) is a special case of matching with contracts. Hatfield and Kojima [2010] proved that weak responsiveness implies bilateral substitutes. Hence, their result that bilateral substitutes is sufficient for the existence of a stable matching subsumes the existence result in Klaus and Klijn [2005]. Hatfield and Kominers [2011] generalized the matching with contracts model in Hatfield and Milgrom [2005] to a many-to-many setting. One special case of the more general model is matching with couples where hospitals can have more than one position. The work in Hatfield and Kominers [2011] showed that for hospitals

---

[25]See also Klaus, Klijn and Nakamura [2009].

[26]More precisely, in a couples market with only restricted strictly unemployment averse couples and at least one couple whose preferences are not weakly responsive, there exist (weakly) responsive preferences for the other couple(s) such that no stable matching exists.

with substitutable preferences (i.e., instead of responsive preferences), the domain of substitutable preferences is a maximal domain of couples' preferences under which a stable matching is guaranteed to exist.

In a decentralized market, stable outcomes can always be reached provided that couples' preferences are weakly responsive. Starting from an arbitrary matching, Klaus and Klijn [2007] showed how to construct a path of matchings obtained from "satisfying" blocking coalitions that yields a stable matching.

Although on the domain of responsive preferences a stable matching always exists, Klaus and Klijn [2005] showed that *even on this domain* other desirable properties of stable matchings for singles markets do not carry over to couples markets:

- there may be no optimal stable matching for either side of the market;
- different stable matchings may fill a different number of positions;
- there is no stable matching mechanism for which stating the true preferences is a dominant strategy for every couple.

One possible reason why couples' preferences might not satisfy weak responsiveness is geographic. For instance, if a couple's most preferred option is two hospitals in one region and its second most preferred option is two hospitals in some other region, then responsiveness is violated. Cantala [2004] considered a model with different regions of hospital locations and assumed that couples prefer to be matched in the same region rather than in distinct ones. He further assumed that couples share a ranking of regions. Despite these strong restrictions on the couples' preferences, Cantala [2004] exhibited a couples market that fails to have a stable matching.

3.6. *Large Markets*

Very recently, Kojima, Pathak and Roth [2010] analyzed data from the Association of Psychology Postdoctoral and Internship Centers (APPIC) matching program. They found that (i) the preferences of the overwhelming majority of couples violate weak responsiveness, but nonetheless (ii) there exists a stable matching for each of the nine years of the psychology market for which they had data. To explain the stability of the psychologists market (and similar large markets), Kojima, Pathak and Roth [2010] built a model for large markets with couples. Using stylized facts from the psychologist market, they placed several conditions on the sequence of random markets (indexed by n, which is the number of hospitals) of growing size (in order to study limit behavior). Informally,[27]

- the length of students' preference lists does not grow when the number of market participants grows;

---

[27]See Kojima, Pathak and Roth [2010] for details, in particular regarding possible relaxations of the assumptions.

- the number of single students does not grow much faster than the number of hospitals. Moreover, the number of couples does not grow at the same rate as the number of hospitals and instead grows at the slower rate of $O(n^a)$, where $0 \leq a \leq \frac{1}{2}$;
- the capacity of each hospital is bounded;
- the popularity of different hospitals (measured by the probability of being listed by doctors as acceptable) does not vary too "much."

A crucial implication of the model is that if n is large enough then it is very likely that there is a large number of hospitals with vacant positions (even if there are more students than positions). This implication partially drives their first main result: the probability that there exists a stable matching converges to 1 when the number of hospitals n approaches infinity. The proof of their first main result consists of showing that the probability that their "sequential couples algorithm" succeeds (i.e., finds a stable matching) converges to one as the market size approaches infinity.[28]

Next, Kojima, Pathak and Roth [2010] proceeded with an analysis of the incentives for manipulation in a market with couples, employing a mechanism that builds on the sequential couples algorithm.[29] They considered a Bayesian game in which single students, couples, and hospitals submit their preference lists to the clearinghouse. Under conditions (on a sequence of Bayesian games) that are very similar to the ones in their first main result, they established their second main result: for any $\epsilon > 0$, truth-telling is an ϵ-Bayes-Nash equilibrium when the number of hospitals n approaches infinity.

