

Matching Theory and Its Applications

Jonathan Levin
Stanford University
March 2011

The theory of matching markets is the study of resource allocation problems in which two sets of agents, or a set of agents and a set of goods, have to be assigned to one another in a way that respects the preferences of the market participants. The theory differs from textbook price theory in that (a) there is an emphasis on heterogeneity --- in matching markets, every agent and good can be distinct from every other; (b) indivisibilities are treated as fundamental; and (c) it may be impossible to use money to adjust the terms of trade and clear the market. Of course, there is some blurring of boundaries and in places matching theory connects closely to traditional price theory, auction theory and search theory.

Models of matching originate in the work of David Gale, Lloyd Shapley, Martin Shubik and Herbert Scarf in the 1960s and 1970s. The theory they began turns out to be both quite beautiful and highly applicable, with the applications being realized progressively over time, notably in the work of Al Roth and collaborators. Indeed, along with the theory of auctions to which it is closely related, matching theory can be viewed as one of the great practical success stories of modern economic theory.

These notes are organized with the theory discussed first and then applications. In describing the theory, I've tried to provide an accessible discussion of the major results and ideas, while conveying a taste of the mathematical logic. The second half of these notes describes applications to entry-level labor markets, school choice programs, kidney exchanges, internet advertising markets and other areas, and some of the insights that have emerged from these applications.

I. Development of Matching Theory

1. Matching Markets and the Deferred Acceptance Algorithm

In 1962, David Gale and Lloyd Shapley published a six and a half page paper in the *American Mathematical Monthly* titled “College Admissions and the Stability of Marriage”. With little mathematical notation but masterful clarity, they began the study of two-sided matching markets, and proved several beautiful and enduring results.

To understand their contribution, we can start with the simplest possible setting. Suppose there are N men and N women. Each man has a strict preference ordering over potential female partners. Each woman has a strict preference ordering over the men. Our goal is to match them into couples. Gale and Shapley argued that one criterion for a good matching is *stability*. A matching is stable if every person prefers his partner to being unmatched, and there is no man and woman who would prefer to be together than with their current partners, and hence might disrupt the matching.

Gale and Shapley proved that a stable matching always exists and can be found with a "deferred acceptance" algorithm. Imagine lining the men up against one wall, and the women against the opposite wall. Each man approaches his most preferred woman. Each woman then tentatively accepts her preferred suitor (if she has one) and sends the others back. The rejected men begin a new round with each approaching the next woman on his preference list. Again each woman tentatively accepts her preferred suitor, potentially rejecting a man she was holding from the prior round. The process continues until each man is paired with a woman, at which point matches are finalized.¹

If one thinks about this procedure for a few minutes, it is clear that it must terminate in no more than N rounds. Each round the set of women with tentative partners can only increase and at some point every woman must receive a proposal. Moreover, the final matching is stable. To see why, suppose a man inquires after the fact about a woman

¹ Here I'm assuming implicitly that all the men prefer being matched to some woman over being single, and the same holds for the women. In this case, every woman must eventually receive a proposal and with equal numbers of men and women, the algorithm will terminate with all participants matched. The algorithm and the results extend straightforwardly to the case where some people find others unacceptable, or to a situation with *unequal* numbers of men and women. In these cases, some individuals can be left unmatched and the algorithm will terminate either when all men are matched or when the unmatched men have run completely through their preference lists.

he prefers to his assigned partner. Because men move sequentially down their preference list during the algorithm, he must have already proposed to this same woman earlier and been rejected. At the point he was rejected, the woman must have had a tentative partner she preferred and her final partner must be either this same man or someone she likes even better. So if she receives the inquiry, she will not want to break her existing match.

Theorem 1. (Gale and Shapley) *In a marriage market where men and women have strict preferences, a stable matching always exists and can be found using the deferred acceptance algorithm.*

Now, if we go back to the original setting, we can see that in general there may be multiple stable matchings. Suppose there are two men m and m' , and two women w and w' . Man m prefers woman w , and man m' prefers w' , but that the women have the opposite preferences: w prefers m' and w' prefers m . There are two stable matchings. One possibility is to match (m, w) and (m', w') . The matching is stable because both men are with their first choice partner. But matching (m, w') and (m', w) is also stable because then both women are with their first choice partner.

Gale and Shapley's algorithm with the men proposing leads to the first stable matching --- the one that the men prefer. If we reverse the roles and have the women propose, we end up at the second stable matching, the one the women prefer. Gale and Shapley prove that this is a general property. The man-proposing deferred acceptance algorithm always terminates at a stable matching that is *optimal* for the men --- in the sense that every man weakly prefers his partner from the algorithm to his partner in *any other stable matching*. The analogous result holds if we reverse roles and run the algorithm with the women proposing to the men.

