DISCUSSION OF “MATCHING MARKETS: THEORY AND PRACTICE”

YEON-KOO CHE

1. Introduction

Market design seeks to offer practical solutions to various resource allocation problems. Still at its relative infancy, the field has already enjoyed impressive successes in applying economics tools and insights to improve the methods for allocating government resources such as radio spectra, for organizing professional labor markets such as those for medical interns and residents, for assigning students to public schools, and for exchanging kidney donors with medical incompatibilities among transplant patients.[1]

An important constraint encountered in many real-world allocation problems is that monetary transfers are limited or unavailable; for instance, public school seats and human kidneys cannot be traded for money.[2] The limitation in the use of monetary transfers means that the classical solution—a competitive market—is excluded from the designer’s tool box. A competitive market, or market-inspired mechanisms such as auctions, or Vickrey-Clarke-Groves (VCG) mechanism, is appealing from both efficiency and fairness standpoints. These mechanisms require agents to outbid other contenders for an object he/she desires; the object will then go to the agent willing to pay the most, typically one who values it most. The resulting allocation is thus efficient. It is also fair in the sense that the objects valued more by the agents are priced higher, requiring higher payment for those wishing to claim them. For instance, the VCG mechanism requires one to pay precisely the social opportunity cost of his/her assignment.

Absence of monetary transfers means that efficiency and fairness must be achieved through other means. The matching theory surveyed by Abdulkadiroglu and Sonmez in this volume

---

Date: April 23, 2011.

Che: Department of Economics, Columbia University, New York City, NY 10027 and YERI, Yonsei University (email: yeonkooche@gmail.com). This paper is prepared as the author’s discussion of the invited paper by Atila Abdulkadiroğlu and Tayfun Sönmez for the 2010 World Congress of Econometric Society. The author is grateful to Fuhito Kojima for very helpful comments and suggestions.

[1] See Roth (2002) for a general survey on market design.

[2] The limitation of monetary transfers arises from moral objection to “commoditizing” objects such as human organs and from fairness considerations (Roth, 2007). Assignment of resources based on prices often favors those with the most wealth rather than those most deserving, and can be regarded as unfair for many goods and services. See Che, Gale and Kim (2011) for an argument along these lines based on utilitarian efficiency.

represents some of the best known ideas in accomplishing these goals. The survey begins with a tribute to two pioneering approaches to allocating resources without monetary transfers—the Gale and Shapley (1962)'s deferred acceptance mechanism and Shapley and Scarf (1974)'s top-trading cycles (suggested originally by David Gale)— and explains how these ideas have been further developed and extended to deal with issues arising in specific problems.

To appreciate how these matching mechanisms seek to achieve desirable outcomes, it is useful to understand what features of these mechanisms replace the role of "prices" in classical markets. Hence, I will briefly interpret the two pioneering mechanisms from this perspective. Doing so will lead to one of my main comments: namely, how one can explicitly introduce features that can harness the beneficial role of prices in matching markets. My second comment suggests expanding the scope of applying matching theory by converting inefficient decentralized procedures into a well-designed centralized mechanism.

2. Matching Mechanisms in Standard Environments

In the standard matching mechanisms, ordinal preference/priority rankings submitted by the participants "price" individual choice, and often this is sufficient to induce desirable outcomes. To illustrate this point, consider how a canonical matching mechanism, Gale and Shapley's deferred acceptance algorithm (henceforth, DA) solves the house allocation problem, that is to say, assign a finite number of agents to a finite number of objects, one for each. Imagine that agents represent students and the objects represent different public schools (each with one seat, for simplicity). The mechanism requires students to rank the schools according to their preferences and the schools to rank the students according to their priorities. (Schools may prioritize students based on their entrance exam scores and/or their residence areas or sibling attendances at the schools.) The algorithm them proceeds in successive rounds: in each round, the students apply to the schools who have not yet rejected them and the schools in turn reject all but at most one applicant. The students apply in the order of their preference rankings and schools decide based on their priority rankings. The algorithm is terminated when no proposal is rejected, at which point the matching is finalized.

