RECENT DEVELOPMENTS IN MATCHING THEORY AND ITS PRACTICAL APPLICATIONS

FUHITO KOJIMA

Abstract. In recent years, many developments have been made in matching theory and its applications to market design. This paper surveys them and suggests possible research directions. The main focus is on the advances in matching theory that try to solve market design problems in practical markets where the classic theory is inapplicable. Specifically, we discuss the recent theory of matching in large markets and "approximate market design," and a new theory of "matching with constraints," as well as their applications.
JEL Classification Numbers: C78, D47, D71.
Keywords: matching, stability, approximation, large markets, matching with constraints, market design

1. Introduction

Matching theory has made a considerable progress ever since the seminal theoretical contribution by Gale and Shapley (1962) and its economic application by Roth (1984). The theory has been used to guide designs of medical match (Roth and Peranson, 1999) and other entry-level labor markets (Roth, 2002), school choice (Abdulkadiroğlu and Sönmez, 2003), course allocation in education (Sönmez and Ünver, 2010; Budish and Cantillon, 2012), and organ donation (Roth, Sönmez, and Ünver, 2004, 2005, 2007), just to name a few examples.

Theory has been put in practical use with much success, but the interaction between theory and practice is not unidirectional. On the contrary, as more practical applications have been made, the accompanying experiences have also uncovered limitations of existing theory. This has opened up opportunities for researchers to not only deepen the analysis of

---

Date: June 30, 2015.
Department of Economics, Stanford University, Stanford, CA 94305, fkojima@stanford.edu. This paper was prepared for an invited talk at the 2015 World Congress of Econometric Society in Montréal. I benefited from discussions with Yeon-Koo Che and Parag Pathak, as well as comments from Indira Puri and seminar participants at Santa Clara, Stanford, and Waseda. Eric Fanqi Shi provided excellent research assistance. I gratefully acknowledge financial support from the Sloan Foundation, as well as financial support from the National Research Foundation through its Global Research Network Grant (NRF- 2013S1A2A2035408).

existing models but also generalize the standard model or introduce new models, thereby enabling us to address challenges that arise in applications.

Much of recent developments in matching theory has been guided by our desire to solve practical problems with features that are absent from the standard model, and yet important in practice. As it turns out, some of these features are so concrete that predictions from theories that incorporate them make sharp predictions, but at the same time they are general enough to enrich our understanding of the markets beyond one specific example.

This paper describes recent developments in matching theory and its applications. The survey by Abdulkadiroğlu and Sönmez (2013) based on their lectures at the last World Congress offers an excellent introduction to the basic models as well as applications.[1] Thus, in this piece, I will focus on several specific topics that represent the kind of recent interaction between theory and practice I alluded to above. That is, my focus will be on how the features of some real markets deviate from the classic models, and how the theory has been modified or expanded to tackle these problems.

In this paper, I will focus on two topics that share the above general theme. The first topic is the analysis of an approach for design that could be called "approximate market design." Classic theory has found that it is often impossible to satisfy some desirable properties, no matter how hard one tries to design a matching mechanism. For example, there exists no two-sided matching mechanism that is strategy-proof and produces a stable matching for every given input. Another example is a result that if a labor matching market includes a married couple who desire a pair of positions for them, then even the existence of a stable matching is not guaranteed. There are many impossibility results like these in the matching literature. Still, for practical design, it is not clear if these negative results pose a big concern. We explore the idea of relaxing the desirable properties by requiring them only approximately: In matching with couples, for instance, we will only require the existence of a stable matching "with high probability." In order to formalize such approximate statements, we typically study the asymptotic behavior of the economy as it becomes large. This modeling approach appears to be especially fit for representative applications such as labor markets for medical residents and school choice in large school districts.

The other topic covered in this paper is the burgeoning literature of "matching with constraints." In the standard theory of matching, the only constraints for an allocation are

[1]Other surveys include Roth (2007, 2008) and Sönmez and Ünver (2009). Roth and Sotomayor (1990) provide a comprehensive account of the early literature.

those given by each agent's ability to consume at most a fixed number of (in many cases, one) good and supply constraints such as limited positions in a firm or limited space in a school. However, many matching markets are subject to other kinds of constraints.[2] For example, medical markets are often constrained by regulations that limit the number of doctors practicing in certain medical specialties. Regulations on geographical distributions of doctors, teachers, and other kinds of workers are widespread. School districts may require each school in the district to maintain a certain balance of student body in terms of socioeconomic status or academic skill. The classic theory is not applicable in such situations, but recent advances in matching with constraints have generalized the theory to analyze these more complicated situations.

These specific topics are of interest by themselves, but I believe that they also enable us to learn a more general lesson. On the one hand, many models in these lines of research are highly tailored to very specific situations so that they can be directly applied to them: This paper offers a lot of discussions of American, British, and Japanese medical residency markets, school choice programs in New York City and Boston, and U.S. clinical psychologist match, for example. On the other hand, many of the models are also general enough to guide the direction of analysis and design beyond any one application. I believe that a good economic theory, or at least good economic theory of matching, is both detailed enough to make a sharp prediction or recommendation in specific applications *and* general and tractable enough to provide general insights and guide our thinking when we face new problems. Thankfully, I think that the effort of the research community has been successful in achieving both of these goals.

The rest of this paper goes as follows. Section 2 presents the basic matching model and classical results. Section 3 describes the research on large matching market models and approximate market design. Section 4 discusses research in matching with constraints. Section 5 concludes.

2. Basic Model

This section introduces a simple two-sided matching model. I describe the model in terms of matching between doctors and hospitals because the theory has been extensively applied to that type of markets. Needless to say, one can use the same model to analyze

[2]Of course, constraints are important in other contexts as well. In the upcoming auction for reallocating spectrum rights in the U.S., there are numerous interference constraints that need to be handled (Milgrom and Segal, 2014). See also Milgrom (2009) who proposes an auction based on the "assignment messages" that can express various preferences and constraints.

many other situations, such as worker-firm matching or heterosexual marriage or school choice (i.e., matching between students and schools) among other things.

Let there be a finite set of doctors and a finite set of hospitals. Let me denote a generic doctor by i, j, and so forth, and a generic hospital by A, B, and so forth. Each doctor i has a strict preference relation $\succ_i$ over the set of hospitals and an exogenously given outside option (denoted by $\emptyset$) which represents being unmatched. For any doctor i and any pair of hospitals (or the outside option) A, B, we write $A \succeq_i B$ if and only if $A \succ_i B$ or $A = B$. We say that hospital A is **acceptable** to doctor i if $A \succ_i \emptyset$.

Notations and terminologies introduced for doctors in the last paragraph are also used for hospitals in analogous manners. We assume each doctor can work for at most one hospital, while allowing for a hospital to hire multiple doctors (a model with this property is called a many-to-one matching model). Thus, naturally we assume that each hospital has a preference relation over the set of subsets of doctors. Although in principle the preferences can be quite arbitrary, throughout this paper we assume that hospital preferences satisfy a condition called **responsiveness**: This condition requires that the hospital has a capacity (i.e., the number of positions) and a linear order over the doctors plus the outside option, and its optimal choice from any set of applicants to it is its most preferred acceptable doctors (with respect to the aforementioned linear order) up to its capacity.

Let us introduce further notation for preferences. Thanks to the responsiveness assumption, it turns out that only the ranking over acceptable partners matters for our analysis. Given this observation, we denote preferences by writing an ordered list of only acceptable partners. For example, the notation $\succ_i$: A, B represents a preference relation of doctor i such that she likes hospital A best, hospital B second, and the outside option third, and finds all other hospitals unacceptable.

A **matching** μ is a mapping that specifies which doctor is assigned to which hospital (or the outside option). A matching μ is **individually rational** if (i) no agent is matched with a partner who is unacceptable to her, (ii) no doctor is matched with more than one hospital, and (iii) no hospital is matched with more doctors than its capacity. Given a matching μ, a **blocking pair** to it is a doctor-hospital pair who prefer to be matched with each other (while possibly rejecting some or all of their partners at μ) rather than being matched according to μ. A matching is **stable** if it is individually rational and there is no blocking pair to it.

A **mechanism** is a function from the set of preference profiles to the set of matchings. A mechanism is said to be **strategy-proof** if reporting the true preferences is a weakly dominant strategy for each agent. Strategy-proofness is regarded as an important property for a mechanism's success. However, it is well-known that strategy-proofness and stability are incompatible to each other. More specifically, there is no strategy-proof mechanism that produces a stable matching for all possible preference profiles (Roth, 1982). Given this impossibility result, it is customary to consider incentive compatibility for doctors only. A mechanism is said to be **strategy-proof for doctors** if truthful preference reporting is a weakly dominant strategy for each doctor.

In this problem, the following **(doctor-proposing) deferred acceptance algorithm** due to Gale and Shapley (1962) plays a crucial role:

- Step 1: Each doctor applies to her first choice hospital. Each hospital rejects its least-preferred doctors in excess of its capacity and all unacceptable doctors among those who applied to it, keeping the remaining doctors tentatively.

In general, for any $t = 1, 2, \dots$

- Step t: Each doctor who is not tentatively kept by a hospital applies to her next highest choice (a doctor does not apply to any hospital if she has been rejected by all hospitals acceptable to her). Each hospital considers these doctors *and* doctors who are tentatively kept from the previous step together, and rejects its least-preferred doctors in excess of its capacity and all unacceptable doctors, keeping the remaining doctors tentatively.

