

## Matching

Muriel Niederle  
*Stanford University and NBER*

Alvin E. Roth  
*Harvard University*

Tayfun Sönmez  
*Boston College*

Preliminary, Feb 11, 2007

Forthcoming in *The New Palgrave Dictionary of Economics*, 2nd edition,  
Palgrave Macmillan

“Matching” is the part of economics that focuses on the question of who gets what, particularly when the scarce goods to be allocated are heterogeneous and indivisible; e.g. who works at which job, which students go to which school, who receives which transplantable organ, etc. Studying how particular matching markets succeed at creating efficient matches, or fail to do so, has yielded insights into how markets in general work well or badly.

Because market failures have sometimes been successfully fixed by devising new rules for both centralized and decentralized market organization, matching has been a major focus of the emerging field of market design. Some designs by economists have included labor market clearinghouses for doctors and other healthcare workers in the United States, both for their first jobs, and as they enter specialties. Clearinghouses have also been implemented in less traditional markets, which do not adjust prices/wages to help clear the market, such as the matching of students to schools in New York City and in Boston. And new clearinghouses are being implemented for the organization of live-donor kidney exchanges, among patients in need of a kidney transplant who have willing donors with whom they are incompatible.

In the next section we review some studies of matching, including some market failures that have been addressed either by introducing appropriate rules to a decentralized market (as in admissions to graduate programs in American universities), or by introducing a centralized clearinghouse (as in the markets for new doctors in the U.S., Canada, and Britain). The subsequent two sections consider the simple theory behind some clearinghouse designs. Then we’ll return to some of the successful market design applications, which build on the theoretical models, but handle practical problems that are sometimes not yet fully understood in theory.

We’ll focus on three kinds of market failure that sometimes impede efficient matching.

1. Failure to provide *thickness*; i.e. to bring together enough buyers and sellers to transact with each other.
2. Failure to overcome the *congestion* that thickness can bring, i.e. that can result when lots of buyers and sellers are trying to transact. That is, failure to provide enough time, or failure to make transactions fast enough that market participants can consider enough alternative possible transactions to arrive at satisfactory ones.
3. Failure to make it *safe* for market participants to reliably reveal or otherwise act on their information.

## **I. Some Market Failures and their Consequences**

### **I.A. Unraveling, congestion, and centralized clearinghouses**

A variety of professional labor markets have suffered from the *unraveling* of appointment dates: from year to year, appointments became earlier and earlier in advance of actual employment. Markets that had once been thick, with many employers and applicants on the market at the same time, became thin, as potential employees faced early offers, dispersed in time, to which they had to respond before they could learn what other offers might be forthcoming. That is, applicants often received “exploding” offers that had to be accepted or rejected without waiting to see if a more desirable offer might be coming. An applicant who accepts such an offer, in case acceptances are binding, will never learn of the more desirable offers that might have become available, but if the offer is reasonably desirable, rejecting it might be very risky. And, when applicants are quickly accepting offers in this way, employers, when they make offers, have to start taking into account whether the offer is likely to be accepted, since by the time an offer is rejected, other desirable applicants may have already accepted offers elsewhere. This often makes unraveling a dynamic process, with offers becoming earlier and shorter in duration from year to year. This kind of unraveling has been described in detail in markets for lawyers (Avery et al. 2001), gastroenterologists (Niederle et al. 2006) and many others (cf. Roth and Xing 1994). A clear example is the market for new doctors (Roth 1984).

The first job for almost all new doctors in the United States and Canada is as an intern or resident at a hospital. In the early 1900’s, medical graduates were hired for such jobs near the end of their fourth year of medical school, just before graduation. By the 1930’s, hiring was largely completed half a year before graduation, and by the 1940’s it had moved to sometimes as much as two years before graduation. That is, in the early 1940’s, students were being hired long before they would begin work, at dispersed times, and without much opportunity to consider alternatives, and long before they had sufficient experience to reveal either to employers or themselves what kinds of medicine they would most prefer and be best able to practice. There was widespread recognition among the participants that the market was often failing to create the most productive matches of doctors to hospitals, both because there was too little opportunity to consider alternatives, and because the matching was being done before important information about students became available.

One way in which many markets tried to address this failure was by attempting to establish rules concerning when offers could be made. In the market for new American doctors, the most concerted attempt at this kind of solution began in 1945, with the help of the medical schools, who agreed not to release any information to hospitals about students until a specified date.

However, the market experienced congestion, and over the next few years, students were called upon to make increasingly prompt decisions whether to accept offers. In 1945 offers were supposed to remain open for 10 days. By 1949 a deadline of 12 hours had been rejected as too long. Hospitals were finding that, if an offer was rejected after even a brief period of consideration, it was often too late for them to reach their next most preferred candidates before they had accepted other offers. Even when there was a long deadline, much of this action was compressed into the last moments, because a student who had been offered a position that wasn't his very first choice would be inclined to wait as long as possible before accepting, in the hope of eventually being offered a preferable position. So hospitals felt pressed to pressure students to reply immediately, and offers conveyed by telegram were frequently followed by phone calls requesting an immediate reply.