What is striking about this is that the preferences of the men (and similarly of the women) are in a sense aligned. When we look at the set of stable matches, there is a unique matching that is (weakly) preferred by every man. Moreover, one can show that this man-optimal matching is also, among the set of stable matchings, the one that is (weakly) worst for every woman. In fact, both of these observations reflect a deeper mathematical property of the model. The set of stable matchings is a lattice, so that if you

take any two stable matchings and give each man his more preferred (or less preferred) woman from these matchings, you arrive at a new stable matching.²

Theorem 2. (*Gale and Shapley*) *In a marriage market where men and women have strict preferences, the man-proposing deferred acceptance algorithm finds the stable matching that is most preferred for all men, and least preferred for all women.*³

The theory of two-sided matching also leads to some other beautiful results. For instance, if a man or a woman is unmatched in some stable matching, he or she must be unmatched in every stable assignment.⁴ To see why, first note that the *number* of matches is identical in all stable assignments. The reason for this is that anyone matched in a stable assignment must prefer his or her partner to being unmatched. Now suppose man m is unmatched in some stable assignment. Then he must be unmatched in the woman-optimal (and man-pessimal) matching. Now the set of total matches cannot be larger in any other stable match because that would mean improving the outcome for some woman who is unmatched in the woman optimal stable match.

Another remarkable result relates to the incentives of the participants. Dubins and Freedman (1981) showed that if men and women are asked to submit their preferences lists to a computer that then assigns partners using the man-proposing deferred acceptance algorithm, it is a dominant strategy for the men to report their true preferences. So at least for the men, the mechanism has the appealing property that there is no need to engage in strategic calculations, and a relatively naive participant that reports truthfully does not suffer for being non-strategic.

Unfortunately, truth-telling is not a dominant strategy for the women. To see why, think about the example above with two stable matches. If the men propose, woman w would do best to declare man m unacceptable. If she does this, she'll reject m in the first

² Roth and Sotomayor (1990) attribute the lattice result to the mathematician John Conway. More recent work has shown that many of the properties of the deferred acceptance algorithm can be derived succinctly using results from lattice theory (see e.g. Hatfield and Milgrom, 2005).

³ The last part of this result is due to Knuth (1976); see also Roth and Sotomayor (1990).

⁴ There are various versions of this result, e.g. McVitie and Wilson (1970) and Roth and Sotomayor (1990).

round of the DA algorithm and subsequently match with m' when the algorithm terminates.

Roth (1982) showed that this incentive issue is in fact very general. It is impossible to create a mechanism for two-sided matching markets that is both strategy-proof (makes it a dominant strategy for men and women to both report their true preferences) and always results in a stable matching. One positive caveat to this result is that if the deferred acceptance algorithm is used in a large market, truth-telling becomes “almost” a dominant strategy for both sides (Immorlica and Mahdian, 2005; Kojima and Pathak, 2009).

Theorem 3. (Dubins and Freedman; Roth) *In a marriage market where men and women have strict preferences, the man-proposing deferred acceptance algorithm is strategy-proof for the men. But there is no mechanism that finds a stable matching and is strategy-proof for both sides of the market.*

These basic results on two-sided matching have been extended in a number of directions. One major extension, pioneered in the original Gale and Shapley paper, is to show that the two-sided matching model can be generalized to allow women to have multiple partners (or in their terminology, “colleges” to have multiple “students”). Roth and Sotomayor’s (1990) book describes many results for the “many-to-one” matching setting.⁵ The key assumption to ensure that stable outcomes exist and the deferred existence algorithm works turns out to be that colleges view students as “substitutes” rather than “complements”. For instance, one useful class of substitutes preferences arises if each college has a fixed potential class size and a rank ordering over students that it wishes to have in its class. A more problematic case would occur if a college wants to field an orchestra, but only if it can attract a sufficient number of strong musicians to play in the orchestra. This “complements” case is still the subject of active research.

2. “House Allocation” Problems and Top Trading Cycles

⁵ Some, but not all, of the results can also be extended to “many-to-many” matching. Recently, Michael Ostrovsky (2009) has shown that stable matches also exist in “supply chains” provided they maintain some of the same structure that is important in the many-to-one case.

A distinctive feature of the above setting is that participants on both sides of the market have preferences and the goal is to find a stable matching that respects these preferences. In a famous 1974 paper, Herbert Scarf and Shapley studied a closely related allocation problem in which only one side of the market has preferences. In this model, the goal is to match “people” with “houses”. The people have (strict) preferences over houses, but the houses do not themselves have preferences.⁶

If we start from a situation where there are no prior ownership rights, there is an obvious mechanism to allocate houses that is strategy-proof and results in an efficient allocation. (Here, I mean a strong form of efficiency that once the allocation is made there is no pair of individuals, or group of individuals, that can profitably trade --- i.e. that the resulting allocation is a *core allocation*.) Each individual is randomly assigned a priority, and then individuals select houses in order of their priority. This is the “random serial dictatorship” (RSD) mechanism.