This algorithm terminates at the first step in which no rejection occurs (clearly, it terminates in a finite number of steps). Gale and Shapley (1962) show that the resulting matching is stable. This algorithm plays an important role throughout, and I discuss it many times, so I refer to it as DA for brevity (the *hospital-proposing* version of the deferred acceptance algorithm can be defined in an analogous manner, but in this paper I use DA to refer to only the doctor-proposing version of the algorithm).

Even though there exists no strategy-proof mechanism that produces a stable matching for all possible inputs, DA is strategy-proof for doctors (Dubins and Freedman, 1981; Roth, 1982).

3. Large Markets and Approximate Market Design

3.1. **Manipulating stable mechanisms.**

As already mentioned, the stability requirement for a matching mechanism necessarily leads to manipulation possibilities, because there is no stable mechanism that is strategy-proof for both sides of the market. Below is a (perhaps unrealistically) simple example illustrating this point for DA.

Example 1. Consider the following market with two hospitals A, B, and two doctors i, j. Suppose that each hospital has only one position, and the hospital and doctor preferences are

$$\succ_A : i, j, \qquad \succ_i : B, A,$$
$$\succ_B : j, i, \qquad \succ_j : A, B.$$

With this input, at the first step of DA, i and j apply to B and A, respectively, and both of them are accepted. Hence, DA terminates at the first step, and the resulting matching is

$$\mu = \begin{pmatrix} A & B \\ j & i \end{pmatrix},$$

where this matrix notation means A is matched to j while B is matched to i.[3]

Now, suppose that hospital A misreports its preferences by declaring that only i is acceptable to it (hence j is unacceptable to it). Then, at the first step of DA, j's application to A is rejected, while i's application to B is tentatively accepted. The algorithm proceeds to step 2, j applies to B, and B tentatively accepts j while rejecting i. Then, at step 3, i applies to A, and this application is accepted. DA terminates at this step, and the resulting matching is:

$$\mu' = \begin{pmatrix} A & B \\ i & j \end{pmatrix}.$$

Comparing matchings μ and μ', one can verify that hospital A is made better off by misreporting its preferences, implying that DA is not strategy-proof. □

As mentioned earlier, the manipulability of DA in the above example can be extended to the general impossibility of strategy-proof and stable mechanisms. This finding has a profound implication for matching theory. It provides a sense in which a desirable market design is impossible for the two-sided matching problem if its participants behave strategically. However, whether this impossibility result has a "bite" in practice or not is far from obvious. This is because the result is based on a highly stylized example under complete information, and it is silent about whether DA or other stable mechanisms are easily manipulable in reality. If anything, there are many real markets that employ stable matching mechanisms despite its incentive incompatibility, and evidence suggests that stable mechanisms work without much problem while unstable ones often fail (Roth, 1991, 2002). A question for economists is to understand why stable mechanisms appear to be working despite the failure of strategy-proofness.

[3] I use similar notation throughout this paper.

To tackle this problem, Roth and Peranson (1999) conduct a series of simulations on data from the NRMP and other medical matching markets as well as randomly generated data. In their analysis of NRMP data, they find that very few agents had any profitable misreporting of preferences. Among the key features of the NRMP data is that there are many participants. For the randomly generated data, Roth and Peranson (1999) find that the proportion of agents who can profitably misreport preferences is smaller in large markets than in small markets, as shown in Figure 1. These findings lead them to

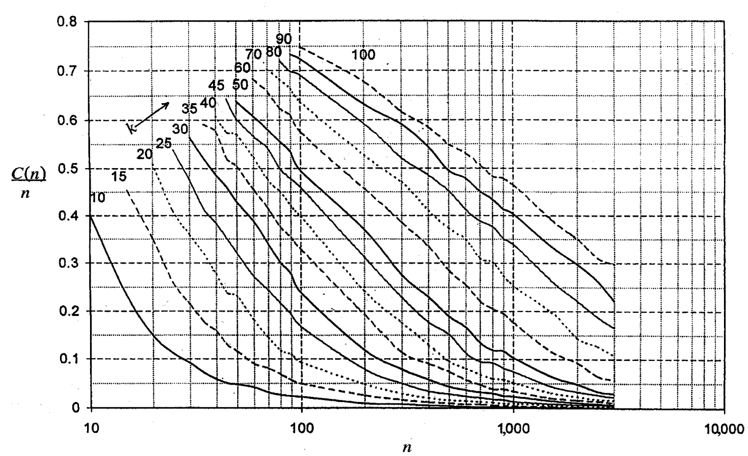

FIGURE 1. The vertical axis measures the proportion of hospitals that have a profitable unilateral manipulation in DA (denoted $C(n)/n$). There are n doctors and n hospitals, each doctor finds a constant number k of hospitals acceptable, each firm ranks all doctors who apply to it, and preferences are uniformly i.i.d. Reproduced from Roth and Peranson (1999).

conjecture that the proportion of participants who can profitably misreport preferences decreases and tends to zero as the size of the market grows.[4]

Theories have been developed to account for this regularity found in numerical analysis, and formalized a sense in which the above conjecture holds. We follow Immorlica and Mahdian (2005) and Kojima and Pathak (2009) to illustrate the main forces behind this phenomenon. Consider a growing sequence of markets indexed by the number of hospitals, which can be interpreted as representing the market size (the number of doctors are allowed –although not required– to grow at the same rate as the hospitals). Each agent's preference is his or her private information. Doctors and hospitals simultaneously report

[4]Similar empirical observations have also been made for other markets such as APPIC, which is an American market for clinical psychologists (Kojima, Pathak, and Roth, 2013), and Boston Public School's student assignment system (Pathak and Sönmez, 2008).

their preferences given their private information. Under certain regularity conditions (which we discuss in Section 3.1.1), Kojima and Pathak (2009) show approximate incentive compatibility of DA as formalized in the following statement.

Theorem 1. *Consider a sequence of markets. For any $\varepsilon > 0$, there exists a positive integer n such that truthtelling is an ε-Bayes Nash equilibrium in the game of preference reporting under DA for each market with at least n hospitals.*[5]

Instead of presenting the formal proof, let us describe an intuition here. To do so, recall first that DA is strategy-proof for doctors, so it suffices to show approximate incentive compatibility for hospitals. More specifically, we shall establish that, for each hospital, truthtelling of its preferences is an approximate best response if everyone else is truthful. For this purpose, we first note that the intuitive reason for a successful misreporting by hospitals is by way of a "rejection chain," which is a chain reaction of applications and rejections that occur during the execution of the algorithm. More specifically, a rejection chain is a sequence of applications and rejections initiated by a strategic rejection which causes the rejected doctor to apply to her next choice hospital, which may displace another doctor, who applies to her next choice hospital, and so forth. Some of the rejected doctors along the rejection chain may apply to the misreporting hospital, and this hospital may be made better off if these new applicants are desirable.[6] In the context of the previous example, hospital A initiates a rejection chain by rejecting j; the rejected doctor j applies to her next choice hospital B, causing it to reject doctor i; and this in turn causes i to apply to the original manipulating hospital A, making it better off.

The above observation helps obtain the intuition for the approximate incentive compatibility result in Theorem 1. In a large market (in the sense of having a large number of participants), there are many hospitals with at least one vacant position at the DA outcome under truthtelling (with high probability).[7] This implies that the doctors who are strategically rejected by a misreporting hospital or the doctors who are rejected by them –directly or indirectly– later in the rejection chain are likely to apply to one of the hospitals with those vacant positions. Once a doctor applies to a vacant position, she is accepted and the rejection chain terminates because no new doctor is rejected. When

[5]We say that a strategy profile is an ε-Bayes Nash equilibrium if no agent gains more than ε in expectation by any unilateral deviation from that strategy profile.

[6]This intuition turns out to be precise. More specifically, Kojima and Pathak (2009) show that if there exists profitable misreporting for a hospital, then that hospital receives an application by some doctor at some step of the rejection chain.

[7]A number of regularity conditions play important roles to guarantee that there are many vacant positions. This point is discussed in some detail in Section 3.1.1.

this happens, the misreporting hospital loses the original doctor because of strategically rejecting her, but it does not get any new doctor. Therefore, in a large market the misreporting hospital is unlikely to be made better off at all, which implies approximate incentive compatibility as stated in the theorem.

Remark 1. Theorem 1 has an implication for a structural property of the set of stable matchings. It is well-known that if there are two stable matchings such that a hospital is matched to different sets of doctors between them, then there is a profitable preference misreporting for that hospital under DA (the converse holds for one-to-one matching, although not in many-to-one matching). Hence Theorem 1 implies a "core convergence" property in the sense that if the market is large, then for most agents, their stable matching partner(s) is unique.

3.1.1. *How robust is approximate incentive compatibility?* While it may be tempting to summarize the aforementioned findings as that "incentive problems go away for DA in large markets," such a general statement can be misleading. In particular, the result of Kojima and Pathak (2009), as well as earlier analyses by Roth and Peranson (1999) and Immorlica and Mahdian (2005), relies on a number of regularity conditions.

Among others things, let me discuss the so-called "limited acceptability" assumption.[8] This assumption requires that each doctor finds only a small fraction of hospitals acceptable among many hospitals in the market. In its simplest and perhaps most restrictive form, this notion is formalized as the assumption that there is a fixed number k such that all doctors find only k hospitals acceptable even as the number of the hospitals, n, grows without a bound.