Ashlagi, Braverman and Hassidim [2011] studied a model very similar to Kojima, Pathak and Roth [2010], but with a weaker notion of stability. They allowed the number of couples to grow at a near-linear rate: $n^{1-\epsilon}$ where $\epsilon > 0$. Under this assumption, they established that the probability that a stable matching exists approaches 1 as the market grows large. To obtain their result, they introduced a new matching algorithm called "sorted deferred acceptance." Moreover, under the induced mechanism, truth-telling is an approximate Bayes-Nash equilibrium for any large enough market. A key element in their analysis is the selection of a particular sequence of proposals by couples that is least likely to generate existence problems. Furthermore, their simulation and theory suggest that asymptotic existence is likely to fail if the number of couples grows proportionally to the market size. Therefore, the growth rate $n^{1-\epsilon}$ is probably the highest one can achieve without running into the non-existence of stable matchings.

---

[28] The sequential couples algorithm is based on the deferred acceptance algorithm and incorporates applications by couples. Obviously, the sequential couples algorithm may fail to find a stable matching (even if there exists one).

[29] The difference is that the mechanism has to specify some matching when the sequential couples algorithm fails (to find a stable matching).

4. Concluding Discussion and Outlook

Studies in the last decade have answered many questions regarding the presence of couples in matching markets. Most of these studies are theoretical, but many of them are motivated by empirical findings. Given the possible non-existence of stable matchings, computational experiments have been very helpful as well to measure and compare the performance of algorithms that aim to find a stable matching whenever one exists.

Obviously, there is still a long road to a full understanding of the complexities introduced by the presence of married couples. One theoretical question is how to weaken stability to obtain an appropriate non-empty solution. For instance, the weak stability proposed in Klijn and Massó [2003] does not seem adequate: Klaus and Klijn [2005] showed that most of their negative results hold even for weak stability. A helpful tool in such an analysis could be a concept of comparing (unstable) matchings according to the number of blocking pairs they admit or number of agents involved in blocking pairs (see e.g. Biró, Manlove and Mittal [2010]), or other measures like the one developed by Biermann [2011].

Another important issue that has received relatively little attention is the matching of couples in decentralized markets. To the best of our knowledge, Klaus and Klijn [2007] is the only paper that explicitly dealt with it. They showed that on their restricted preference domain, under rather mild conditions a decentralized market converges to a stable matching with probability 1. They also showed that in general, even if a stable matching exists, such convergence is not always possible (see also Example 4). In order to obtain positive results it therefore seems necessary to study and model the behavior of couples in more *specific* real-life decentralized markets.

To a large extent, the literature on matching with couples belongs to the rather new field of *market design*. "Economic engineering" (coined by Roth [2002]), however, typically takes a trial and error approach, in which laboratory experiments play a prominent role. We are not aware of any laboratory experiments that deal with couples in matching markets, but they would be crucial in the study of couples' preference formation and decision making.

The analysis of matching markets with couples is a first small step in understanding markets with complementarities. Until recently, these markets were difficult to tackle. However, the tools developed in large markets studies can turn out to be very helpful. An important example of this is a result obtained by Ashlagi, Braverman and Hassidim [2011]. They adapted their analysis of large markets with couples to a more general worker-firm model studied by Roth and Sotomayor [1990]. Roth and Sotomayor (1990) proved that if firms' preferences are substitutable a stable matching exists; and even if only one firm does not have substitutable preferences, a stable matching need not exist. Ashlagi, Braverman and Hassidim [2011] showed that if the market is large and random, as long as the number of firms that have non-substitutable preferences is growing at a near-linear rate, the probability that a stable matching exists tends to 1 as the market size tends to infinity.

To put our survey in perspective, we note that there are some new interdisciplinary research directions that incorporate our particular subject. *Algorithmic mechanism design* (see e.g. Nisan and Ronen [1999]) is one of them. In both the economics and computer science communities it is now widely accepted that a mechanism can be successful only if both strategic and algorithmic considerations form part of the design process. More general interdisciplinary areas such as *algorithmic game theory* (see e.g. Nisan, Roughgarden, Tardos and Vazirani [2007]) and *computational social choice* (see e.g. Chevaleyre, Endriss, Lang and Maudet [2007]) have already shown the fruitful interaction between economics and computer science.

References

Abdulkadiroğlu, A. and Sönmez, T. [2010] Matching Markets: Theory and Practice, Prepared for the 2010 Econometric Society World Congress, Shanghai, China.

Ackermann, H. Goldberg, P.W. Mirrokni, V. Röglin, H. and Vöcking, B. [2008] Uncoordinated Two-Sided Matching Markets, *Proc. 9th ACM Conference on Electronic Commerce (EC'08) Chicago, USA*, 256-263.

Aldershof, B. and Carducci, O.M. [1996] Stable Matchings with Couples, *Discrete Applied Mathematics*, **68**, 203–207.