Theorem 4. *In house allocation problems with no initial endowments, the random serial dictatorship is strategy-proof and Pareto efficient.*

It is not immediately obvious if an analogous mechanism can be created that respects prior ownership. For instance, if each individual starts with a house and has to choose whether to vacate it and participate in a RSD mechanism, an individual that has a house she likes, but nevertheless is not her first-best, might play it safe and opt out. So potential gains from trade may not be realized.

This situation, where each individual starts with a house, but there are potential gains from trade, is the one considered by Scarf and Shapley. They showed that a core allocation exists, and can be found using an elegant “top trading cycle” (TTC) algorithm. They attribute the algorithm to David Gale. It works as follows. Each individual points to his preferred house. Each house points to its owner. We look for cycles in the resulting directed graph. At least one cycle must exist, and perhaps many cycles. (To see why, start with an individual, move to the house to which he is pointing, then to the owner of the

⁶ Somnez and Unver (2010) is a recent survey that covers many results on the house allocation problem and also the two-sided matching model.

house, and so forth. Because there are a finite number of individuals, the resulting chain must form a cycle.⁷) Each person in a cycle is assigned to the house to which he or she is pointing, and these individuals and their assigned houses are removed. The process repeats with the remaining people and houses, and continues until everyone is assigned.

The top trading cycle algorithm has a number of remarkable properties. First, it leads to a core allocation, exploiting all possible gains from trade in the market. To see why, note that for a coalition to improve on the TTC outcome, each individual would have to “trade up” to a house she prefers. But this means each individual has to trade up to a house that was assigned at an earlier round in the algorithm than she received her TTC house --- an impossibility. Second, the mechanism is strategy-proof. Given the option to misrepresent their preferences, individuals have no incentive to do so.

Theorem 5. (Scarf and Shapley) *In a housing market where participants have strict preferences, there is a unique core allocation and it can be found with the Top Trading Cycles algorithm. Also, the TTC mechanism is strategy-proof for participants.*

So like the deferred acceptance algorithm, the top trading cycle mechanism provides a simple, constructive and attractive solution to a generic but seemingly complex allocation problem.

There is also an interesting extension of top trading cycles that connects it to random dictatorship. Abdulkadiroglu and Sonmez (1999) consider a variant of TTC that deals with the case where some individuals are endowed with houses but others are not, and some houses may be vacant. Each individual is assigned a priority. People then choose houses in order of their priority. If a person chooses a vacant house, he or she takes the house and exits. If a person chooses an occupied house, the owner moves to the front of the line and gets to select her preferred house (of the houses remaining). This continues until either someone chooses a vacant house, or a cycle is formed. If a cycle forms every in the cycle is assigned the house they have chosen and exits. We proceed to the next person and so on. This generalized version of top trading cycles nests both TTC

⁷ Of course one possibility is that a person’s most preferred house is the one he’s already in. In that case, he retains his house. Another possibility is that two people point to each other’s houses. In that case, they trade. It is also possible to have cycles of length three, four or more.

and random serial dictatorship --- and like these mechanisms, is both efficient and strategy-proof.

3. The Assignment Market and Connections to Auction Theory

A characteristic feature of the above models is that there is no money and no prices. This makes sense for certain applications but also sets the models apart from many standard analyses in economics. There are, however, close connections between matching markets and standard price theory. One of the pioneering papers in this regard is Shapley and Martin Shubik's (1971) paper on the assignment market. (An even earlier and seminal analysis of the assignment problem is Koopmans and Beckmann, 1957.)

In the assignment model, there is a finite set of agents and a finite set of indivisible goods. Agents have preferences over the goods, and in the simplest case want to purchase only a single good. In this case, we are interested in a one-to-one matching of matching individuals and goods. Unlike the models above, however, there is also money and preferences are quasi-linear, so it is possible to use prices to clear the market.

Shapley and Shubik develop some striking and beautiful properties of this model. They show, for instance, that the core is equivalent to the set of competitive equilibrium allocations. Interestingly, there are typically a range of competitive equilibria, because although all competitive equilibria involve the same (total-value-maximizing) assignment of goods to individuals, there are usually a range of supporting prices. To take the simplest example, suppose there is a single good, one agent with a value of two and another with a value of one. Then the first agent should buy the good, but any price between one and two clears the market. Note that in this example, there is both a maximum and minimum market-clearing price. This attractive property extends to the case of many goods. There is always a minimum and maximum vector of competitive equilibrium prices, and more generally the set of market-clearing prices is a lattice.⁸

⁸ Shapley and Shubik also characterize the core as the solution to a particular linear programming problem (also the approach taken by Koopmans and Beckman, cited above). The connections between competitive equilibrium, the game theoretic solution of the core, and constrained linear optimization runs through much of the subsequent work on assignment markets.