Limited acceptability is widely assumed in the literature despite its restrictiveness, so it seems to warrant discussion here. Roth and Peranson (1999) motivate this assumption by an empirical observation. Specifically, they note that the numbers of employers listed as acceptable in most applicants' submitted preferences are very small in practice. For example, most applicants in NRMP apply to 15 or fewer hospital residency programs out of about four thousand programs (Roth and Peranson, 1999); in Japanese Residency Matching Program, the average number of programs listed by applicants is between 3 and 4 out of more than 1000 residency programs (Kamada and Kojima, 2015); in NYC high-school match, more than 70 percent of the students list 11 or fewer school programs out of about 500 (Kojima and Pathak, 2009); in APPIC (an American matching market

[8]To my knowledge, the term "limited acceptability" was coined by Lee (2012).

for clinical psychologists), both the median and mean number of programs listed in the applicants is between 7 and 8 out of just over 1000 programs.

Another motivation for this assumption is that under *un*limited acceptability (i.e., the assumption that every doctor finds every hospital acceptable and vice versa), predictions from large market models do not match stylized facts. Roth and Peranson (1999) conduct simulations in markets in which the numbers of doctors and hospitals are equal to each other, and they find that the average proportion of hospitals that have a profitable misreporting opportunity *increases* under unlimited acceptability, as shown in Figure 2. In

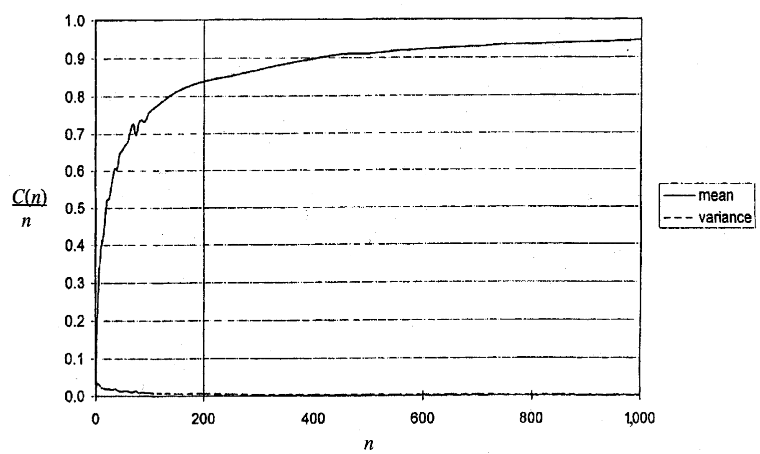

FIGURE 2. The vertical axis measures the proportion of hospitals that have a profitable unilateral manipulation in DA (denoted $C(n)/n$). There are n doctors and n hospitals, each doctor finds each hospital acceptable and vice versa, and preferences are i.i.d. uniformly distributed. Reproduced from Roth and Peranson (1999).

fact, Knuth, Motwani, and Pittel (1990) have shown earlier that the expected proportion of such hospitals converges to 100 percent as the market size goes to infinity.[9] These results cast some doubts on the external validity of the above approximate incentive compatibility results: Does it hinge on very restrictive assumptions or is it a robust feature of most markets in practice?

[9]Kojima and Pathak (2009) note that the extreme form of limited acceptability such as constant k is not needed for their conclusion and show sufficient conditions for the growth rate of k as a function of n, but the growth rate that is allowed in their analysis is as slow as $o(\ln(n))$.

To my knowledge, Lee (2012) is the first to make a significant theoretical advance to shed light on this issue. He sets up a model of one-to-one matching[10] where the numbers of doctors and hospitals are equal to each other, and each doctor finds each hospital acceptable and vice versa– an environment where each hospital has a profitable misreporting under DA with high probability. Agents are ex ante identical and draw cardinal utility functions from a certain distribution. Lee (2012) shows that, for any stable matching mechanism, the expected cardinal utility gain from optimal preference misreporting becomes vanishingly small as the market becomes large. Note that this result does not contradict the earlier manipulability results under unlimited acceptability by Knuth, Motwani, and Pittel (1990) and Roth and Peranson (1999). This is because what these earlier papers show is that there is a large probability that a misreporting results in some strictly positive utility gain, whereas what Lee (2012) shows is that for most instances of profitable manipulation, the cardinal utility gain from a manipulation is very small. His result shows that, even with unlimited acceptability, there is a sense in which stable mechanisms are hard to exploit by preference misreporting.

Another major advance has been made by a recent contribution by Ashlagi, Kanoria, and Leshno (2014). They consider one-to-one matching markets with unlimited acceptability, just as earlier studies such as Knuth, Motwani, and Pittel (1990), Roth and Peranson (1999), and Lee (2012). In contrast to these contributions, however, they assume that the market is unbalanced in the sense that the numbers of doctors and hospitals are different from each other.[11] Surprisingly, they find that unbalanced markets behaves very differently from balanced ones. They establish that in an unbalanced market, the expected number of hospitals who can profitably misreport becomes small as the market becomes large, just as in the markets with limited acceptability studied by Kojima and Pathak (2009). In other words, the manipulability problem under unlimited acceptability may be a knife edge phenomenon for balanced markets, and in a "typical" market (even with a small imbalance between the demand and supply of doctors), stable mechanisms are robust to strategic manipulations.

3.2. **Complementarity and existence.** Another problem that a labor matching organizer often faces is caused by preferences that exhibit complementarity. A canonical example is preference complementarity of married couples who need two jobs, one for each

[10]Appendix E of Lee (2012) extends his results to the many-to-one matching model, but the analysis for this case is significantly more complicated, and it does not fully establish approximate incentive compatibility for the many-to-one case.

[11]Some other studies such as Immorlica and Mahdian (2005) and Kojima and Pathak (2009) allow for balanced as well as unbalanced markets, and their results hold in both cases.

member. For most couples, many pairs of positions are complements because they want to have two jobs close to each other.

Match organizers such as NRMP and APPIC try to accommodate couples by letting them submit preferences over pairs of hospital positions and treating them differently from singles in the algorithm.[12] Couples pose a more fundamental problem, however, because their presence in the market can lead to non-existence of stable matchings. This, of course, implies that it is impossible to find a mechanism that produces a stable matching for all possible preferences. To see these facts clearly, consider the following simple example.

Example 2. There are two hospitals A and B with one position each, one single doctor s and one couple (m, w) composed of the man m and the woman w. Their preferences are

$$\succ_A : m, s, \qquad \succ_s : A, B,$$
$$\succ_B : s, w, \qquad \succ_{(m,w)} : (A, B),$$

where the notation for the couple (m, w) means that the only acceptable outcome for this couple is for m to be matched to A and w to be matched to B, and everything else is less preferred to the outside option (the outcome in which both members of the couple are unmatched). To show that there exists no stable matching in this market, consider the following (exhaustive) cases.

(1) Suppose m and w are matched to A and B, respectively. Then s is unmatched because the seats of both hospitals are filled by the couple. Given this, s and B block this matching because s prefers B to her current outcome (the outside option) and B prefers s to its current match, w.

(2) Suppose m and w are both unmatched.
   (a) If s is matched to A, then the couple (m, w) and hospitals A, B block this matching since the couple prefers (A, B) to the current outcome (the outside option), A prefers m to its current match s, and B prefers w to its current outcome (the outside option).
   (b) Suppose that s is matched to B or s is unmatched. Then s and A block this matching because A is the first choice for s while A prefers s to its current outcome (the outside option).

[12] Other match organizers such as Scottish Foundation Allocation Scheme (the medical matching market in Scotland) do not ask couples to submit preferences over pairs of positions, but have the algorithm follow an exogenously given rule to guarantee that the two positions for each couple are not too far away from each other.

(3) Every other matching is individually irrational for the couple (m, w), so it is unstable.

Cases (1)–(3) are exhaustive, so this establishes that there exists no stable matching in this market.

This conclusion immediately implies that there exists no algorithm which is guaranteed to produce a stable matching in the presence of couples. To get intuition, however, consider an algorithm that tries to modify the doctor-proposing DA to accommodate couples. More specifically, imagine that the algorithm lets doctors apply as in DA, but lets each couple apply to pairs of positions, and if one member of the couple is rejected, then it lets its remaining member withdraw her application to her tentative partner, and lets other applicants withdraw their current applications to apply to the newly open position, and so on (these features are important for addressing preference complementarity). If one applies such a modified algorithm to the present example, then it would proceed as follows: In the first step, s applies to A while m and w apply to A and B, respectively. Given these applications, A rejects s in favor of m. Then, in the second step, s applies to his second choice, B. Now, given the new application, B rejects w and keeps s. Because w has been rejected, the couple withdraws the application of m to A. This withdrawal creates a vacancy at A, so now s withdraws the application to B and applies to A. Finally, m and w apply to A and B, respectively, which results in the same situation as in the very first step of the algorithm. Thus, this heuristic algorithm follows a cycle of applications and rejections and fails to terminate, never producing a stable matching in this market. □

Despite this theoretical possibility, some matching clearinghouses in practice regularly enjoy high rates of participation and produce matchings that are honored by participants. In fact, these clearinghouses have almost always found stable matchings. For instance, Roth and Peranson (1999) run several variants of DA on submitted preferences from 1993, 1994 and 1995 in NRMP and find that these algorithms produce a stable matching for each of these years. Kojima, Pathak, and Roth (2013) similarly study data from APPIC during 1999-2007 and find a stable matching for each of these years.