Congestion can be a problem in any market in which transactions take some time, but it is especially visible in entry-level professional labor markets in which many workers and jobs become available at the same time (e.g. after graduation from university, medical school, law school, etc.)

In the face of congestion, many markets unravel, as employers try to gain more time to make offers by starting to make offers earlier. But the market for new doctors found a solution in the form of a centralized clearinghouse. Starting in the early 1950's, the various medical groups organized a centralized clearinghouse, which remains in use today, having undergone some changes over the years. Nowadays, a medical student applies to hospitals and goes on interviews in the winter of the final year of medical school, and then in February submits an ordered preference list of positions to the centralized clearinghouse, the National Resident Matching Program (NRMP). At the same time, the residency programs (the employers) submit an ordered preference list of candidates. Once all the preference lists are collected, the clearinghouse uses an algorithm to suggest a match, and residency programs and applicants are informed to whom they have been matched. Although this began as an entirely voluntary program, it has been so successful that today it is virtually the only way that most residencies are filled. As we will see below, that success depends critically on the matching algorithm.

The NRMP, and clearinghouses like it, also make very clear the kinds of issues involved for a marketplace to make it safe for participants to reveal their information. In a clearinghouse in which you are asked to state your preferences, the question is simply, is it a good idea to state your true preferences, or would you do better otherwise? For the NRMP we'll see that stating true preferences is indeed both safe and sensible. We'll also discuss clearinghouses that failed this test, like the one for placing students into schools in Boston, that consequently failed to accomplish their objectives in other ways also, and were redesigned.

Before presenting some formal models that will allow us to start to explain which matching algorithms and clearinghouses have been successful and which have failed, it will be helpful to think about several different kinds of matching markets.

### **I.B. Two-sided and One-sided matching markets**

Labor markets, like the market for new doctors, are usually modeled as two-sided markets, in which agents on one side of the market (workers) need to be matched with agents on the other side (employers), and each agent has preferences over possible matches. We'll see below that this two-sided structure allows strong conclusions to be drawn about the properties of matchings, and matching mechanisms.

In many markets this two-sided structure is absent. One way this occurs is when any participant in the market can be matched with any other. For example, if a group of people want to form pairs to be roommates or bridge partners, any one of them can in principle be matched with any other, although not all matchings would be efficient. We'll encounter markets of this kind when we speak of kidney exchange.

Another way in which markets can be one-sided is if the agents in the market need to be matched to objects, e.g. when people need to be assigned rooms in a dormitory, or places in a public school that doesn't itself have preferences or take strategic actions (unlike in a two sided matching market). That is, such a market matches people to places, but only one side, the people, are active participants in the market. Some markets can also be hybrids, with both two- and one-sided properties (as when schools aren't strategic players, but still have priorities over students).

Below we consider some static models of two- and one-sided matching that have proved useful in the design of clearinghouses, and in understanding what they do. In the section on design, we'll also speak about some decentralized design solutions to various market failures, such as unraveling. While there has been some good initial progress on formal models of decentralized markets, and dynamic models in which phenomena like unraveling can play out over time (see e.g. Li and Rosen 1998), these areas are still in need of development, and have not yet received the theoretical attention commensurate with their importance in the study of markets generally.

## **II. Formal models of matching**

### **II.A. Two-sided matching models**

The workhorse models of two sided matching come in several varieties. The simplest, presented in detail below, is the "marriage model" in which each firm seeks to hire only a single worker, and wages and other kinds of price adjustment are represented simply in the preferences that workers and firms have for each other (e.g. in these models, wages are part of the job description that determines preferences). However the kinds of results

we present here generalize to models in which wage and price formation is explicitly included, some pointers to such models are included in the references.

The marriage model consists of two disjoint sets of agents, men =  $\{m_1, \dots, m_n\}$  and women =  $\{w_1, \dots, w_p\}$ , each of whom has complete and transitive preferences over the agents on the other side (and the possibility of being unmatched, which we model as being “matched to yourself”). Preferences can be represented as rank order lists, e.g. if man  $m_i$ 's first choice is  $w_3$ , his second choice  $w_2$  [ $w_3 >_{m_i} w_2$ ] and so on, until at some point he prefers to remain unmatched, that is  $P(m_i) = w_3, w_2, \dots, m_i, \dots$ . If agent  $k$  (on either side of the market) prefers to remain single rather than be matched to agent  $j$ , i.e. if  $k >_k j$ , then  $j$  is said to be *unacceptable* to  $k$ . If an agent is not indifferent between any two acceptable mates, or between being matched and unmatched, we'll say he/she has *strict* preferences.