Aldershof, B. and Carducci, O.M. [1999] Stable Marriage and Genetic Algorithms: A Fertile Union, *Journal of Heuristics*, **5**, 29–46.

Ashlagi, I., M. Braverman and Hassidim, A. [2011] Stability in Large Matching Markets with Complementarities, Working Paper, Sloan School of Management, MIT.

Biermann, F.M. [2011] A Measure to Compare Matchings in Marriage Markets, Fondazione ENI Enrico Mattei, Working Paper 592.

Biró, P. [2008] Student Admissions in Hungary as Gale and Shapley Envisaged, Technical Report no. TR-2008-291 of the Computing Science Department of Glasgow University.

Biró, P. and Fleiner, T. [2010] Fractional Solutions for NTU-games, Technical Report no. TR-2010-315 of the Computing Science Department of Glasgow University.

Biró, P., Fleiner, T. and Irving, R.W. [2011] Matching Couples with Scarf's Algorithm, Mimeo, Institute of Economics, Hungarian Academy of Sciences.

Biró, P., Fleiner, T. Irving, R.W. and Manlove, D. [2010] The College Admissions Problem with Lower and Common Quotas, *Theoretical Computer Science*, **411**, 3136–3153.

Biró, P., Irving, R.W. and Schlotter, I. [2011] Stable Matching with Couples: An empirical study, *Journal of Experimental Algorithmics*, **16**, Article No.: 1.2.

Biró, P., Manlove, D. and Mittal, S. [2010] Size versus Stability in the Marriage Problem, *Theoretical Computer Science*, **411**, 1828–1841.

Biró, P. and McDermid, E.J. [2010] Matching with Sizes (or Scheduling with Processing Set Restrictions), *In Proceedings of ISCO: International Symposium on Combinatorial Optimization, volume 36 of Electronic Notes on Discrete Mathematics*, 335–342.

Bogomolnaia, A. and Jackson, M.O. [2002] The Stability of Hedonic Coalition Structures, *Games and Economic Behavior*, **38**, 201–230.

Cantala, D. [2004] Matching Markets: The Particular Case of Couples, *Economics Bulletin*, **45**, 29–46.

Chevaleyre, Y., Endriss, U. Lang, J. and Maudet, N. [2007] A Short Introduction to Computational Social Choice, *In Proceedings of SOFSEM-2007: the 33rd Conference on Current Trends in Theory and Practice of Computer Science, LNCS 4362, Springer-Verlag*, 51–69.

Dutta, B. and Massó, J. [1997] Stability of Matchings when Individuals have Preferences over Colleagues, *Journal of Economic Theory* **75**, 464–475.

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

Gusfield, D. and Irving, R.W. [1989] The Stable Marriage Problem: Structure and Algorithms, MIT Press, Boston, Mass.,1989.

Hatfield, J.W. and Kojima, F. [2010] Substitutes and Stability for Matching with Contracts, *Journal of Economic Theory*, **145**, 1704–1723.

Hatfield, J.W. and Kominers, S.D. [2011] Contract Design and Stability in Matching Markets, Working Paper, Graduate School of Business, Stanford University.

Hatfield, J.W. and Milgrom, P. [2005] Matching with Contracts, *American Economic Review*, **95**, 913-935.

Immorlica, N. and Mahdian, M. [2005] Marriage, Honesty, and Stability, *In Proceedings of SODA 2005*, 53-62.

Klaus, B. and Klijn, F. [2005] Stable Matchings and Preferences of Couples, *Journal of Economic Theory*, **121**, 75–106.

Klaus, B. and Klijn, F. [2007] Paths to Stability for Matching Markets with Couples, *Games and Economic Behavior*, **58**, 154–171.

Klaus, B., Klijn, F. and Massó, J. [2007] Some Things Couples Always Wanted to Know about Stable Matchings (but Were Afraid to Ask), *Review of Economic Design*, **11**, 175–184.

Klaus, B., Klijn, F. and Nakamura, T. [2009] Corrigendum: Stable Matchings and Preferences of Couples, *Journal of Economic Theory*, **144**, 2227–2233.

Klijn, F. [2008] Mechanism Design in School Choice: Some Lessons in a Nutshell, *Boletín de la Sociedad de Estadística e Investigación Operativa*, **24**, 11–22.

Klijn, F. and Massó, J. [2003] Weak Stability and a Bargaining Set for the Marriage Model, *Games and Economic Behavior*, **42**, 91-100.

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

Kojima, F., Pathak, P.A. and Roth, A.E. [2010] Matching with Couples: Stability and Incentives in Large Markets, Working Paper, Department of Economics, Stanford University.