These observations beg the question: why do these matching clearinghouses produce stable matchings even though the standard theory suggests there may exist no stable matchings when couples participate? Kojima, Pathak, and Roth (2013) suggest that the market size may be the answer. More specifically, under certain regularity conditions similar to those in Kojima and Pathak (2009), the probability of the existence of a stable matching is high in large markets:

Theorem 2. *Consider a sequence of markets. For any $\varepsilon > 0$, there exists a positive integer n such that the probability that a stable matching exists is at least $1 - \varepsilon$ for each market with at least n hospitals.*

This result can be shown using a constructive algorithm, which shows not only the existence of a stable matching, but also how to find it. The algorithm is a close variant of the one used in NRMP and works as follows. First, we run DA for a market composed of hospitals and single doctors only, while excluding couples. Then we add couples to the market one by one, allowing them to apply to pairs of positions from their top choices as described in Example 2. The additional applications by a couple may displace some doctors, creating rejection chains just as those created by a strategic rejection analyzed in Section 3.1. If a rejection chain reaches a hospital which a member of some couple is tentatively matched to, then he or she may be dislocated and cause withdrawal of an application by the remaining member of the couple, which may lead to a failure of an algorithm, as we have observed in Example 2. In a large market with many hospitals, however, it is with high probability that the rejection chains created by the couples are terminated as a doctor applies to a hospital with a vacant position, just like the rejection chain created by strategic manipulation in the last section. Therefore, in large markets with a relatively small number of couples, the probability of the existence of a stable matching is high.

3.2.1. *How robust is approximate existence?* As in the case of the approximate incentive compatibility, it is important to note that the proof of asymptotic existence by Kojima, Pathak, and Roth (2013) depends on a number of assumptions such as limited acceptability. One of the most important– and perhaps the most restrictive– assumptions is that the number of couples is very small compared to the market size. More specifically, the authors assume that the number of couples grows at a rate that is strictly smaller than the square root of the number of the hospitals (whereas the number of single doctors can grow at the same rate as the number of hospitals), and this assumption plays a crucial role in their proof. Roughly speaking, this assumption requires that the proportion of couples becomes vanishingly small, and does so fast enough. This is unfortunate because there is no obvious reason that a large market size makes most applicants singles. Kojima, Pathak, and Roth (2013) complement their theoretical results using data from APPIC and simulations on randomly generated data, verifying that in these markets, the probability of finding a stable matching is indeed close to one. Still, the reason for such high probability of existence remained somewhat unresolved.

Ashlagi, Braverman, and Hassidim (2014) address this issue. They demonstrate that if the number of couples increases strictly more slowly than the market size–not the square root of the market size–, then the probability of the stable matching existence converges to one as the market size grows.[13] The proof is based on a novel modification of the algorithm used in Kojima, Pathak, and Roth (2013), which enables them to declare failure for fewer instances and find a stable matching more often than Kojima, Pathak, and Roth (2013). In addition, the authors find that if the number of couples increases at the same rate as the market size, then the non-existence probability can stay bounded away from zero even in large markets. Therefore, for the guaranteed existence of a stable matching, it is *necessary* for the couple population to grow more slowly than the market as a whole. Given this necessity result, we still do not seem to fully understand whether matching markets work well in real large markets when the proportions of couples are high.[14]

3.2.2. *Complementarity.* Married couples offer a leading example of complementarity, but they are not the only source of it. To the contrary, complementarities are prevalent in various markets. Firms often seek to hire workers with complementary skills. Professional sport teams demand athletes that complement one another in skills or in their positions. Public schools may need to satisfy diversity of their student bodies in terms of skill levels or socio-economic statuses.

Che, Kim, and Kojima (2015) and Azevedo and Hatfield (2014) study this problem and find that, similarly to the case of matching with couples, there is a sense in which the market size helps establish the existence of a stable matching.[15] Both papers study a model of continuum of agents and find that a stable matching exist in the continuum

[13]Strictly speaking, there are some other differences in assumptions between Ashlagi, Braverman, and Hassidim (2014) and Kojima, Pathak, and Roth (2013), so a direct comparison is not obvious. However, these differences appear to be unimportant for the main mechanics of the model and analysis, so I do not elaborate on them here.

[14]In that regard, a series of studies in computer science have made notable advances. Biró, Irving, and Schlotter (2011) consider a certain matching problem with couples that is motivated by medical matching in Scotland, and compare various algorithms including the one used in Scotland as well as the one used in NRMP. Biró, Fleiner, and Irving (2011) and Biró, Manlove, and McBride (2014) provide algorithms which are based on Scarf's algorithm and integer programming, and show that stable matchings are found by these algorithms with high probability even when the proportion of couples is high. For these contributions, see an interdisciplinary survey by Biró and Klijn (2013).

[15]There are several differences between the models of Che, Kim, and Kojima (2015) and Azevedo and Hatfield (2014), so their results are independent of each other. I do not go into detail here and discuss the main features that are common to both papers unless otherwise noted.

economy.[16] Based on that result, these papers study markets with a large but finite number of participants. While the existing impossibility results in the literature (Hatfield and Milgrom, 2005; Hatfield and Kojima, 2008; Sönmez and Ünver, 2010) imply that a stable matching does not always exist in any finite market, there exists a matching that is "almost stable" in the sense that the utility gain from blocking it is negligibly small, and in addition the blocking activity leads to only a small change in the overall matching pattern of the economy.

In addition to the substantive contribution, a novel methodology may be worth mentioning. In Che, Kim, and Kojima (2015), existence of a stable matching is shown based on two findings: (i) a stable matching is characterized as a fixed point of a certain mapping over a functional space that, like the classical tâtonnement process in general equilibrium theory, adjusts demand and supply of positions, and (ii) such a fixed point exists given the continuity of the tâtonnement mapping, where the continuum population assumption plays a crucial role for establishing the continuity property.[17] The existence of a fixed point is established by the Kakutani-Fan-Glicksberg fixed point theorem, a generalization of Kakutani fixed point theorem to function spaces. While this type of technique is not uncommon in economics and game theory, it appears to be new to the matching literature.

A recent paper by Nguyen and Vohra (2014) explores another approach. They consider a situation in which the social planner can force some hospitals to accommodate an excessive number of applicants by increasing their capacities. The authors provide a bound on the increase of positions such that the existence of a stable matching is guaranteed. In the case of couples, even if a given instance has no stable matching, there is another instance with a small change of seats – at most 4 seats in each hospital and at most 9 seats in total– where a stable matching exists. Thus, as long as the social planner can change the number of positions of some hospitals, say through mandate or subsidy, then a stable matching exists.[18]

[16]It is worth mentioning the preceding research by Azevedo and Leshno (2014). They, like Che, Kim, and Kojima (2015), set up a model with a finite number of hospitals and a continuum of doctors. The main difference is that Azevedo and Leshno (2014) assume that hospitals have responsive preferences, which implies the absence of complementarity.

[17]The construction of the tâtonnement-like mapping is inspired by those in the literature such as Adachi (2000), Hatfield and Milgrom (2005), and Echenique and Oviedo (2006), but some modifications are needed to account for the continuum population.

[18]Nguyen, Peivandi, and Vohra (2014) use a similar idea of violating feasibility in a random object allocation setting and obtain a result of a similar flavor.

3.3. **Object allocation.** We now turn our attention to the object allocation setting, i.e., the problem in which the social planner has a number of indivisible objects and aims to allocate them to the members of the society. In this setting, stability is not the main requirement because no agent has a resource with which she can block the prescribed assignment.[19] Instead, other properties such as efficiency, fairness, and incentive compatibility are of the primary importance.

In this setting, one of the most well-known mechanisms is the **random priority** (also known as random serial dictatorship) mechanism, **RP** henceforth. In this mechanism, a serial order over all the applicants is determined by a fair lottery. Then, under the realized serial order, the first applicant receives her stated first choice object, the second applicant receives his stated first choice among the remaining objects, and so forth. In addition to being very simple to describe and easy to implement, RP has desirable properties. First, it is strategy-proof. Second, this mechanism satisfies the "equal treatment of equals" property, a form of fairness which means that two applicants whose stated preferences are identical to each other receive the same random assignment. Third, it is ex post efficient in the sense that the final (deterministic) outcome after the lottery has been resolved is Pareto efficient.

However, despite its ex post efficiency, the RP mechanism may produce a random assignment that is inefficient from an *ex-ante* point of view. To see this point, consider the following example due to Bogomolnaia and Moulin (2001) with a slight modification.

Example 3. Let there be four applicants with preferences given by

$$\succ_{i_1} : A, B,$$
$$\succ_{i_2} : A, B,$$
$$\succ_{i_3} : B, A,$$
$$\succ_{i_4} : B, A,$$

and two hospitals A and B, each with only one seat to be allocated. As we are considering the object allocation setting in this section, let these hospitals be mere objects to be consumed by the applicants, and thus have no intrinsic preferences.[20]

[19]However, note that stability can be interpreted as a fairness property if the objects are endowed with priorities over the agents as in the case of the school choice problem (Abdulkadiroğlu and Sönmez, 2003).