An *outcome* of the game is a *matching*:  $\mu: M \cup W \rightarrow M \cup W$  such that  $w = \mu(m)$  iff  $\mu(w) = m$ , and for all  $m$  and  $w$  either  $\mu(w)$  is in  $M$  or  $\mu(w) = w$ , and either  $\mu(m)$  is in  $W$  or  $\mu(m) = m$ . That is, a matching matches agents on one side to agents on the other side, or to themselves, and if  $w$  is matched to  $m$ , then  $m$  is matched to  $w$ .

A matching  $\mu$  is *blocked by an individual*  $k$  if  $k$  prefers being single to being matched with  $\mu(k)$ , i.e.  $k >_k \mu(k)$ . A matching  $\mu$  is *blocked by a pair of agents*  $(m, w)$  if they each prefer each other to the partner they receive at  $\mu$ , i.e.  $w >_m \mu(m)$  and  $m >_w \mu(w)$ .

A matching  $\mu$  is *stable* if it isn't blocked by any individual or pair of agents.

A stable matching is efficient, and in the core, and in this simple model the set of (pairwise) stable matchings equals the core.

### **Theorem 1 (Gale and Shapley, 1962)**

A stable matching exists for every marriage market.

Gale and Shapley approached this problem from a purely theoretical perspective, but proved this theorem via a constructive algorithm of the kind that has subsequently turned up at the heart of a variety of clearinghouses.

Deferred Acceptance Algorithm, with men proposing  
(roughly the Gale-Shapley 1962 version)

0. *If some preferences are not strict*, arbitrarily break ties (e.g. if some  $m$  is indifferent between  $w_i$  and  $w_j$ , order them consecutively in alphabetical order. Different agents may break ties differently: i.e. tie-breaking can be decentralized by having each agent fill out a strict preference list...)

1 a. Each man  $m$  proposes to his 1st choice (if he has any acceptable choices).

b. Each woman rejects any unacceptable proposals and, if more than one acceptable proposal is received, "holds" the most preferred and rejects all others....

k a. Any man who was rejected at step k-1 makes a new proposal to its most preferred acceptable mate who hasn't yet rejected him. (If no acceptable choices remain, he makes no proposal.)

b. Each woman holds her most preferred acceptable offer to date, and rejects the rest.  
STOP: when no further proposals are made, and match each woman to the man (if any) whose proposal she is holding.

Note that the proof of the theorem now follows from the observation that the matching produced in this way is itself stable. If some man would prefer to be matched to a woman other than his assigned mate, he must, according to the algorithm, have already proposed to her, and she has rejected him, meaning she has a man she strictly prefers, hence they cannot form a blocking pair.

Roth (1984) showed that the algorithm adopted by the medical clearinghouse in the 1950's was equivalent to the hospital proposing deferred acceptance algorithm. Gale and Shapley observed that which side of the market proposes in a deferred acceptance algorithm has consequences.

**Theorem 2 (Gale and Shapley 1962)**

When all men and women have strict preferences, there always exists an M-optimal stable matching (that every man likes at least as well as any other stable matching), and a W-optimal stable matching. Furthermore, the matching  $\mu_M$  produced by the deferred acceptance algorithm with men proposing is the M-optimal stable matching. The W-optimal stable matching is the matching  $\mu_W$  produced by the algorithm when the women propose.

Note that the algorithm has been stated as if people take actions in the course of the algorithm, and we can ask whether those actions would best serve their interests. To put it another way, is it possible to design a clearinghouse in which a matching is produced from participants' stated rank order lists in such a way that it will never be in someone's interest to submit a rank order list different from their true preferences? The following theorem answers that question in the negative.

**Theorem 3 Impossibility Theorem** (Roth –see Roth and Sotomayor 1990):

No stable matching mechanism exists for which stating the true preferences is a dominant strategy for every agent.

However it is possible to design the mechanism so that one side of the market can never do any better than to state their true preferences.

**Theorem 4** (Dubins and Freedman, Roth – see Roth and Sotomayor 1990)

The mechanism that yields the M-optimal stable matching (in terms of the stated preferences) makes it a dominant strategy for each man to state his true preferences.

The conclusions of Theorems 1-3 also hold for a variety of related models (in which firms employ multiple workers, and wages are explicitly allowed to vary, see e.g. Shapley and Shubik 1971, and Kelso and Crawford, 1982 for notable early models of matching with money, and see Roth and Sotomayor 1990, and Hatfield and Milgrom 2005). However when we look at many-to-one matching models (in which firms employ multiple workers, but workers seek just one job), we have to be careful. It turns out that no procedures exist that give firms a dominant strategy, but that a worker proposing deferred acceptance algorithm still makes it a dominant strategy for workers to state their true preferences (see Roth and Sotomayor 1990 for more details and further references). (These results are closely connected to related results in **auction theory**, see in particular Hatfield and Milgrom 2005, and Milgrom 2004.)