Kojima, F. and Troyan, P. [2011] Matching and Market Design: An Introduction to Selected Topics, forthcoming in *Japanese Economic Review*.

Kojima, F. and Ünver, M.U. [2006] Random Paths to Pairwise Stability in Many-to-Many Matching Problems: A Study on Market Equilibration, *International Journal of Game Theory*, **36**, 473–488.

Knuth, D.E. [1976] Marriages Stables. Les Presses de l'Université, Montreal.

Marx, D. and Schlotter, I. [2011] Stable assignment with couples: parameterized complexity and local search, *Discrete Optimization*, **8**, 25–40.

McDermid, E.J. and Manlove, D.F. [2010] Keeping Partners Together: Algorithmic Results for the Hospitals/Residents Problem with Couples, *Journal of Combinatorial Optimization*, **19**, 279–303

Nisan, N. and Ronen, A. [1999] Algorithmic mechanism design, *In Proceedings of STOC 1999*, 129-140.

Nisan, N., Roughgarden, T. Tardos, É. and Vazirani, V.V. Eds. [2007] Algorithmic Game Theory, Cambringe University Press, 2007.

Robards, P.A. [2001] Applying Two-Sided Matching Processes to the United States Navy Enlisted Assignment Process, Master's Thesis, Naval Postgraduate School Monterey CA, 2001.

Ronn, E. [1990] NP-Complete Stable Matching Problems, *Journal of Algorithms*, **11**, 285–304.

Roth, A.E. [1982] Incentive Compatibility in a Market with Indivisible Goods, *Economics Letters*, **9**, 127–132.

Roth, A.E. [1984a] Stability and Polarization of Interests in Job Matching, *Econometrica*, **52**, 47–58.

Roth, A.E. [1984b] The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory, *Journal of Political Economy*, **92**, 991–1016.

Roth, A.E. [1985] The College Admissions Problem is not Equivalent to the Marriage Problem, *Journal of Economic Theory*, **36**, 277–288.

Roth, A.E. [1986] On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets, *Econometrica*, **54**, 425–427.

Roth, A.E. [1990] New Physicians: A Natural Experiment in Market Organization, *Science*, **250**, 1524–1528.

Roth, A.E. [1991] A Natural Experiment in the Organization of Entry-level Labor Markets: Regional Markets for New physicians and Surgeons in the United Kingdom, *American Economic Review*, **81**, 415–440.

Roth, A.E. [2002] The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics, *Econometrica*, **70**, 1341-1378.

Roth, A.E. [2008] Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions, *International Journal of Game Theory*, **36**, 537–569.

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

Roth, A.E. and Sotomayor, M.A.O. [1990] Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis. Econometric Society Monograph Series. Cambridge University Press, New York.

Roth, A.E. and Xing, X. [1994] Jumping the Gun: Imperfections and Institutions related to the Timing of Market Transactions, *American Economic Review*, **84**, 992–1044.

Roth, A.E. and Vande Vate, J.H. [1990] Random Paths to Stability in Two-Sided Matching, *Econometrica*, **58**, 1475-1480.

Scarf, H.E. [1967] The Core of an N-Person Game, *Econometrica*, **35**, 50-69.

Sethuraman, J., Teo, C.-P. and Qian, L. [2006] Many-To-One Stable Matching: Geometry and Fairness, *Mathematics of Operations Research*, **31**, 581-596.

Short, M.M. [2000] Analysis of the Current Navy Enlisted Detailing Process, Master's Thesis, Naval Postgraduate School Monterey CA, 2000.

Sönmez, T. [1997] Manipulation via Capacities in Two-Sided Matching Markets, *Journal of Economic Theory*, **77**, 197-204.

Sönmez, T. [1999] Can Pre-arranged Matches Be Avoided in Two-Sided Matching Markets?, *Journal of Economic Theory*, **86**, 148-156.

Sönmez, T. and Ünver, M.U. [2011] Matching, Allocation, and Exchange of Discrete Resources. Handbook of Social Economics, Vol. 1A, pp. 781-852, Alberto Bisin, Jess Benhabib, and Matthew Jackson, eds., North-Holland, 2011.

Sudarshan, A. and Zisook, S. [1981] National Resident Matching Program, *New England Journal of Medicine*, **305**, 525–526.

Yang, W., Giampapa, J.A. and Sycaratech, K. [2003] Two-Sided Matching for the U.S. Navy Detailing Process with Market Complications, Technical Report CMU-RI-TR-03-49, Robotics Institute, Carnegie Mellon University, 2003.