[20]Although hospitals have their own preferences in many doctor-hospital matching markets, we keep using the term "hospital" here for convenience. Note also that in some markets such as U.K., hospitals preferences are not elicited and hospitals are treated as mere objects to be assigned to doctors (Biró,

Let me now describe how RP works in this example. If the realized serial order is i_1, i_2, i_3, i_4, then i_1 receives A, i_2 receives B, and i_3 and i_4 receive the outside option. If the realized order is i_4, i_3, i_2, i_1, then i_4 receives B, i_3 receives A, while i_2 and i receive the outside option. Observe that, under the first serial order, doctor i_2 receives the seat at B although it is her less preferred hospital. Also observe that, under the second serial order, doctor i_3 receives the seat at A although it is his less preferred hospital. Moreover, B is the first choice for i_3 while A is the first choice for i_2. Therefore, if they could exchange probability shares of receiving these hospitals with each other at the ex ante stage, then both doctors would be better off. In other words, RP leaves some room for mutually beneficial trades in probability shares (although not in the final pure outcomes), and thus it is not ex ante efficient. □

To improve efficiency in this problem, Bogomolnaia and Moulin (2001) propose a new mechanism. It is called the **probabilistic serial (PS)** mechanism, and it is based on a procedure called the "simultaneous eating algorithm" as follows. Imagine that each doctor reports her preferences, and the doctors (or more precisely, proxies on the doctors' behalf) simultaneously engage in "eating" at the speed of one in time interval $[0, 1]$. What do the doctors eat? Each doctor eats an infinitesimal amount of probability shares to get into a hospital of her choice. That is, at each point in time interval $[0, 1]$, each doctor increases the probability to receive her first choice hospital by a unit speed. If many doctors desire the same hospital, the sum of the probability shares of that hospital eaten away may reach its supply of the seats before the ending time of the algorithm, i.e., time one. If such a situation occurs, then the algorithm prohibits doctors from eating further from that hospital, and has them eat from their most preferred remaining hospitals. At time one, each doctor has eaten probability shares from various hospitals. These probability shares indeed specify a valid probability distribution over the objects because by construction each of them is nonnegative and they sum up to one. The PS mechanism returns this probability share profile as its output.

Bogomolnaia and Moulin show that PS eliminates the kind of inefficiency of RP we saw in Example 3. At the PS outcome in the above example, doctors i and j are admitted to hospital A with probability $1/2$ while the other two doctors get into hospital B with probability $1/2$. These probabilities are more preferred to the RP assignments by all doctors, and it is efficient in the sense that no further improvement for everyone is possible. More generally, Bogomolnaia and Moulin (2001) show that the PS assignment

Irving, and Schlotter, 2011). In other applications such as school choice, it is common for schools to be treated as objects (Abdulkadiroğlu and Sönmez, 2003).

always satisfies an efficiency property called ordinal efficiency (or sd-efficiency), which is the property that it is not first-order stochastically dominated by any other random assignment.[21] Moreover, the PS assignment is fair in the sense that an agent's own random assignment (weakly) first-order stochastically dominates those of others with respect to her own preference. This property is called envy-freeness, and it implies equal treatment of equals, the main fairness property of RP (by contrast, RP does not necessarily satisfy envy-freeness).

PS has one major drawback, however: This mechanism is not strategy-proof. Unfortunately, Bogomolnaia and Moulin (2001) demonstrate that there is no mechanism that satisfies strategy-proofness, ordinal efficiency, and equal treatment of equals. Both RP and PS satisfy two of these three properties, and there is no way to improve upon them to satisfy all of them.

Given this impossibility result, it is important to study the tradeoffs between RP and PS in the kind of applications that are of practical interest. Pathak (2007) conducts simulations based on the data from New York City's high-school match. He computes the RP and PS assignments under the submitted preferences in NYC's supplementary round of school seat assignment. He finds that the PS assignment places more applicants to their preferred choices than RP, but the difference is very small: For example, out of the total of 8,255 students, about 4,999 receive their top choice in RP, while 5,016 receive their top choice in PS, a difference of mere 17. Based on this finding and the fact that PS is not strategy-proof, he recommends using RP rather than PS in this application.

On the other hand, Kojima and Manea (2010) show that PS becomes non-manipulable in large markets. More specifically, given any doctor and her preferences, if hospital capacities are sufficiently large, then PS makes it a dominant strategy to report her true preferences (note that truthtelling is an exact, as opposed to approximate, dominant strategy in this statement). To get an intuition for this result, observe that misreporting in PS has two effects. First, given the same set of available hospitals, reporting false preferences may prevent a doctor from eating her most preferred available hospital. Second, reporting false preferences can delay the expiration date of some hospital (that is, the period at which the hospital gets completely eaten away) by affecting the overall eating schedules of the doctors. The first effect always weakly hurts the misreporting doctor because it may prevent her from eating her most preferred remaining hospital, but the second effect can benefit the doctor by allowing her to eat probability shares of her preferred hospital later in the algorithm than under truthtelling. But the second effect becomes small in

[21]The term "sd-efficiency" is an abbreviated expression for stochastic-dominance efficiency.

large markets because the expiration dates of the hospitals are pinned down by eating behavior of many doctors. On other other hand, the cost of eating suboptimally does not become small. Thus, the first negative effect dominates the second (potentially positive) effect, which establishes the large-market incentive compatibility of PS.

These studies are aimed at understanding the tradeoffs between RP and PS and, interestingly, they show that shortcomings of both mechanisms become small in large markets. A natural question is why this is the case, and whether one can still make a recommendation between these mechanisms on the basis of some other properties.

Che and Kojima (2010) study this issue and show that the tradeoffs between these mechanisms vanish in large markets. More precisely, they demonstrate that RP and PS assignments converge to the same limit as the market size becomes large. Therefore, in the school choice application, as long as the district is large as in the case of NYC, there is little difference between RP and PS. In fact, this result implies that both mechanisms approximately achieve all the desirable properties mentioned above, i.e., ordinal notions of efficiency, incentive compatibility, and fairness. Just as in the two-sided matching case, market size helps the social planner achieve desirable properties that are impossible in an exact form for markets of arbitrary sizes. For instance, consider replica economies of Example 3, where in the q-fold replica economy, there are supplies of q seats in each hospital, as well as q copies of each doctor in the base economy. The probability that a doctor is matched to her less preferred hospital is positive for all q but approaches zero as $q \to \infty$, as can be seen in Figure 3.

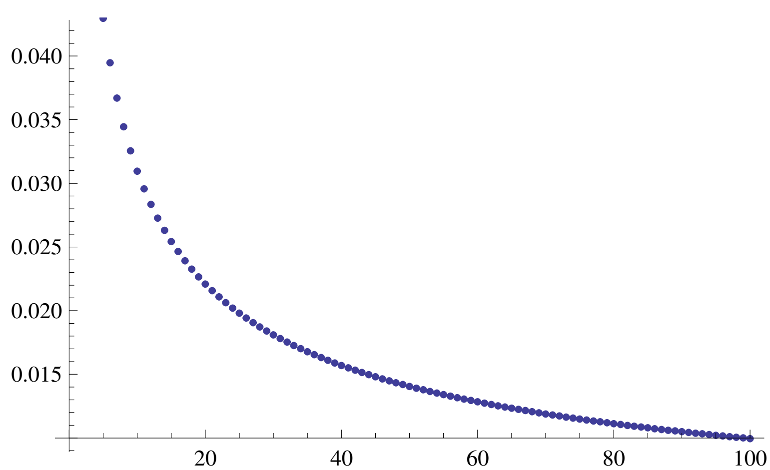

FIGURE 3. Horizontal axis: Market size q. Vertical axis: probability for each doctor to be matched to her second choice hospital. Reproduced from Che and Kojima (2010).

3.3.1. *Further developments.* While the above contributions have identified some senses in which RP and PS mechanisms both achieve desirable properties in certain large markets, they are far from being the last word on this problem.

A notable contribution is made by Manea (2009). He considers allocation of n different types of objects to n agents, where each agent should be allocated exactly one object. In this setting, he shows that the proportion of preference profiles under which the RP random assignment is ordinally inefficient converges to one as n goes to infinity. This result may sound contradictory to Che and Kojima (2010), so two points are worth some discussion in order to to understand these results better. First, there are different ways in which we can model large markets: In Kojima and Manea (2010) and Che and Kojima (2010), the number of the types of the objects are fixed and the supply of each type of object increases, while in Manea (2009), it is the *types* of objects that increases while the supply of each type of good is fixed at one. Which of the "large markets" describes a given market better depends on the application in question, and thus the policy implication should be derived with some caution. Second, the notion of ordinal efficiency used by Manea (2009) is binary in the sense that a random assignment is judged to be efficient if and only if there is *no room for improvement at all*, while the other contributions such as Pathak (2007) and Che and Kojima (2010) use approximate efficiency, i.e., they judge a random assignment to be close to being efficient as long as the room for an improvement is small.

Another recent contribution is related to the asymptotic equivalence result of Che and Kojima (2010). Their approach leaves unanswered the question of whether the equivalence depends crucially on the algorithmic features of these mechanisms or there is a deeper reason for the equivalence. Moreover, from a practical point of view, it is reasonable to ask whether there are other mechanisms that satisfy good finite-market properties in efficiency, incentives, and fairness (or at least some of them) and behave differently from RP or PS in large markets, which may leave some scope for improvements over RP and PS. A recent study by Liu and Pycia (2013) investigates this issue. They demonstrate that asymptotically efficient, symmetric, and asymptotically strategy-proof mechanisms converge to the same limit as the market size becomes large. Therefore, mechanisms such as RP and PS exhaust most of the good properties, and at least in large markets, little room for further improvement is left. Theoretically, it also shows that the asymptotic equivalence as identified by Che and Kojima (2010) is not an isolated coincidence for these two mechanisms, but an instance of a more general phenomenon.