When the market for medical residents was redesigned (Roth and Peranson 1999), a number of practical complications had to be dealt with, such as the fact that about 1,000 graduates a year go through the match as couples who wish to be matched to nearby jobs, and hence have joint preferences over pairs of residency programs. While this can cause the set of stable matchings to be empty, in practice this has not proved to be a significant problem (see also Roth 2002 on engineering aspects of economic design)

## II.B. One-Sided Matching Models

### Shapley and Scarf's "House" Markets

One of the basic models in matching markets was introduced by Shapley and Scarf (1974). Their model is a simple barter economy in which each one of  $n$  agents owns an indivisible good (which they call a house) and has preferences over all houses in the economy. Each agent has use for only one house and trade is only feasible in houses (i.e. there is no money in their model). An **allocation**  $\mu$  in this context is a matching of houses and agents so that each agent receives one and only one house. An exchange in this market does not need to be bilateral. An allocation  $\mu$  is in the **core** if no coalition (including single agent coalitions) of agents can improve upon it (in the sense that all are weakly better off and at least one is strictly better off) by swapping their *own* houses. Shapley and Scarf attribute to David Gale the following **Top Trading Cycles** algorithm (TTC) which can be used to find a core allocation for any housing market:

**Step 1:** Each agent points to the owner of her most preferred house (which could possibly be herself). Since there are finite number of agents there is at least one **cycle** (where a cycle is an ordered list  $(i_1, i_2, \dots, i_k)$  of agents with each agent pointing to the next agent in the list and agent  $i_k$  pointing to agent  $i_1$ ). In each cycle the implied exchange is carried out and the procedure is carried out with the remaining agents.

In general, at

**Step  $r$ :** Each remaining agent points to the owner of her most preferred house among the remaining houses. There is at least one cycle. In each cycle the implied exchange is carried out and the procedure is carried out with the remaining agents.

The algorithm terminates when each agent receives a house.

**Proposition 1** (Shapley and Scarf 1974): The TTC algorithm yields an allocation in the core for each housing market.

The core has some remarkable properties in the context of housing markets. The following propositions summarize the most notable of these results.

While exchange is feasible only in houses, a **competitive allocation** of a housing market can be defined via “token money.” There is an important relation between the core and the competitive allocation for this very basic barter economy.

**Proposition 2** (Roth and Postlewaite 1977): There is a unique allocation in the core (which can be obtained with the TTC algorithm) when agents have strict preferences over houses. Moreover the unique core allocation coincides with the unique competitive allocation.

Another remarkable feature of this model is that the top trading cycles mechanism makes it safe for agents to reveal their true preferences.

**Proposition 3** (Roth 1982): The core as a mechanism is **strategy-proof** when agents have strict preferences over houses. That is, truth-telling is a dominant strategy equilibrium of the preference revelation game induced by the core.

Moreover it is essentially the only mechanism to do so among those that are Pareto efficient and **individually rational** (in the sense that an agent never receives a house inferior to her own).

**Proposition 4** (Ma 1994): The core is the only mechanism that is Pareto efficient, individually rational and strategy-proof.

### House Allocation Problems

Hylland and Zeckhauser (1977) introduced the **house allocation problems** which only differs from housing markets in *property rights*: There are  $n$  houses to be allocated for  $n$  agents where each agent has use for only one house and has strict preferences over all houses. Unlike in housing markets no agent owns a specific house. The mechanism known as the **random serial dictatorship** (RSD) is widely used in real-life applications of these problems. Under RSD agents are randomly ordered in a list with uniform distribution and the first agent in the list is assigned her top choice house, the next agent is assigned her top choice among the remaining houses, and so on. In addition to its popularity in practice, RSD has solid incentives and efficiency properties.

**Proposition 5**: RSD is ex-post Pareto efficient and strategy-proof.

Recall that the only difference between house allocation problems and housing markets is the initial property rights and the core is very well-behaved in the context of the latter. This observation motivates the mechanism **core from random endowments (CRE)**: Randomly assign houses to agents with uniform distribution, interpret the resulting matching as initial allocation and pick the core of the resulting housing market. It turns out, CRE is equivalent to RSD.

**Proposition 6** (Abdulkadiroğlu and Sönmez 1998): For any house allocation problem CRE and RSD yields the same lottery and hence they are equivalent mechanisms.

### **House Allocation with Existing Tenants**

Housing markets and house allocation problems have very different structures on property rights. The former is a pure private ownership economy where each house “belongs” to a specific agent whereas the latter is a collective ownership economy where no strict subset of the grand coalition has claims on any house. Abdulkadiroğlu and Sönmez (1999) introduced the following hybrid **house allocation with existing tenants** model. There are two kinds of agents: *Existing tenants* each of whom owns a house and *newcomers* none of whom have claims on a specific house. In addition to the *occupied houses* owned by existing tenants, there are also *vacant houses*. As in house allocation problems no specific person or group has claims on any vacant house. Suppose that the number of newcomers is equal to the number of vacant houses and hence the number of agents is equal to the number of houses. Agents have strict preferences over all houses and each existing tenant is allowed to keep her current house.