Just as with competitive markets, matching can be seen as a result of agents choosing optimally from opportunity sets available to them. For instance, in the student-proposing deferred acceptance algorithm, the match picks the best school for each student among all schools that have not yet rejected her, and it picks, for each school, the best student among all students who eventually apply to it.[3] Unlike competitive markets, the opportunity set facing each agent is not guided by market prices. Rather, the opportunity set is determined

---

[3] Hatfield and Milgrom (2005) make this sense precise.

by the rankings submitted by all the other agents and the schools. Remarkably, these rankings do “aggregate” the other agents’ preferences and schools’ priorities efficiently. For any student, the schools that are available to her are precisely those that have applications from students worse than her (from the schools’ standpoint), and the schools that are unavailable to her are precisely those that have applications from better students than her; and likewise for the schools. Hence, the algorithm ends up pricing individual choices efficiently and produces a matching that is Pareto efficient if we were to interpret priorities as schools’ preferences. The matching is also fair in the sense of eliminating “justified envy”: a student can never lose a school seat to another student with a lower priority at that school (Balinski and Sönmez, 1999; Abdulkadiroğlu and Sönmez, 2003). Even focusing solely on students’ welfare, the outcome is constrained efficient in that the assignment Pareto dominates all other assignments eliminating justified envy. Further, the matching is strategy-proof, making it a dominant strategy for students to rank the schools truthfully.

Next, consider how another canonical matching mechanism, Top Trading Cycles (henceforth, TTC), solves the house allocation problem. Given strict rankings on both sides, the TTC runs recursively in multiple rounds as follows: In any round, each (remaining) agent is endowed with the object that ranks it highest, each agent points to the agent who owns her most preferred object, a cycle of the pointing then emerges, and the agents in a cycle trade away their endowments for their favorite objects. The trading agents and their objects are then removed, and the algorithm iterates to the next round. The algorithm terminates when all such trading occurs, at which point the assignment is complete.

The analogy with markets is even more apparent here. The TTC assignment can be seen as that of a competitive market where the objects within the same trading cycle are priced the same and strictly higher than those objects traded in later cycles.[4] This mechanism is strategy-proof, and the resulting assignment is Pareto efficient. It is also constrained-fair in that no other Pareto efficient and strategy-proof mechanisms eliminate strictly more incidences of justified envy.[5]

3. HARNESSING THE BENEFITS OF MARKETS IN MATCHING MECHANISMS

Ordinal preference rankings are sufficient for implementing desirable matching allocations in many standard settings. This is not the case, however, if the priority/preference rankings are non-strict for some objects or agents. To illustrate, consider again the school choice problem but one in which schools have no priorities, or equivalently, the students enjoy the same priority at all schools. In this case, there is a priori no reason for the agents to be treated

---

[4] The interpretation is made in Shapley and Scarf (1974) and Roth and Postlewaite (1977).

[5] See Abdulkadiroğlu, Che and Tercieux (2011).

differently by the schools. It is then natural and fair, especially given the unavailability of monetary transfers, to employ a fair lottery to break ties whenever many agents compete for a limited seats at a school. A standard practice is therefore to employ either Gale and Shapley's agent-proposing deferred acceptance algorithm or the top trading cycles algorithm, but with the priority ranking of agents at each object generated by a fair lottery. Note that the former mechanism coincides with the popular assignment method known as the random serial dictatorship (RSD) which randomly orders the agents, then lets them choose objects one at a time in order.[6] These mechanisms implement the assignments that are Pareto efficient ex post, but the assignments may not be Pareto efficient ex ante—namely, it may be possible to reallocate the probability shares of the agents to benefit all agents.[7]

To illustrate, suppose there are three schools a, b, c (with one seat) and three students, 1, 2, 3, each evaluating the lotteries with von-Neumann Morgentern (vNM) utilities described as follows:[8]

	v_j^1	v_j^2	v_j^3
$j = a$	4	4	3
$j = b$	1	1	2
$j = c$	0	0	0

Here, the agents have the same ordinal rankings of the objects. Hence, in DA and TTC with with random priorities, since the students rank the schools truthfully, they are assigned the three schools with equal probability, and receive the expected utility of $EU_1^{DA} = EU_2^{DA} = EU_3^{DA} = \frac{5}{3}$.