3.4. **Discussion on the "large market methodology".** As detailed in the previous sections, considering large markets allows us to overcome many existing impossibility results by establishing approximate properties in various contexts, ranging from incentives to existence and efficiency. Given that so many positive results hold in large markets, some researchers have raised the following concern to me: Isn't the large market approach "too permissive," and this methodology will result in a theory with the "anything goes" property, i.e., any conclusion can be drawn? I think that these criticisms are unfounded. This section explains why.

The above two criticisms are closely related to each other, but for the sake of argument, let me begin with the first criticism that the large market modeling strategy is "too permissive". This criticism appears to suggest that there are "too many" positive results in large markets while there are negative results in finite markets. I do not think that this is a valid criticism because there is no reason that there should be a "right proportion" of positive or negative results except for the amount dictated by reality. If many problems of matching mechanisms in small markets indeed vanish in large markets, then that fact simply uncovers the nature of reality, and not any drawback of the modeling approach. Unless the positive results in those large market studies are shown to depend on unrealistic assumptions or contradict data, this line of criticism does not have any ground.

This leaves us with the second concern, which may be summarized as the "anything goes" claim. That is, the degree of freedom in modeling choice is too high, and any positive result one wishes to obtain can be obtained by choosing a model in a clever manner.

While this claim could be valid in principle, it does not appear to be supported by what has happened in the research in large matching markets so far. For instance, take a claim that strategic misreporting vanishes in large markets. As we have seen before, this claim has been shown to be true for stable mechanisms in various settings. However, this statement highly depends on the nature of the mechanism (stability), and thus it is not a blanket claim that incentives are unimportant in large markets. To see this point, we consider the so-called **Boston mechanism**, which was used in Boston for student placement before it was replaced by DA. The mechanism proceeds as follows:

- Step 1: Each doctor applies to her first choice hospital. Each hospital rejects the lowest-ranking doctors in excess of its capacity and all unacceptable doctors.

In general, for any $t = 1, 2, \dots,$

- Step t: Each doctor who was rejected in the last step proposes to her next highest choice. Each hospital considers these doctors, *only as long as* there are vacant positions not filled by doctors who are already matched by the previous steps, and rejects the lowest-ranking doctors in excess of its capacity and all unacceptable doctors.

This algorithm terminates when every unmatched doctor (if any) has been rejected by every hospital acceptable to her. The main difference of this mechanism from DA is that in each step of the algorithm, doctors who are not rejected are accepted *immediately* rather than in a deferred manner.[22]

Under this mechanism, doctors may have incentives to misrepresent their preferences even in large markets. To see the reason for this, consider a situation where a doctor's first choice is so popular that she is unlikely to be matched if she truthfully ranks it as her top choice in her reported preference. Also assume that her second choice is reasonably popular and thus she is unlikely to be matched to it if she applies to it as her second choice, but there is significantly better chance to be accepted if she applies to it in the first step of the algorithm (note that this is a likely scenario under the Boston mechanism while not in DA). Then she is made better off by misreporting her preference by listing her true second choice as her stated first choice.

Note that in order to make an inference just illustrated above, the doctor only needs to know the overall popularity of the hospitals. Such aggregate information is often available to each doctor even in large markets, if with some noise. And note that the argument presented above holds even if the doctor's information is noisy. Indeed, an example of Kojima and Pathak (2009) shows that this argument goes through with preferences being private information of each doctor, showing that manipulation incentives does not vanish even as the market size goes to infinity.

The preceding discussion demonstrates that it is *not true* that all problems go away in large markets. On the contrary, the large market modeling approach allows us to *distinguish mechanisms* based on approximate desiderata. By contrast, if researchers are constrained to use only traditional binary criteria such as exact strategy-proofness or exact efficiency, too many mechanisms fail such stringent requirements, and we end up being unable to capture the crucial difference between mechanisms that are close to achieving good properties and those that are far from doing so. Failing to capture such difference is problematic because the distinction between good and bad mechanisms seems to be as important as (if not more important than) the distinction between perfect and less-than-perfect ones. Therefore, contrary to the view that the large market methodology makes

[22]Due to this difference from the deferred acceptance algorithm, the Boston mechanism is sometimes referred to as the immediate acceptance algorithm.

market design irrelevant (i.e., all problems go away anyway), this methodology *informs us better* to tell good mechanisms from bad ones.

A recent paper by Hatfield, Kojima, and Narita (2015) makes this point in a different context. The question studied is how the design of a school choice mechanism affects the competitive pressure on schools for improving themselves. They investigate whether a matching mechanism has the property that a school is matched to a weakly more preferred set of students whenever it becomes more preferred by students.[23] This concept formalizes the requirement that a mechanism should never punish a school for becoming more desirable to students. They find that all standard mechanisms, including all stable mechanisms as well as the Boston and the top trading cycles (TTC) mechanisms, fail this property. In a sharp contrast to this negative result, they then consider an approximate version of this property in the large markets and establish that all stable mechanisms satisfy this approximate property, while both the Boston and the TTC mechanism fail even this property. Their result is another case in point, suggesting that considering approximate properties in the large market setting enables us to make sharper statements about desirability of competing mechanisms.

More generally, modeling a market with a large number of participants can serve as a useful modeling device. In studies of decentralized matching markets by Avery and Levin (2010) and Che and Koh (2015), for instance, models with a continuum of students facilitate equilibrium characterization. Echenique and Pereyra (2014) analyze a finite population model for decentralized markets, but analyze its limit behavior as the market size converges to infinity, which allows the authors to tell which predictions of the finite population model are robust and remains relevant in large markets. For these papers, a large market assumption is a natural modeling approach.

For all these reasons, criticisms toward the large market methodology seem to be unfounded. For market design, this approach allows us to distinguish good mechanisms from bad ones based on their approximate performances. More generally, it allows researchers to tell which of the properties found in stylized finite markets are likely to be relevant in real markets. As such, I think that the large market approach should be regarded as a standard tool in matching theory.

[23]This concept is called respecting improvements of school quality. It is an adaptation of a notion due to Balinski and Sönmez (1999) who consider respecting improvements of student quality.

4. Matching with Constraints

Another major issue in matching theory is also related to the fact that the standard theory is not always applicable in practical applications that interest us. In this section, we discuss recent research in matching with constraints. As in the last section, the issue arises in both two-sided matching and object allocation settings.

4.1. **Two-sided matching with constraints.** Real matching markets are often subject to constraints. Medical residency programs may be subject to regulations which restrict the number of positions for different medical specialties. In school choice, it is a common practice to impose certain balance requirements on the student composition in terms of socioeconomic status or academic achievement (Abdulkadiroğlu and Sönmez, 2003).

A leading example of matching with constraints is the medical residency match in Japan. The Japanese government introduced a centralized matching mechanism for medical residents in 2003, and the initial mechanism was the doctor-proposing DA. However, some critics asserted that DA placed too many doctors to hospitals in urban areas such as Tokyo while causing doctor shortages in rural areas. In order to address this complaint, the government introduced a new regulation, which we call a "regional maximum quota" policy. Under this policy, for each region of the country, the number of residents who are placed in that region is required to be at most the regional maximum quota granted for that region. The idea of this regulation is to restrict the number of doctors matched in urban areas such as Tokyo, thereby preventing what some see as too much imbalance of doctor distributions.

In the Japanese application, each region corresponds to one of the 47 prefectures that partition the country, such as Tokyo, Kyoto, and Osaka. Figure 4 presents the magnitude of the regional maximum quota policy. Urban areas suffer from stringent constraints under the policy. For instance, the total numbers of positions advertised by hospitals in Tokyo and Osaka were 1,582 and 860 positions in 2008, respectively, while the government set the regional maximum quotas of 1,287 and 533. Kyoto received an even larger proportional reduction of positions, which offered 353 positions in 2008 while receiving the regional maximum quota of 190, which is almost a 50 percent decline.[24] In total, 34 out of the country's 47 prefectures are subject to binding regional maximum quotas in the sense that the imposed regional maximum quotas are smaller than the total numbers of advertised positions in 2008.

[24]In fact, the magnitude of the planned changes were so large that the government implemented a temporary provision that limited per-year reductions within a certain bound in the first few years of operation.