Abdulkadiroğlu and Sönmez introduced the following **you request my house – I get your turn** algorithm (YRMH-IGYT) which generalizes TTC as well as RSD. Under YRMH-IGYT, agents are randomly ordered in a line and initially only the vacant houses are **available**. The first agent in the line is assigned her top choice provided that it is either her own house or an available house (in which case her own house becomes available) and the process continues with the next agent in the line. If however her top choice is an occupied house, the line is adjusted and the owner of the requested house is carried right in front of the requester. The process continues in a similar way with either him getting assigned his own house or an available house (making his own house available), or otherwise him requesting an occupied house and upgrading its owner at the top of the line. When the process continues in a similar way there will either be a **cycle** of existing tenants (as in TTC) who can swap their own houses or a **chain**  $(i_1, i_2, \dots, i_k)$  of agents where agent  $i_1$  is assigned an available house and each of the following agents is assigned the preceding agent’s house.

The resulting mechanism inherits the plausible properties of its “parents.”

**Proposition 7** (Abdulkadiroğlu and Sönmez 1999): The YRMH-IGYT mechanism is strategy-proof, ex-post Pareto efficient, and individually rational (in the sense that no existing tenant receives a house inferior to her own).

## Kidney Exchange

Living donors are an important source of kidneys for transplantation. But a patient with a willing living donor may not be able to receive a transplant because of an immunologic incompatibility between her and her donor. Recently transplant centers around the world developed **pairwise kidney exchange** programs where two such pairs exchange donors in case donor in each pair is compatible with the patient in the other. Another interesting program is **indirect kidney exchange** where the patient of an incompatible pair receives priority in the deceased donor waitlist if her incompatible donor donates a kidney to that waitlist. Motivated by these practices, Roth, Sönmez and Ünver (2004) introduce the following **kidney exchange** model: There are a number of patients each with a (possibly) incompatible donor. For each patient a subset of donors can feasibly donate a kidney and the patient has strict preferences over these donors and his own donor (who may or may not be compatible with him). In addition to ranking all compatible donors, each patient also ranks a “waitlist option” which represents trading his donor’s kidney with a priority in the waitlist. An **allocation** in this context is a matching of patients and donors such that:

1. each patient is matched with either a donor or the waitlist option, and
2. each donor can be matched with at most one patient while the waitlist option may be matched with multiple patients.

(The donors who remain unmatched are offered to the waitlist in exchange for the equal number of priorities awarded by the allocation). We are only interested in **individually rational** allocations where patients neither receive a donor nor the waitlist option unless it is indicated to be at least as good as his donor’s kidney. If the waitlist option is ranked inferior to his donor for a patient, that means the patient is not interested in such an exchange. As in the case of house allocation with existing tenants model, an allocation consists of cycles and chains where

1. each patient in a cycle receives a kidney from the donor of the next patient in the cycle, and
2. all but the last patient in a chain receives a kidney from the donor of the next patient in the chain whereas the last patient in the chain receives a priority in the waitlist.

If the waitlist option is infeasible, then the resulting problem is formally equivalent to a housing market and therefore has a unique allocation in the core which can be obtained via the TTC algorithm. In this simpler model an allocation (including the one in the core) consists of only cycles. When the waitlist option is feasible an allocation can also have chains (which are indirect exchanges and their more elaborate versions). In this more general model Roth, Sönmez and Ünver (2004) introduce a class of **Top Trading Cycles and Chains** (TTCC) algorithms each of which extend the TTC. Among these algorithms Roth, Sönmez and Ünver (2004) identify one that is Pareto efficient and strategy-proof:

**Proposition 8** (Roth, Sönmez and Ünver 2004): There exists a TTCC mechanism that is Pareto efficient and strategy-proof.

The deferred acceptance algorithm (for two-sided markets) also has some uses in one-sided allocation problems in which children are to be allocated to schools, if the schools, although not active strategic players, have priorities over students that need to be treated like preferences (Abdulkadiroglu and Sonmez 2003).

### **III. Design and engineering**

#### **III.A. Introducing a centralized stable match**

Of the several dozen markets and submarkets we know of that established clearinghouses in response to unraveling in a (two-sided) labor market, those that produce stable matchings have been more successful. Of particular note in this regard are the markets used in the various regions of the British National Health Service. In the 1960's, these markets suffered from the same kind of unraveling that had afflicted the American medical market in the 1940's. A Royal Commission recommended that each region organize a centralized clearinghouse, and the various regions each invented their own matching algorithms, some of which were stable and some of which were not (an example of such unstable algorithms will be given later). Those clearinghouses that produced stable matches succeeded, while those that did not most often failed and were abandoned. But over a broad range of markets, the correlation between stability and success in halting unraveling isn't perfect; some unstable mechanisms remain in use, and some stable mechanism have occasionally failed, as we will discuss later. And there are other differences between markets than the way their clearinghouses are designed. This is why, in order to establish that producing a stable outcome is an important feature for the success of a match, controlled experiments in the laboratory can be informative.