Theorem 6. (Shapley and Shubik) *In an assignment market, core allocations exist and coincide with the set of competitive equilibria. In addition, the set of competitive equilibria includes equilibria with (component-wise) maximal and minimal prices.*

As in the models discussed above, there are natural algorithms for finding core outcomes (or equivalently, market-clearing prices) in the assignment model. For instance, Crawford and Knoer (1981) propose a version of the deferred acceptance algorithm that works effectively as an auction. In their model, there are a finite set of salary levels that “firms” can pay to “workers”. The workers care only about salary, the firms have preferences over the workers as well as quasi-linear preferences over money. Each firm can hire at most one worker, so if there are N workers and L salary levels, the firm can obtain one of NL allocations in addition to getting nothing. Suppose the firm starts with a rank-order list of these outcomes.

We then run the deferred acceptance algorithm, with each firm proposing its best outcome (i.e. proposing the lowest salary to its preferred worker in the first period). Each worker tentatively accepts its best offer (randomly breaking ties if necessary), then rejected firms move to their next best offer – which might mean raising its offer to the same worker, or moving to a new worker. The outcome of this process is an efficient (surplus-maximizing) matching of firms to workers, with salaries equal to their lowest competitive equilibrium levels (modulo tie-breaking and discrete bid increments). In addition (again modulo tie-breaking and discrete bid increments), it is a dominant strategy for firms to behave truthfully in the algorithm.⁹

Theorem 7. *In the assignment market with unit demands, there exist auction mechanisms that are strategy-proof and lead to the minimum price competitive equilibrium outcome.*

⁹ Readers familiar with the theory of mechanism design will recall the Green-Laffont-Holmstrom Theorem, which says that in quasi-linear environments, any mechanism that is strategy-proof and leads to an efficient allocation must be equivalent to a Vickrey-Groves-Clarke mechanism. Indeed, with unit demands, the minimal competitive prices are exactly the prices that emerge from a Vickrey auction. Demange, Gale and Sotomayor (1986) discuss algorithms to find the lowest competitive prices, bidder incentives in the auction and the connection to Vickrey pricing.

The analysis also generalizes to the case where the firms can match with multiple workers (or equivalently the “bidders” can buy multiple “goods”). The key condition, identified by Kelso and Crawford (1982), is that each bidder view the goods as “substitutes,” meaning that an increase in the price for good j weakly raises the bidder’s willingness to pay for good i . In the assignment market with multi-good demands, the substitutes condition ensures the existence of market clearing item prices, and in a certain sense is also necessary for such prices to exist (Gul and Stacchetti, 1999; Milgrom, 2000).

Kelso and Crawford extend the deferred acceptance algorithm to the case where firms want multiple workers. Hatfield and Milgrom (2005) provide a more general analysis that nests both the two-sided matching market model with no money (the Gale-Shapley model) and the assignment model with quasi-linear preferences and money (the Shapley-Shubik model), and in fact allows for general forms of “contracts” between firms and workers.

This work also turns out to have a close connection to auction theory. Starting in the 1990s, there has been a surge of interest in multi-good auctions, spurred in part by the auctions for radio spectrum pioneered by the U.S. Federal Communications Commission. The standard auction used for selling radio spectrum is the simultaneous ascending auction. Roughly, the auction consists of a series of rounds, and prices rise from round to round until there is no excess demand for any of the items.

Milgrom (2000) shows that (subject to some small caveats about discrete price increments), this auction will end at the lowest market clearing prices so long as bidders have substitutes preferences and bidding is “straightforward” --- that is, bidders respond to prices by bidding according to their true demands. The idea is that at the beginning of the auction, all goods will be in excess demand. If the price of good i increases, the demand for all goods $j \neq i$ will weakly increase, so as long as the demand for i does not “undershoot” supply, we will continue to have weakly excess demand for all goods, until prices rise enough that the market just clears.¹⁰

¹⁰ Milgrom (2000) shows an “approximate” result with discrete price increments. Gul and Stacchetti (2001) look in at how exact market clearing can be obtained with continuous price increases. Note that the straightforward behavior of bidders is assumed. If bidders have multi-unit demands, they generally will have an incentive to bid less than true demand in order to reduce the price they pay for inframarginal units.

Most of the positive results about assignment markets, starting with those of Shapley and Shubik, depend on individuals wanting only a single good, or viewing goods as substitutes. The situation becomes more complicated if bidders want to assemble packages of items that they view as complementary. This remains a frontier problem in auction theory. Milgrom (2007) provides a recent assessment of the state of the art.

D. Conclusion on the Theory

In this section I have described three related types of allocation problems: two-sided matching markets where prices cannot be used to adjust the terms of trade; markets without money where we desire to assign or re-allocate indivisible goods; and assignment markets where we want to find prices that clear markets for potentially distinct and heterogeneous indivisible goods.

What links the analysis of these problems is that in each case we started with a set of criteria for what constitutes a “good “ solution to the problem --- e.g. a stable or core outcome, no incentive for participants to behave strategically, etc. --- and demonstrated that practical mechanisms exist that provide such a solution. In the next section, I describe a broad set of real-world applications in which allocation problems arise that correspond to the abstract models described above, and in which the mechanisms considered here, or variants of those mechanisms, have been put to work.