This assignment is Pareto dominated by the following ex ante assignment: Student 3 is assigned b for sure, and agents 1 and 2 are each assigned between a and c with probability 1/2. They all receive the expected utility of $EU_1^B = EU_2^B = EU_3^B = 2$, more than 5/3. This alternative assignment does better since even though their ordinal preferences are the same, their cardinal preferences—preferences intensities—differ: When b is substituted for a, students 1 and 2 suffer relatively more than agent 3 does (notice an agent's vNM values sum to the same amount). This difference means that the first two students are willing to trade off their probability shares of b in exchange for shares of a, with respect to student 3. The ordinal mechanisms such as DA and TTC simply do not allow these cardinal preferences to be reflected in the assignment.

---

[6] The RSD is used for assigning university housing, public school seats, after school programs, and parking spaces. Note also that the RSD also coincides with the TTC with random endowment (see Abdulkadiroğlu and Sönmez (1998) and Pathak and Sethuraman (2010)).

[7] In fact, they may generate assignments that are ex ante inefficient even in the ordinal sense. See Bogomolnaia and Moulin (2001).

[8] The example here is borrowed from Abdulkadiroğlu, Che and Yasuda (2008).

In order to achieve an ex ante efficient allocation, a mechanism must allow the agents to communicate cardinal preferences and aggregate them efficiently into their opportunity set. This can be done by explicitly introducing markets for lotteries of goods, that is, without monetary transfers. The pseudo-market mechanism proposed by Hylland and Zeckhauser (1979) takes precisely this approach. The mechanism "simulates" the outcome that would emerge were there markets for lotteries: each agent is endowed with the same budget in "fictitious" currency (e.g., 100 tokens), and uses the currency to purchase optimal probability shares of alternative objects, at prices that clear the markets. The resulting allocation determines the lotteries of goods, one for each agent,[9] and the allocation is ex ante Pareto efficient by the first welfare theorem. In the above example, it can be seen that, given the budget of 100 tokens, prices of $(p_a, p_b, p_c) = (200, 100, 0)$ implements the good allocation.

The pseudo market mechanism requires agents to formulate their cardinal preferences. This can be challenging for the agents, and a mistake in formulating preferences may result in misallocation. In practice, cardinal preferences can be expressed and aggregated by adding a simple feature to the ordinal mechanisms. Abdulkadiroğlu, Che and Yasuda (2008) suggest a simple way to modify the DA. In their choice-augmented deferred acceptance (CADA) algorithm, the agents must submit the name of a "target" school, in addition to the rankings of schools. This additional information is then used to break ties (possibly together with a lottery), to favor those who targeted a school. More precisely, once agents' priorities are randomly drawn, each school reshuffles the ordering to elevate the priorities of those who targeted that school over those who have not. Other than the tie-breaking, the CADA coincides with the standard DA. This simple feature improves the capacity of the mechanism to aggregate the agents' cardinal preferences, thus improving the overall ex ante efficiency. In particular, in the above example, the CADA implements the good allocation since student 3 will target school b, whereas the other two will target a.[10] Job market signaling in the market for economics assistant professors is expected to have a similar effect (Coles, Kushnir and Niederle, 2010).

---

[9] Since the mechanism determines the marginal distribution of random assignment, i.e., the lotteries of objects for the agents, one must design the lottery over pure assignments to resolve the randomness in a way that implements the intended allocation. This requirement is ensured by Birkhoff-von Neumann theorem (Birkhoff, 1946; von Neumann, 1953).

[10] All students submit their rankings truthfully, so $a - b - c$ in this order. School a will then rank students 1 and 2 ahead of 3 (but randomly between the first two), and school b will rank student 3 ahead of 1 and 2 (again randomly between these two). In the first round of CADA, all students apply to a which will then choose one of the first two, rejecting the others. They then apply to school b in the second round, and school b will admit student 3. Hence, student 3 is assigned b for sure, and 1 and 2 are assigned between a and c with equal probability, thus implementing the superior assignment.

4. Exploring New Applications

Market design as a general field, and matching theory as a specific tool, has made tangible contributions in solving real-life problems, in areas such as medical matching, school choice and kidney exchanges. The many issues associated with these problems continue to occupy the attention of market design theorists. At the same time, it seems worth exploring new applications, for many ideas developed from these experiences can also be fruitfully applied beyond the set of problems studied so far.