The laboratory experiments reported by Kagel and Roth (2002) help to verify the influence of a stable or unstable matching mechanism. After unraveling had begun in a small laboratory market, a clearinghouse was introduced using either the stable deferred acceptance algorithm or the unstable algorithms that failed in various regions of the British National Health service (Roth, 1991). In the lab, as in the field, participants learned to wait for and use the stable algorithm, but learned to arrange their matches early and thus avoid using the unstable algorithm. Note that a laboratory market is quite different from a naturally occurring labor market, but it has the advantage that it allows the effect of the different algorithms to be observed in an environment in which everything else is the same.

Centralized clearinghouses that yield stable outcomes have sometimes been introduced to organize markets suffering from failures other than unraveling, but related to congestion or the safety of revealing private information.

Examples of algorithms that produce unstable outcomes, but have been used in a number of market clearinghouses, are so called priority algorithms, used e.g. by some British clearinghouses, and also in several school choice problems in the US. A priority algorithm classifies different matches in terms of priorities, based on the rank orders submitted, and then makes feasible matches in order of priority. In Boston, for example,

the centralized system attempted to give as many students as possible their first choice school. The difficulty with the system was that someone who did not get assigned to their first choice was much less likely to be assigned to the school they had listed as their second choice than they would have been if they had listed it as their first choice, since those schools often get filled by students who list them as their first choice. This means participants have strong incentives to not report their preferences truthfully, if there is a good chance that they would not be admitted to their true first choice school; it might be wiser to list their second choice school as their first choice. The newly adopted Boston clearinghouse fixes this problem using a deferred acceptance algorithm (Abdulkadiroglu et al 2005, 2006).

Some markets manage to halt unraveling, but still suffer from congestion. The market for clinical psychologists (before it reorganized through a modified deferred acceptance algorithm see Roth and Xing 1997) and the match of students to New York City high schools before it was redesigned (Abdulkadiroglu et al 2005, 2006) are good examples. Clinical psychologists tried to run a deferred acceptance algorithm over the phone in the course of a day, “match day”, from 9:00 AM to 4:00 PM. All offers had to remain open until 4 PM, and students were supposed to hold only one offer at a time. Even though turnaround time in this market was very fast (offers took about 5 minutes, rejections about 1 minute), simulating a deferred acceptance algorithm in real time, for a market with about 2000 positions in 500 programs, takes much longer than the 7 hours of match day. (And making the market longer may increase the effects of congestion, if it means that participants can no longer stay by the phone for the whole market, so that the time for an offer to be made and rejected becomes disproportionately longer.) Congestion is an issue whenever a large number of offers have to be made. The system used to assign students to New York City high schools used to be carried out through the mails, and over 30,000 students a year were “stranded” on waiting lists and had to be assigned to a school for which they had expressed no preference. The new New York City clearinghouse is able to process preferences quickly, and in the years since its adoption, fewer than 3,000 students have had to be assigned each year to a school for which they had expressed no preference.

### **III.B. What are the effects of a centralized match?**

Centralized clearinghouses can help make markets thick and uncongested, and avoid unraveling. Studying their effect on various markets can also help us understand how clearinghouses, and the timing of the market (e.g. how far a labor market operates in advance of employment), influences the outcome of the market in other respects. For example, the market for gastroenterology fellows provides us with a natural case study of the effects of a clearinghouse not only on hiring practices (namely the timing of the market, and the kinds of offers that are made), but also employment opportunities, job placement, and the potential impact on wages.

Gastroenterology fellows are doctors who have completed three years of residency in internal medicine, and are now employed in a fellowship that will result in their becoming board certified subspecialists in gastroenterology. The market in which

gastroenterology fellows are hired operated in a decentralized way for many years, and experienced the problems of congestion, unraveling and exploding offers, as described above in connection with the market for medical residents. In 1986, various internal medicine subspecialties organized a clearinghouse called the Medical Specialties Matching Program (MSMP), sponsored by and organized along the same lines as the NRMP (which operates the resident match). But in the mid 1990's, gastroenterology fellowship programs, and applicants, started to defect from the match, and the gastroenterology market again unraveled. A match was successfully re-established only in 2006 (Niederle, Proctor and Roth, 2006). In those intervening years, as the market unraveled, the national market broke up into more local markets (Niederle and Roth 2003b). Fellowship programs, particularly smaller ones, had a larger tendency to hire their own residents than under a centralized match.

A second aspect of the outcome that received prominence in 2002, is the question whether a match affects wages: An anti-trust lawsuit against the NRMP and numerous other defendants was brought in 2002 by 16 law firms on behalf of 3 former residents seeking to represent the class of all former residents (and naming as defendants a class including all hospitals that employ residents).