II. Real-World Applications of Matching Theory

For twenty years after Gale and Shapley’s initial paper, research on matching markets was essentially theoretical in nature. In retrospect, it seems fair to say that the field did not get that much attention from economists. This has changed in recent years and a main reason for the change has been the realization that the theory is not only mathematically appealing but has many practical and important real world applications.

1. The National Residency Matching Program

One of the origins of recent work applying matching theory to real-world markets is Alvin Roth's 1984 *Journal of Political Economy* paper on the National Residency Matching Program. In the United States, the NRMP is a centralized clearinghouse that each year handles the task of assigning graduating medical students to residency positions in hospitals. In the fall, students apply to hospitals and are interviewed. Each student then submits to the clearinghouse a rank order list of his or her desired residency positions and the hospitals submit corresponding preference lists of students. The clearinghouse uses an algorithm to match the students to positions and on "Match Day" announces the results.¹¹

Roth (1984) pointed out that the clearinghouse, since 1951, had used essentially the Gale-Shapley deferred acceptance algorithm to assign residents to hospitals. While the independent discovery of a stable matching algorithm is itself fairly remarkable, the history makes it even more interesting. Prior to the clearinghouse being adopted, medical students had found positions informally through a decentralized process. Over time, however, the market unraveled and students were accepting positions years before finishing medical school. The adoption of a centralized match appeared to solve this unraveling problem.

In a series of papers, Roth argued theoretically, empirically and experimentally that the stability of the clearinghouse and its success in stopping market unraveling was related to the stability of the matching mechanism. In one study, Roth (1990) examined residency matches in different cities in the United Kingdom. Some used versions of the deferred acceptance algorithm, or "stable mechanisms", while others used mechanisms such as "priority matching" that need not result in a stable matching. In virtually every case, the stable mechanisms had remained in use, while the unstable processes had broken down, in many cases because hospitals and residents were forming advance agreements.

Roth and Xing (1994) expanded the argument by noting that many entry-level labor markets, not just the medical residency market, have a tendency to "unravel". Two salient examples in the United States are early offers made by judges to prospective law

¹¹ An apparently unusual feature of this market is that it does not allow much leeway for negotiations about salaries and benefits. In the last decade, this became the impetus for an antitrust case filed by residents. Bulow and Levin (2006) present a model of equilibrium wage setting in markets such as the NRMP where salaries are attached to positions, rather than being negotiated separately for individual workers filling those positions. They argue that equilibria in such markets can be quite efficient, although wages may be low.

clerks (Avery, Jolls, Posner and Roth, 2000), and the use of early admission programs by selective colleges and universities (Avery and Levin, 2010). A particularly nice feature of the Roth and Xing paper is that although the paper is motivated by a specific set of empirical examples, and provides enormous detail on these case studies, the discussion opens up a broad line of theoretical inquiry into the timing of market clearing, an area which even now has not been fully explored.

Subsequently, Roth became involved personally in the NRMP in the mid-1990s when he assisted in re-designing the algorithm to better handle the increasing prevalence of married couples that sought to be matched together. Roth and Peranson (1998) describes the re-design. One of their interesting observations during the re-design was that in simulations of the residency market, the number of stable matches appeared to be very small. This is important because theoretical work connects the incentives to misreport preferences to the number of stable matches. Roth and Peranson hypothesized that in large markets, the set of stable matches might be “generically” small and temptations to misreport preferences minimal. This idea is developed in the papers by Immorlica and Mahdian (2005) and Kojima and Pathak (2010) mentioned earlier.

2. *School Choice*

Another application of matching theory that has been particularly high-profile in the United States has been to school choice in urban school districts. In many school districts, children are asked to submit their preferences as to the elementary or high school they would like to attend. Children may have priority at certain schools based on the location of their home, whether their siblings already attend the school, their grades, or other factors, but the assignments can depend on whether the school they want to attend is in excess demand.

In an influential article in the *American Economic Review*, Abdulkadiroglu and Sonmez (2003) observed that this problem could be viewed as a two-sided matching problem, in which one side (students) had preferences and the other might or might not (depending on the extent to which schools “prioritized” students). They also pointed out that the common mechanisms used by cities appeared to have some undesirable

properties. In particular, they did not necessarily encourage parents to truthfully report their preferences and did not obviously lead to stable or Pareto efficient assignments.

Abdulkadiroglu and Sonmez (2003) proposed instead that cities use either the deferred acceptance algorithm or a version of top trading cycles. The former generates a stable allocation, which in a setting with school priorities can be re-interpreted as an allocation with “no justified envy”. (An allocation has no justified envy if there is no student who would prefer school A to his own school, and who has a higher priority at school A than some student assigned there.) If students can trade seats, however, the deferred acceptance algorithm might leave students with potential gains from trade. In contrast, the top trading cycles leads to an outcome that is Pareto efficient for students.