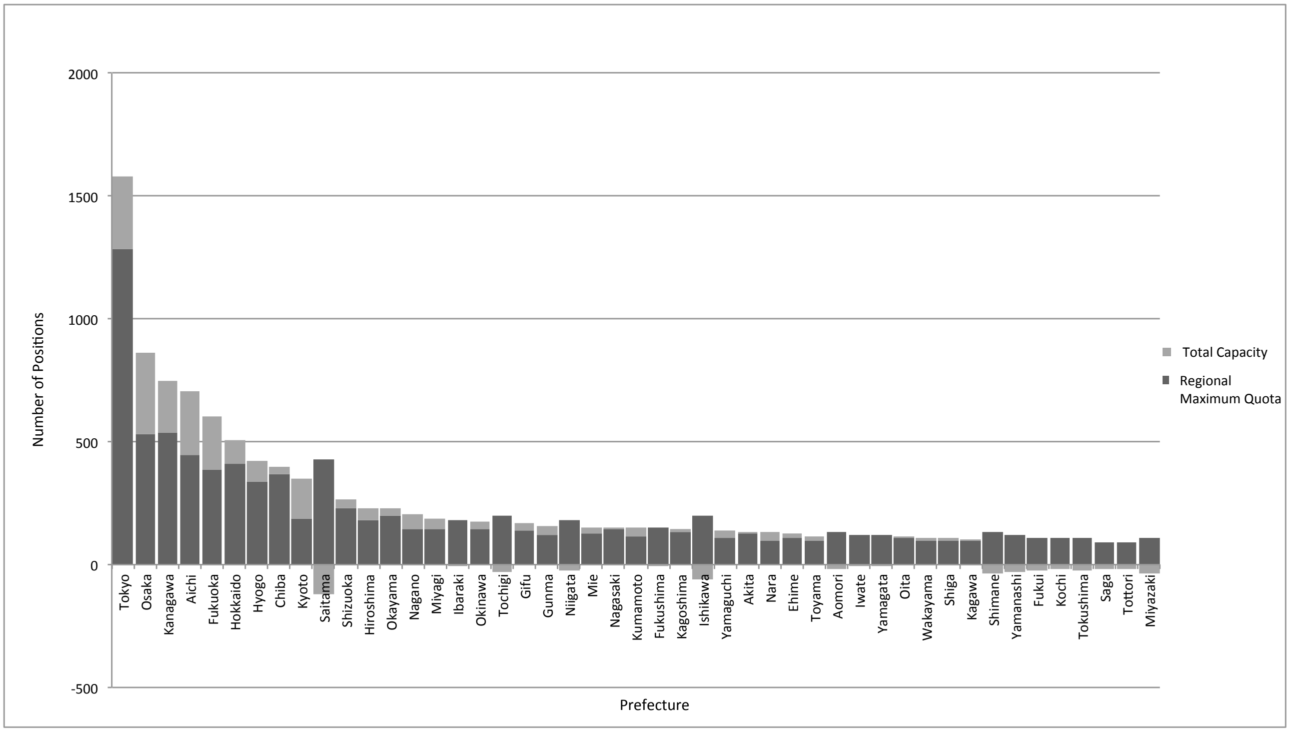

Figure 4. For each prefecture, the total capacity is the sum of advertised positions in hospitals located in the prefecture in 2008. The regional maximum quotas are based on the government's plan in 2008 (Ministry of Health, Labour and Welfare, 2009). Negative values of total capacities in some prefectures indicate the excess amount of regional maximum quotas beyond the advertised positions. Reproduced from Kamada and Kojima (2015).

Policies that are mathematically isomorphic to regional maximum quotas exist in various contexts beyond the Japanese medical match example. For example, Chinese graduate admission is regulated by a policy that imposes a maximum quota on academically oriented masters programs, which is part of an attempt by the government to increase masters educated in professionally oriented programs. In many countries, there are state-sponsored college seats which are subject to the state's budget constraint, so the "region" of the state-sponsored college seats is subject to a maximum quota constraint (examples of Hungary and Ukraine are documented by Biró, Fleiner, Irving, and Manlove (2010) and Kiselgof (2012)). Medical match in the U.K. works in two rounds, the first of which is an allocation of doctors to different regions of the country while the second is the allocation of doctors to specific hospitals within the region. The number of doctors matched in each region is bounded by the capacity for that region in the first round of the allocation scheme.

As illustrated in the previous paragraphs, there are many real matching problems with constraints. Unfortunately, existing solutions used in those markets suffer from inefficiency and instability. The following example taken from Kamada and Kojima (2015) illustrates a typical problem of existing mechanisms in the constrained matching problem, using the Japanese mechanism as a concrete case (they also find similar drawbacks for mechanisms in other constrained matching markets in practice).

Example 4 (An existing mechanism may lead to inefficiency and instability)**.** The Japanese mechanism, called Japan Residency Matching Program (JRMP) mechanism by Kamada and Kojima (2015), is perhaps the simplest and the most prevalent mechanism aimed at addressing regional maximum quotas. The mechanism runs the doctor-proposing DA except that it replaces the actual hospital capacity with an exogenously given artificial capacity, called the target capacity, which is weakly smaller than each hospital's real capacity and, for each region, sums up to the regional maximum quota.

To see the shortcomings of the JRMP mechanism, consider the following (overly simplified) example. Let there be one region, and its regional maximum quota be 10 positions. The capacity of each hospital is 10. Suppose that there are 10 doctors, $i_1, \ldots, i_{10}$. Both hospitals prefer i_1 to i_2 to ... to i_{10} to the outside option. Regarding the doctor preferences, i_1, i_2, and i_3 prefer A to the outside option to B, while all other doctors prefer B to the outside option to A.

Now consider the JRMP mechanism, assuming the target capacities for each hospital to be 5. At the first step of the algorithm, doctors i_1, i_2 and i_3 apply to hospital A, and all the remaining doctors apply to hospital B. Hospital A rejects no one at this round, because the number of applicants is less than its target capacity and all the applicants are acceptable to it. By contrast, hospital B rejects i_9 and i_{10} while accepting other applicants, because the number of applicants exceeds the *target capacity* while not the actual hospital capacity. Given that i_9 and i_{10} find A unacceptable, this algorithm terminates at this point. Thus, the the outcome of this algorithm is given by

$$\mu = \begin{pmatrix} A & B & \emptyset \\ i_1, i_2, i_3 & i_4, i_5, i_6, i_7, i_8 & i_9, i_{10} \end{pmatrix}.$$

Consider a matching μ' defined by,

$$\mu' = \begin{pmatrix} A & B \\ i_1, i_2, i_3 & i_4, i_5, i_6, i_7, i_8, i_9, i_{10} \end{pmatrix}.$$

Since the regional maximum quota is still respected, μ' is feasible. Moreover, every agent is weakly better off with i_9, i_{10}, and B being strictly better off than at μ. Hence we conclude that the JRMP mechanism results in an inefficient matching in this example.[25] □

Remark 2. In the above example, the outcome of JRMP is unstable in a certain sense as well. For instance, hospital B and doctor i_9 form a blocking pair while the regional maximum quota is not binding at μ. That is, even after i_9 is matched with B, the total number of doctors in the region is 9, which is less than the regional maximum quota of 10. Although it is beyond the scope of this paper to provide and analyze a formal definition of stability under constraints, the outcome of the JRMP mechanism clearly violates any reasonable notion of stability. See Kamada and Kojima (2015, 2014b) for detail.

As mentioned earlier, Kamada and Kojima (2015) study other examples such as graduate school admission in China, medical match in the United Kingdom, and teacher matching in Scotland. They find that all these environments are isomorphic to the problem with regional maximum quotas and mechanisms used in these markets fail efficiency and stability.

The problem with the existing solutions is that they modify a standard mechanism such as DA in one way or another, but they do not appropriately maintain desirable properties of the original mechanism. Thus, instead of making a modification in an undisciplined manner, it is useful to formalize normative properties that need to be satisfied and seek a mechanism that satisfies those properties. Kamada and Kojima (2015) take this approach and present a new mechanism that, for any given input, achieves efficiency, stability, and incentive compatibility, among others. To do so, they define the **flexible deferred acceptance (FDA) algorithm**. This algorithm resembles DA and in fact identical to JRMP except that it has a kind of "waitlist processing" phase during each step of the algorithm, where the target capacities allocated across hospitals in a region are modified in a flexible manner. More specifically, the FDA algorithm proceeds in multiple steps of doctors' applications and hospitals' acceptances/rejections like DA. The algorithm begins with the empty matching, and each step $t = 1, 2, \ldots$ proceeds as follows:

[25]In this example, not all hospitals are acceptable to all doctors. One may wonder whether this is an unrealistic assumption because doctors may be so willing to work that any hospital is acceptable. However, the example can be easily modified so that all hospitals are acceptable to all doctors while some doctors are unacceptable to some hospitals (which may be a natural assumption because, for instance, typically a hospital only lists doctors who they interviewed). Also, in many markets doctors apply to only a small subset of hospitals. In JRMP, for instance, a doctor applies to only between 3 and 4 hospitals on average (Kamada and Kojima, 2015).

- Step t: Each doctor who is not tentatively kept by a hospital applies to her next highest choice (if any).
  - Phase 1: Each hospital considers both new applicants and doctors who are temporarily held from the previous step together, and tentatively accepts its most-preferred acceptable doctors up to its target capacity, put the next preferred acceptable doctors on its waitlist up to its true capacity, and rejects every other doctor.
  - Phase 2: Take the first hospital with respect to an exogenously fixed order. If there is any applicant on its waitlist and the total number of doctors tentatively matched in the hospital's region is strictly smaller than its regional maximum quota, then let the hospital tentatively accept its most preferred applicant from its waitlist. Apply the same procedure to the second hospital (again, with respect to the exogenously fixed order), and so forth (and after the last hospital, return to the first hospital). When there is no more applicant to be processed, reject all the remaining doctors on the waitlist and proceed to the next step.

The algorithm terminates at a step in which no rejection occurs (it is straightforward to verify that this algorithm terminates in a finite number of steps). We define the **FDA mechanism** to be the mechanism which produces the matching at the termination of the above algorithm.