Niederle and Roth (2003a) showed empirically that in fact there is no difference in wages between medicine subspecialties that use a match and those that don't. The suit was dismissed in 2004 following legislation intended to clarify that the medical match is a marketplace and does not violate antitrust laws.

The theory of the complaint was that a match holds down wages for residents and fellows. Bulow and Levin (2006) present a very stylized theoretical model providing some logical support for this possibility, by comparing a market with impersonal prices (to present the NRMP) to perfectly competitive prices at which each worker is paid his marginal product. Subsequent theoretical papers have shown that the conclusion about wage suppression doesn't necessarily follow if the model is expanded to include the possibility of firms hiring more than one worker (Kojima 2007), or when the model incorporates the actual procedures by which the medical match is conducted (Niederle 2006).

### **III.C. Beyond centralized matching: Why do some markets work well, while others do not?**

We have seen that stability is an important feature for a centralized match to remain in use. However, the history of the gastroenterology market shows that producing a stable outcome is not sufficient to guarantee a successful clearinghouse. For a centralized match to work well, participants need to have incentives to participate in the match. McKinney, Niederle and Roth (2005) observed that the collapse of the gastroenterology fellowship match seems to have been caused by an unusual shock to the supply of highly qualified gastroenterology fellows, a kind of shock that was not observed in other internal medicine subspecialties that continued to use a match. Furthermore, market conditions

seemed to have stabilized, so that a centralized match would work well once again, if it could be successfully reinstated.

However, many gastroenterology fellowship programs were concerned that, while they were willing to refrain from making the early offers that had become customary, and wait for the match, their main competitors would continue to make early exploding offers to promising applicants. Such concerns could effectively prevent a successful restart of a centralized clearinghouse.

This raises the more general question as to why some markets unravel and experience congestion problems in the first place, while others do not. Empirically, most markets that experience congestion also experience that employers (hospitals, federal judges, colleges...) make short term offers, with a binding deadline, and in which the acceptance of an offer is often effectively binding (Niederle and Roth 2006, for descriptions in the markets for law graduates, and for college admissions, see e.g. Avery, Jolls, Posner and Roth 2001, and Avery, Fairbanks and Zeckhauser 2003).

On the other hand there are markets that do not unravel, such as the market for graduate school admission. In this market, a policy (adopted by the large majority of universities) states that offers of admission and financial support to graduate students should remain open until April 15. Furthermore, a student faced with an earlier deadline is explicitly encouraged to accept this offer, and in case a better one is received before April 15, to renege on that former acceptance. This of course makes early exploding offers much less attractive. Niederle and Roth (2006) explore environments in which either eliminating the possibility of making exploding offers, or making early acceptances non-binding, helps prevent markets from operating inefficiently early.

These insights were used to help reorganizing the gastroenterology fellowship match. To reduce the concerns of programs that their competitors would start making exploding offers before the match, a resolution was adopted by the four main Gastroenterology professional organizations that stated that acceptances made before the match were not to be considered as binding, and such applicants could still change their minds and participate in the match.

#### **IV. Directions for future research**

As economists' understanding of the matching function of markets increases, and as economists are more often called upon to help design markets, one challenge will be to understand better how decentralized markets work well or badly, and not only in the final hiring stages.

For example, a common problem in many entry level labor markets (and in dating and marriage markets) is that participants do not have well formed preferences over potential matching partners, and forming those preferences is often very costly. For example, in the market for assistant professors, Economics departments receive hundreds of applications for any position, but in general only interview about 30 candidates at the annual winter

meetings. From among those they interview, they must decide whom to fly out for extended campus visits and seminars, and it is from among this latter set of candidates that they eventually choose to whom to make an offer. Because this is a time consuming and costly process, and because it is difficult to identify promising candidates after the initial interviews and flyouts (after which many candidates will have already accepted positions) many departments have to take care to interview applicants who have a good chance not only of receiving their offer, but of accepting it. This often amounts to a coordination problem: not all departments should interview the same applicants. Allowing applicants to credibly submit information about their interest in particular schools can help alleviate this coordination problem, and in 2007 the American Economic Association implemented a signaling mechanism of this sort in the market for economists.

In general, the study of the matching function of markets has directed attention at the design of rules and procedures of both centralized and decentralized markets. The goal of the growing interest among economists in matching and market design is to understand the operation of markets, both centralized and decentralized, well enough so we can fix them when they're broken.