The article attracted immediate attention from several U.S. cities. Abdulkadiroglu, Pathak, and Roth (2005) and Abdulkadiroglu, Pathak, Roth and Sonmez (2005) describe subsequent re-designs of the school choice programs in New York and Boston. In both cases, the cities adopted the deferred acceptance algorithm. Early evidence suggests that both re-designs have been successful, particularly in New York in increasing the number of students matched during the principal phase of the matching process.

In addition to making an important practical contribution, the application of deferred acceptance to school choice has spurred new economic theory. Several papers have looked at the matching mechanism previously used in Boston and other cities. In the Boston mechanism, students submit rank-order lists of schools. The algorithm tries to assign all students to their first choice school. If a school is under-demanded, all students choosing it get a seat. If a school is over-demanded, the seats are assigned by priority. Rejected students then apply to their second choice school, and so on. The key difference relative to deferred acceptance is that if a student gets a seat in a given round, he or she cannot be displaced even if later the school could get a student that has higher priority.

The Boston mechanism is not strategy-proof. A student who wants to attend a highly demanded school but would be willing to attend another school where he or she has higher priority might do better to rank the second school first, to avoid the risk of missing out entirely on one of her top two schools. Pathak and Sonmez (2008) point out that this is potentially unfair, because “sophisticated” students may be able to benefit relative to “naive” students from strategic preference reporting. Abdulkadiroglu, Che and

Yasuda (2009), however, argue that the Boston mechanism has some advantages because it induces students to make choices that reflect the *intensity* of their preferences and not just their ordinal ranking. For example, a student that vastly prefers his or her top-ranked school is less likely to play it safe than one who is almost indifferent between her first and second choice.

Another new development in the theory relates to tie-breaking. In the models described above, we assumed that preferences were strict. This has been the standard assumption in matching theory (Roth and Sotomayor, 1990), and although it seems innocuous, it is not. In cities such as Boston, the schools assign students into broad priority classes. In running the deferred acceptance algorithm, there can be ties where two students with the same priority seek a single seat. Ergin and Erdil (2008) point out that if one breaks these ties randomly, it may be possible, after the algorithm is complete, to make trades that preserve stability but are Pareto improving for the students. But if we augment the deferred acceptance algorithm to implement such trades at the end, it alters student incentives. So with ties, there is a trade-off between getting a student-optimal stable matching and having a strategy-proof mechanism.

3. Kidney Exchanges

In another application of the theory, Roth and collaborators have pioneered the use of matching algorithms to improve the process for organ donations in the United States. The motivating problem is that in the United States there is a large shortage of organ donors and hence a long waitlist for cadaveric donations. Some patients are able to find willing donors, but because of blood and tissue type incompatibilities, not every donor is suitable for every patient.

The specific pattern of incompatibilities suggests that it might be desirable institute exchanges, with one patient's donor donating to another patient, and the second patient's donor donating to the first patient. Longer chains are also possible. For instance, there are four common blood types: 0, A, B, and AB. Patients with blood type 0 can only accept kidneys from donors with blood type 0. Patients with blood type A can accept a type 0 or a type A kidney, while a patient with blood type AB can accept any kidney (subject to tissue type compatibility).

Now suppose that a patient with blood type 0 is able to find an incompatible type A donor. Another patient with blood type A finds an incompatible AB donor. And a type AB patient finds a (compatible) type 0 donor. Absent an exchange, there can be exactly one successful transplant. But with a three way exchange, the type 0 patient can receive the type 0 kidney, the A patient the type A kidney, and the AB patient the type AB kidney. Three patients receive transplants!

In a first analysis of this problem, Roth, Sonmez and Unver (2004) connected the problem of finding efficient kidney exchanges to the housing economy analyzed by Shapley and Shubik. They suggested that if a clearinghouse could be formed for patients and their donors, a Top Trading Cycles algorithm could be used to find efficient exchanges. In a subsequent 2005 paper, they realized that the pattern of incompatibilities permitted some simplifications, stemming from the fact that compatibility could be treated as a yes/no variable rather than as a continuous probability of success. They then showed that substantial progress could be made using only pairwise exchanges.

In the last few years, these articles and follow-on work by the authors and medical collaborators have helped spur the creation of at least two regional clearinghouses in the United States (in New England, and Ohio), and may ultimately lead to a national kidney exchange. Roth (2007) contains an interesting discussion of the ethical issues involved in these exchanges, and some of the objections to compensating donors in order to relieve the shortage of kidneys available to transplant patients.

4. Internet Advertising Auctions

One of the most economically significant applications of matching theory in recent years has been to new auction markets. Perhaps the most famous example is Google's matching auction for sponsored search keywords. Remarkably, Google generates most of its over \$25 billion in annual revenue from running billions of auctions a week, each auction generating on average a few pennies. The same approach is used by Microsoft's Bing search engine and the Chinese search engine Baidu.