Important cases in point are college and graduate program admissions. In the US and many other countries, admissions into colleges and graduate schools are conducted through decentralized procedures. The procedures differ across countries and specific programs, but, roughly speaking, they involve three steps: Students first apply to (often multiple) schools they wish to enter. Programs then evaluate the applications and select from among the applicants, offering admissions to those selected. Finally, those who have received admissions decide which offers to accept.

These procedures are rife with problems that are all too familiar, and more importantly can be addressed with the tools already available, to matching theory. When students apply to programs, they do not know which programs will admit them, and likewise programs do not know which students will accept their admissions offers. Both sides make mistakes. A student may aim too high and miss out on schools she should apply to or may, for fear of rejection, fail to apply to a school that would be happy to admit her. Likewise, a program may fail to admit a student it should, either because it is too optimistic about its offers (made to better students) being accepted or too pessimistic about the student not accepting its offer. Either way, there are many ways for justified envy to arise.

Justified envy makes the assignment unfair for both students and programs, but it also entails other social costs. For instance, many college applicants in Korea and Japan who feel under-placed wait for a year to re-apply (often without foregoing their seats in the original programs that have admitted them, in which case the programs' resources are wasted).[11] Many of these students essentially "waste" a year to get another crack at the process. Of course, the mistakes also entail other forms of inefficiencies: some students may be unassigned, some schools may have unfilled seats, and these schools would like to fill those seats with the unassigned students and the students would be happy to enroll in these schools.

---

[11] In Korea, the number of re-applicants—the applicants who had spent at least one year after graduating from a high school—participating in the 2010 college admissions process reached about 170,000 (21%) out of about 710,000 applicants in total (see http://www.asiae.co.kr/news/view.htm?idxno=2011030711285711758). In Japan, the number of re-applicants was about 110,000 out of 1.18 million participants in 2010. I thank Fuhito Kojima for pointing out similar practices and the figures on Japan.

Decentralized procedures cope with these problems in several ways. A common method is for programs to over-admit students often at several multiples of their capacities. This reduces the risk of students being unfairly under-placed, but creates the risk of programs exceeding or falling short of their capacity targets. Another method is to improve the "yield" of programs' admissions decisions by revealing students' idiosyncratic preferences prior to general admissions. For instance, early admissions programs in the US and Korea allow students to apply to a limited number of schools (usually one) before the general admissions.[12] Even in general admissions, colleges in Korea are partitioned into several groups and a student can apply to only one school in each group; similarly applicants in UK must choose between Oxford and Cambridge. Forcing students to reveal their preferences through these methods improves the yield of schools' admissions, but limits students' choices. Yet another method is to admit students sequentially in multiple rounds, wait-listing some students early on and upgrading their status to full admissions later if there are rejections from earlier admittees. These coping methods may alleviate, but do not eliminate, the problems mentioned above and may even create new problems. Plainly, they do not achieve what one can achieve from a well-designed centralized procedure. We already know that Gale and Shapley's deferred acceptance (DA) algorithm completely eliminates justified envy and attains Pareto efficiency on both sides, and constrained efficiency on the proposing side subject to no justified envy.

This raises a question: Why are college (or graduate program) admissions not centralized?[13] One reason could be that (at least some) programs do not wish to participate in even a very well-designed centralized procedure such as DA. Even though such a procedure may generate overall benefits in terms of efficiency and fairness, these benefits may not be shared uniformly by all participants. To illustrate, suppose that students have the same ordinal preferences of schools. While extreme, this assumption is not unrealistic since in many countries schools are clearly ranked in respects that students care about. Suppose next that the schools have the same ordinal preferences of the students. Again the assumption is special but not completely unrealistic in many countries where programs admits students based on their scores from national standardized tests. In such a case, the matching produced by both DA and TTC is assortative; i.e., best students match with best program, second-best matches with second-best, and so on. A shift to such a mechanism will clearly benefit students and programs at the top tier but could very well hurt those at a lower tier since these

---

[12] Avery and Levin (2010) focus on a related but slightly different role for early admissions; they suggest that colleges wish to select those students who have specific enthusiasm for them, all else equal, and that early admissions allow students to signal their enthusiasm for colleges.

[13] In fact, several countries, including China, employ centralized college admissions procedures

latter participants have some chance of attracting better partners from the current decentralized procedure. For this reason, despite the clear overall benefits, many participants may not have incentives to participate in the centralized mechanisms such as DA and TTC.