### **References and further reading**

1. Abdulkadiroglu, Atila, Parag A. Pathak, and Alvin E. Roth, "The New York City High School Match," *American Economic Review, Papers and Proceedings*, 95,2, May 2005, 364-367.
2. Abdulkadiroglu, Atila, Parag A. Pathak, and Alvin E. Roth, "Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," December, 2006.
3. Abdulkadiroglu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez, "The Boston Public School Match," *American Economic Review, Papers and Proceedings*, 95,2, May 2005, 368-371.
4. Abdulkadiroglu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sonmez, "Changing the Boston School Choice Mechanism," May 2006.
5. A. Abdulkadiroglu and T. Sönmez, Random serial dictatorship and the core from random endowments in house allocation problems, *Econometrica* 66 (1998), 689-701.
6. A. Abdulkadiroglu and T. Sönmez, House allocation with existing tenants, *J. Econ. Theory* 88 (1999), 233-260.
7. Abdulkadiroglu, Atila and Tayfun Sönmez. 2003. "School choice: A mechanism design approach." *American Economic Review*, 93 pp. 729-747.
8. Avery, Christopher, Andrew Fairbanks and Richard Zeckhauser, *The Early Admissions Game: Joining the Elite*, Harvard University Press, Cambridge, MA 2003.
9. Avery, C., Jolls, C., Posner, R. A., and Roth, A. E. (2001), "The Market for Federal Judicial Law Clerks," *University of Chicago Law Review*, 68:793-902.

10. Bulow, Jeremy and Jonathan Levin. 2006. "Matching and Price Competition." *American Economic Review*, 96(3), 652 - 668.
11. Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 69, 9-15.
12. Hatfield, John, and Paul Milgrom. 2005. "Matching with Contracts." *American Economic Review*, 95(4), 913-935.
13. A. Hylland and R. Zeckhauser, The efficient allocation of individuals to positions, *J. Polit. Econ.* 87 (1979), 293-314.
14. Kagel, J.H., Roth, A.E., 2000. The dynamics of reorganization in matching markets: a laboratory experiment motivated by a natural experiment. *Quarterly Journal of Economics* 115 (1), 201-235.
15. Kelso, Alexander S., and Vincent P. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50(6), 1483 - 1504.
16. Kojima, Fuhito, "Matching and Price Competition When Firms Can Hire More Than One Worker" forthcoming, *American Economic Review*.
17. Li, Hao, and Rosen, Sherwin. "Unraveling in Matching Markets." *American Economic Review*, 88 (June 1998), 371-87.
18. Ma, Jinpeng, "Strategy-Proofness and the Strict Core in a Market with Indivisibilities" *International Journal of Game Theory*, 1994(23), 75-83.
19. Milgrom, Paul, *Putting Auction Theory to Work*, Cambridge University Press, 2004.
20. McKinney, C. Nicholas, Muriel Niederle, and Alvin E. Roth. 2005. "The collapse of a medical labor clearinghouse (and why such failures are rare)." *American Economic Review*, 95(3), 878-889.
21. Niederle, Muriel, "Competitive Wages in a Match with Ordered Contracts", December 2006. <http://www.stanford.edu/~niederle/Niederle.OrderedContracts.pdf>
22. Niederle, Muriel, Deborah D. Proctor and Alvin E. Roth. 2006. "What will be needed for the new GI fellowship match to succeed?" *Gastroenterology*, 130, 218-224.
23. Niederle, Muriel, and Alvin E. Roth. 2003a. "Relationship Between Wages and Presence of a Match in Medical Fellowships." *JAMA, Journal of the American Medical Association*, 290(9), 1153-1154.
24. Niederle, Muriel, and Alvin E. Roth. 2003b. "Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match." *Journal of Political Economy*, 111(6), 1342 -- 1352.
25. Niederle, Muriel, and Alvin E. Roth. 2006. "Making Markets Thick: How Norms Governing Exploding Offers Affect Market Performance."
26. Roth, A.E., "Incentive Compatibility in a Market with Indivisible Goods," *Economics Letters*, Vol. 9, 1982, 127-132.
27. Roth, Alvin E. 1984. "The evolution of the labor market for medical interns and residents: A case study in game theory." *Journal of Political Economy*, 92, 991-1016.
28. Roth, A.E., "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*, Vol. 81, June 1991, 415-440.
29. Roth, Alvin E. 2002. "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics." *Econometrica*, 70(4), 1341-1378.

30. Roth, Alvin E. and Marilda Sotomayor. 1990. Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis. Econometric Society Monograph Series, Cambridge University Press.
31. Roth, Alvin E. and Elliot Peranson. 1999 "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 89(4), 748 - 779.
32. Roth, Alvin E. and Andrew Postlewaite. 1977. "Weak versus strong domination in a market with indivisible goods." *Journal of Mathematical Economics*, 4, pp. 131-137.
33. Roth, Alvin E., Tayfun Sönmez and M. Utku Ünver, "Kidney Exchange," *Quarterly Journal of Economics*, 119, 2, May 2004, 457-488.
34. Roth, A.E. and X. Xing "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions," *American Economic Review*, 84, September, 1994, 992-1044.
35. Roth, A.E. and X. Xing "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists," *Journal of Political Economy*, 105, April 1997, 284-329.
36. L. Shapley and H. Scarf, "On cores and indivisibility", *J. Math. Econ.* 1 (1974), 23-37.
37. L. Shapley and M. Shubik, "The assignment game I: The core", *International Journal of Game Theory*, Vol.1, Number 1, December, 1971.