In internet search, users enter queries and search engines return two types of search results. One set of results is generated by mining data about the underlying link structure of the web and consumer browsing behavior. But another set of results (on

Google, displayed down the right hand side of the page and sometimes at the very top), are paid placements and these are allocated via auction. Advertisers are willing to pay for these positions because user queries can often signal a strong intention that they are interested in a product or service --- e.g. a life insurance company would be eager to have its ad seen by a consumer who had just entered “life insurance” into Google.

In two parallel papers, Edelman, Ostrovsky and Schwarz (2007) and Varian (2007) show that the problem of allocating positions on an internet search results page can be mapped into the Shapley-Shubik analysis of the assignment market. As Edelman et al. and Varian model the problem, there are N positions that can be ranked in terms of the attention and clicks they will receive. There are N advertisers who can be ranked in terms of their value per click. The efficient (surplus-maximizing) assignment gives the high value advertiser the highest position and so forth.

From the assignment market analysis discussed above, we know that there are market clearing prices for the positions, that the competitive allocations are efficient (and hence assortative) and that they coincide with the core. In this case, they also involve increasing prices --- the position that garners the most clicks must have the highest price (on a per-click basis). Moreover, from the theoretical discussion above, we also know that there are a variety of auction designs that would yield these prices. One example would be a Vickrey-Clarke-Groves auction.

It turns out Google and the other search auctions use a novel auction design --- a “generalized second price auction”. In this auction, each advertiser submits a “per click” bid --- a maximum amount the advertiser would be willing to pay if its ad is clicked on. The bids are ranked in descending order, bidder k is assigned to position k and if the consumer clicks on ad k , the advertiser pays the per-click bid of the $k+1^{\text{st}}$ bidder. This auction does not give the bidders an incentive to bid “truthfully” but it does have a natural Nash equilibrium that results in each advertiser being assigned to its efficient position, and paying (in expectation) the lowest market clearing price for that position.

Apart from the massive scale on which search auctions are run, what is neat about these auctions is that, like the Residency Match, they evolved organically through a series of incremental steps to a mechanism that accords with the ideas coming from the abstract economic theory. In the case of sponsored search auctions, the original market for

internet advertising involved bidders paying directly for positions, this was modified to paying for clicks which established a natural “conversion rate” between positions. And the auction itself went from being a “pay-your-bid” auction which does not have a pure strategy Nash equilibrium to a “second-price” auction which has equilibria that correspond to the competitive equilibria of the associated assignment problem.

5. Further Applications and the Field of Market Design

I have touched on a few notable applications of matching theory, but there are many others. Versions of the deferred acceptance algorithm are used in a variety of entry-level labor markets, in some fraternity and sorority rushes, and in at least one online dating website. House allocation problems arise in assigning students to courses with limited seats, or (not surprisingly!) to scarce housing. Doubtless there are many other applications of which I am not aware, or have yet to be discovered.

Insights from the assignment model also have been important in recent auction applications. As mentioned earlier, some of the theory supporting the simultaneous ascending auction used to allocate radio spectrum is based on the idea that the auction works as an algorithm for discovering market clearing prices in an assignment setting. Related ascending or descending price auctions also have been used in selling electricity generation units, emissions permits, rough diamonds, and other commodities (Milgrom, 2004). Klemperer (2009) and Milgrom (2009) discuss applications of sealed bid assignment auctions to Central Bank auctions, electricity markets and other settings.

More generally, the applications described above have spurred the creation of what Roth has called an "engineering" branch of economics that has applied the tools of matching theory and auction theory to the design and improvement of a broad range of economic markets (Roth, 2002, 2008; Wilson, 2002; Milgrom 2004). For the last ten to fifteen years, this area of "Market Design" has been one of the most active and exciting areas in economic, bridging modern theory and practical application.

IV. References

- Abdulkadiroglu, Atila, Yeo-Koo Che and Yosuke Yasuda, "Resolving Conflicting Preferences in School Choice: the "Boston" Mechanism Reconsidered," *American Economic Review*, forthcoming.
- 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.
- Abdulkadiroglu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sonmez, "The Boston Public School Match," *American Economic Review, Papers and Proceedings*, 95,2, May, 2005, 368-371.
- Abdulkadiroglu, Atila and Tayfun Sonmez, "House Allocation with Existing Tenants," *Journal of Economic Theory*, 1999, 88, 233-260.
- Abdulkadiroglu, Atila and Tayfun Sonmez, "School Choice: A Mechanism Design Approach," *American Economic Review*, 2002, 93, 729-747.
- Avery, Christopher, Christine Jolls, Richard A. Posner, and Alvin Roth, "The Market for Federal Judicial Law Clerks" *University of Chicago Law Review*, 2001, 68(3) 793-902.
- Avery, Christopher and Jonathan Levin, "Early Admissions at Selective Colleges," *American Economic Review*, 2010, 100(5), 2125-2156.
- Bulow, Jeremy and Jonathan Levin, "Matching and Price Competition," *American Economic Review*, 2006, 96, 652-668.
- Crawford, Vincent P. and Elsie Marie Knoer, "Job Matching with Heterogeneous Firms and Workers," *Econometrica*, 1981, 49, 437-450.
- Demange, Gabrielle, David Gale and Marilda Sotomayor, "Multi-Item Auctions," *Journal of Political Economy*, 1986, 94(4), 863-872.
- Dubins, L. E. and D.A. Freedman, "Machiavelli and the Gale-Shapley Algorithm," *American Mathematical Monthly*, 1981, 88, 485-494.
- Edelman, Benjamin, Michael Ostrovsky and Michael Schwarz, "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars of Keywords," *American Economic Review*, 2007, 97(1), 242-259.
- Erdil, Aytek and Haluk Ergin, "What's the Matter with Tie-Breaking?: Improving Efficiency in School Choice," *American Economic Review*, 2008, 98(3), 669-689.