This argument suggests that a new perspective for matching market design is necessary when one considers introducing a centralized procedure into decentralized markets. The standard qualities one strives for in centralized matching—efficiency and fairness—may conflict with the incentives for participation, the very condition required to build the consensus for centralization. In order to implement a successful transition into a centralized mechanism, one may have to compromise on these values. One sensible way to make such a compromise is to consider a hybrid procedure which gives programs an option to retain a decentralized procedure.[14] Specifically, one could redesign college admissions so that programs admit as many students as they wish through early admissions (operated according to the current decentralized procedure), but they fill the remaining seats through a centralized mechanism such as DA or TTC (or even CADA). Such a hybrid procedure makes centralized admissions an option, as opposed to a mandatory requirement, for the programs. Under such a regime, programs can experiment with centralized admissions on a small scale, and as they gain favorable experience, they may gradually increase their use of the centralized procedure. For this reason, such a hybrid procedure is unlikely to face resistance from the programs. Of course, one should clearly understand the welfare and distributional implications of such a mechanism before it is recommended. And the recommended procedure must be further vetted by extensive research in theory and lab and field experiment before being implemented into practice. Nevertheless, I believe that an innovation such as this deserves economists' attention and creativity, for its payoff seems substantial.

References

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 1998. "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems." *Econometrica*, 66: 689–701.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. "School Choice: A Mechanism Design Approach." *American Economic Review*, 93: 729–747.

---

[14] Centralized procedures often do not allow for an option to match outside the procedures. The National Resident Matching Program (NRMP) used to allow for residency programs to hire part of their residents from the NRMP centralized match while others in a decentralized manner, but recently made a change so that programs are required to hire everyone from the centralized match (see http://www.nrmp.org/). Such a prohibition against outside matching seems sensible in circumstances where the consensus toward centralization already exists among relevant parties. The argument made here applies to settings in which no such consensus is available so that an extra incentive (in the form of the decentralization option) need to be provided for transiting an centralized procedure.

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Olivier Tercieux. 2011. "The Role of Priorities in Assigning Indivisible Objects: A Characterization of Top Trading Cycles." Unpublished mimeo.

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2008. "Expanding 'Choice' in School Choice." mimeo.

Avery, Christopher, and Jonathan Levin. 2010. "Early Admissions at Selective Colleges." *American Economic Review*, 100: 2125–2156.

Balinski, Michel, and Tayfun Sönmez. 1999. "A tale of two mechanisms: student placement." *Journal of Economic Theory*, 84: 73–94.

Birkhoff, Garrett. 1946. "Three Observations on Linear Algebra." *Revi. Univ. Nac. Tucuman, ser A*, 5: 147–151.

Bogomolnaia, Anna, and Herve Moulin. 2001. "A New Solution to the Random Assignment Problem." *Journal of Economic Theory*, 100: 295–328.

Che, Yeon-Koo, Ian Gale, and Jinwoo Kim. 2011. "Market versus Non-Market Assignment." Unpublished mimeo.

Coles, Peter, Alexey Kushnir, and Muriel Niederle. 2010. "Preference Signaling in Matching Markets." NBER Working Paper 16185.

Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 69: 9–15.

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

Hylland, Aanund, and Richard Zeckhauser. 1979. "The Efficient Allocation of Individuals to Positions." *Journal of Political Economy*, 87: 293–314.

Pathak, Parag A., and Jay Sethuraman. 2010. "Lotteries in Student Assignment: An Equivalence Result." *Theoretical Economics.*

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: 37–58.

Roth, Alvin E., and Andrew Postlewaite. 1977. "Weak Versus Strong Domination in a Market with Indivisible Goods." *Journal of Mathematical Economics*, 4: 131–137.

Shapley, Lloyd, and Herbert Scarf. 1974. "On Cores and Indivisibility." *Journal of Mathematical Economics*, 1: 22–37.

von Neumann, John. 1953. "A certain zero-sum two-person game equivalent to the optimal assignment problem." In *Contributions to the theory of games, Vol. 2.* , ed. H. W. Kuhn and A. W. Tucker. Princeton, New Jersey:Princeton University Press.