Gale, David and Lloyd Shapley "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 1962.

Gul, Faruk and Ennio Stacchetti, "Walrasian Equilibrium with Gross Substitutes," *Journal of Economic Theory*, 1999, 87(1), 95-124.

Gul, Faruk and Ennio Stacchetti, "The English Auction with Differentiated Commodities," *Journal of Economic Theory*, 2001, 92(1), 66-95.

Hatfield, John and Paul Milgrom, "Matching with Contracts," *American Economic Review*, 2005, 95, 913-935.

Immorlica, Nicole and Mohammad Madhian, "Marriage, Honesty and Stability," *SODA* 2005, 53-62.

Kelso, Alexander and Vincent P. Crawford, "Job Matching, Coalition Formation and the Gross Substitutes," *Econometrica*, 1982, 50, 1483-1504.

Klemperer, Paul, "The Product-Mix Auction: A New Auction Design for Differentiated Goods," *Journal of the European Economic Association*, 2010.

Knuth, Donald E. *Marriage Stables*, University of Montreal Press, 1976.

Kojima, Fuhito and Parag Pathak, "Incentives and Stability in Large Two-Sided Matching Markets," *American Economic Review*, 2009, 99(3), 608-627.

Koopmans, Tjalling and Martin Beckmann, "Assignment Problems and the Location of Economic Activities," *Econometrica*, 1957, 25(1), 53-76.

Milgrom, Paul "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 2000, 108, 245-272.

Milgrom, Paul, *Putting Auction Theory to Work*, Cambridge University Press, 2004.

Milgrom, Paul, "Package Auctions and Package Exchanges," *Econometrica*, 2007, 75, 935-966.

Milgrom, Paul, "Assignment Messages and Exchanges," *American Economic Journal: Microeconomics*, 2009, 1, 95-113.

Pathak, Parag and Tayfun Somnez, "Leveling the Playing Field: Sincere and Strategic Players in the Boston Mechanism," *American Economic Review*, 2008, 98(4), 1636-1652.

Roth, Alvin E., "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 1982, 7, 617-628.

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

Roth, Alvin E. (1991), "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*, 81, 415-440.

Roth, Alvin E. (2002), "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics," *Econometrica*, 70, 1341-1378.

Roth, Alvin E. (2007), "Repugnance as a Constraint on Markets," *Journal of Economic Perspectives*, 21(3), 37-58.

Roth, Alvin E. "What have we learned from market design?" *Economic Journal*, March 2008, 118, 285–310.

Roth, Alvin E. and Elliott Peranson (1999) "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 89, 748-780.

Roth, Alvin E., Tayfun Sonmez, and M. Utku Unver, "Kidney Exchange," *Quarterly Journal of Economics*, 119, 2, May, 2004, 457-488.

Roth, Alvin E., Tayfun Sonmez, and M. Utku Unver, "Pairwise Kidney Exchange," *Journal of Economic Theory*, 125, 2, December 2005, 151-188.

Roth, Alvin E. and Marilda Sotomayor, *Two-Sided Matching: A Study in Game-Theoretical Modeling and Analysis*, Cambridge University Press, 1990

Roth, Alvin E. and Xioalin Xing, "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions," *American Economic Review*, 84, September, 1994, 992-1044.

Scarf, Herbert and Lloyd Shapley (1973) "On Cores and Indivisibilities," *Journal of Mathematical Economics*, 1, 23-28.

Shapley, Lloyd and Martin Shubik (1972) "The Assignment Game I: The Core," *International Journal of Game Theory*, 1, 111-130.

Sonmez, Tayfun and Utku Unver (2011), "Matching, Allocation and Exchange of Discrete Resources," in J. Benhabib, A. Bisin and M. Jackson eds. *Handbook of Social Economics*, Elsevier.

Varian, Hal R. (2007), "Position Auctions" *International Journal of Industrial Organization*.

Wilson, Robert (2002), "The Architecture of Power Markets," *Econometrica*, 2002.