FDA is similar to DA, but guarantees a feasible outcome in the presence of constraints. It differs from the JRMP mechanism in that it lets hospitals accept doctors in a flexible manner, using the idea of waitlists. Note that Phase 1 of the FDA algorithm is identical to a phase of JRMP, but it is followed by Phase 2, where hospitals can tentatively accept additional doctors from the waitlist as long as its region's maximum quota is not full. This phase helps FDA correct the efficiency loss of the JRMP mechanism while still following DA closely. In fact, Kamada and Kojima (2015) demonstrate that this algorithm finds a stable and constrained efficient matching. In Example 4, for instance, the source of inefficiency in JRMP was that doctors i_9 and i_{10} are rejected from hospital B because of the artificially specified target capacity even though neither the hospital's capacity nor the regional maximum quota necessitates rejection. The FDA algorithm eliminates this inefficiency by allowing hospital B to accept doctors i_9 and i_{10} from the waitlist in Phase 2.

Efficiency and stability are arguably among the most important goals for matching market design, but they are by no means the only ones. As discussed in previous sections, incentive compatibility is important given that the matching organizer needs to elicit preference information from the participants. In that regard, Kamada and Kojima (2015) show that FDA is strategy-proof for doctors just as the original DA is.

Finally, when considering changing practical mechanisms, the effect of different mechanisms needs to be carefully traded off, because most policy changes do not cause Pareto-improvement. It turns out that the outcome of FDA is a Pareto-improvement over JRMP for doctors, although not necessarily for all hospitals. That is, each doctor weakly prefers the outcome of FDA to the outcome of JRMP. An interesting corollary of this result in the medical matching context is that the total number of the medical residents who are matched to some hospital residency programs will weakly increase if the mechanism is changed from JRMP to FDA.

One important question is whether the magnitude of improvement by a better mechanism is substantial, and whether the loss from the constraints are significant enough to warrant attention in the first place. As it seems elusive to establish a sharp and general prediction on this question by theory alone, Kamada and Kojima (2015) supplement their theoretical analysis with simulations. In their simulations, preferences of market participants are randomly generated, with the model parameters calibrated to closely match the publicly available data for Japanese medical match. One of their simulations, reproduced in Figure 5, reports the number of matched and unmatched doctors under the unconstrained DA (i.e., DA with no regional maximum quotas), JRMP, and FDA. The simulation shows that the existence of regional maximum quotas can cause substantially more doctors to be unmatched: About 800 out of about 8300 doctors are unmatched in DA while the corresponding number is about 1400 in JRMP. Meanwhile, FDA alleviates much of additional unmatched doctors. FDA leaves about 1000 doctors unmatched, which is an increase of mere 200 additional doctors compared to the case with no constraints, as opposed to 600 additional unmatched doctors in JRMP. See Kamada and Kojima (2015) for more detail including the simulation methods and additional numerical results.

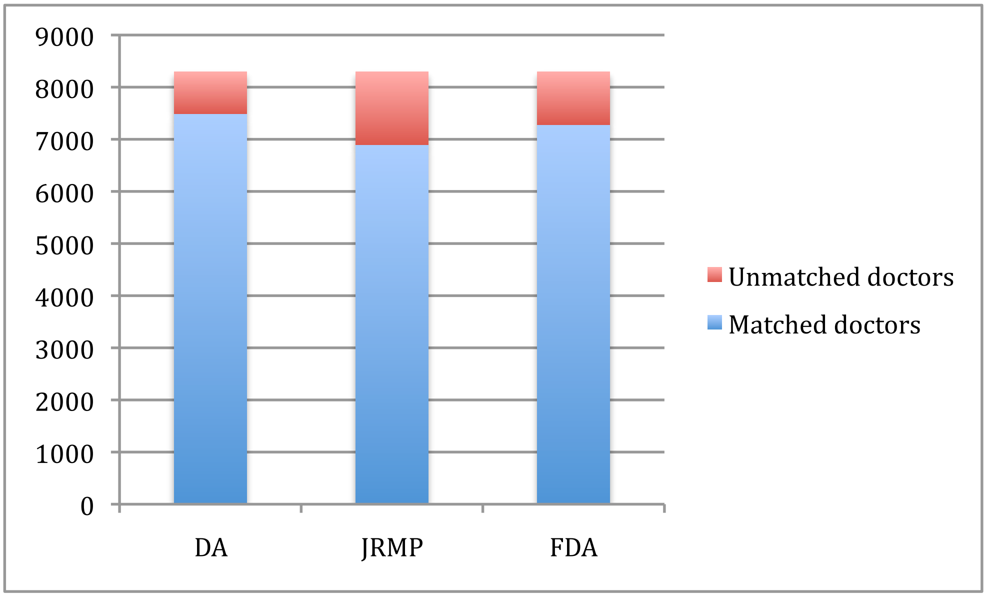

FIGURE 5. The numbers of matched and unmatched doctors under different mechanisms. "DA" refers to the DA outcome under no regional maximum quota constraints.

4.1.1. *Further contributions.* Systematic studies in matching with constraints are still in an early stage of development, so there are a variety of unsolved topics.

One of the basic questions in this area is: What matching should be deemed "desirable"? For efficiency, this question is pretty straightforward because the standard concept of constrained Pareto efficiency is well-defined and has an obvious appeal. For stability, however, the answer to this question is not trivial at all: On the one hand, the standard notion of stability is too demanding because there are cases in which every stable matching in the standard sense violates a constraint; On the other hand, ignoring stability completely seems to be very unappealing given existing evidence that stability is important for a success of matching markets. Is there a reasonable weakening of stability that takes constraints into account while maintaining the spirit of the original concept?

Kamada and Kojima (2014b) study this issue. They define two concepts of stability under constraints. The first concept, which the authors call strong stability, requires (in addition to feasibility and individual rationality) that for any blocking pair, satisfying that blocking pair by letting them match with each other will result in a violation of a constraint. This definition is based on the requirement that when a blocking pair exists, it is "inevitable" because satisfying them would result in infeasibility.

Although this is perhaps the most natural definition of stability under constraints, the authors find that a strongly stable matching does not always exist: In fact, a strongly stable matching can be guaranteed to exist *if and only if* the constraints are trivial.[26] Motivated by this result, they define a less demanding concept, weak stability. They show that weak stability implies efficiency, that it is still strong enough to exclude some unappealing matchings like the JRMP outcome in Example 4, and that it is characterized by several standard axioms. In addition, a weakly stable matching always exists: in fact, FDA produces a weakly stable matching.[27]

Although weak stability has some normative appeal, this concept may still be too permissive for some applications. In fact, FDA satisfies not only weak stability, but also certain distributional balance among hospitals in the same region. Motivated by this fact, Kamada and Kojima (2014a) consider a model in which the social planner has a policy goal on distributional balance of doctors beyond the hard regional maximum quota constraints. They define an intermediate requirement called stability and show that a generalization of FDA is strategy-proof for doctors and finds a stable matching.

Although matching with constraints is a new research topic, related models have been studied in the literature in several contexts. In the medical match in 20th century United Kingdom, some hospitals preferred to hire exactly one female doctor, and an algorithm used there accommodated such hospital preferences (Roth, 1991). In school choice, schools are often subject to diversity constraints in terms of socioeconomic status and academic performance (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu, 2005; Westkamp, 2013; Kojima, 2012; Hafalir, Yenmez, and Yildirim, 2013; Ehlers, Hafalir, Yenmez, and Yildirim, 2014; Echenique and Yenmez, 2013). Models with constraints on *sets* of agents have also been explored in the context of student-project allocations (Abraham, Irving, and Manlove, 2007; Sönmez and Ünver, 2006) and college admission (Biró, Fleiner, Irving, and Manlove, 2010). Kojima, Tamura, and Yokoo (2014) provide a framework based on discrete convex analysis, a branch of discrete mathematics, to obtain various results in these related but different models with a unified technique.

An important and actively studied problem with constraints is the case with *minimum quotas*. Clearly, there are many matching problems in which the constraints take the form of minimum quotas rather than maximum quotas. For example, school districts may require at least a certain number of students are assigned to each school (Biró, Fleiner, Irving, and Manlove, 2010). The cadet-branch matching program organized by

[26]More specifically, the necessary and sufficient condition is that each constraint involves just one hospital or the associated regional maximum quota is zero.

[27]The weak stability concept can be defined under more general constraint structures. In the general case, FDA is not defined, but the authors show the existence of a weakly stable matching using a different method.

United States Military Academy imposes minimum quotas on the number of cadets who must be assigned to each branch (Sönmez and Switzer, 2012; Sönmez, 2013). Minimum quotas are studied by recent contributions such as Ehlers, Hafalir, Yenmez, and Yildirim (2014), Ueda, Fragiadakis, Iwasaki, Troyan, and Yokoo (2012), Goto, Hashimoto, Iwasaki, Kawasaki, Ueda, Yasuda, and Yokoo (2014), and Fragiadakis and Troyan (2013). A stable matching does not always exist in the presence of minimum quotas, but these studies have found various mechanisms that improve upon existing mechanisms while satisfying desirable incentive compatibility properties.

4.1.2. *Connection with matching with contracts.* There is a connection between matching with constraints and matching with *contracts* as defined by Hatfield and Milgrom (2005).[28] This section illustrates how they are related.

Given the additional complications caused by constraints, it is often difficult to see which of the desirable properties in the standard model can be generalized to the model with constraints. Even if a generalization is possible, it may still be unclear to see what kind of mechanism achieves those properties. In the context of matching with regional maximum quotas, it proves useful to consider a hypothetical matching model where doctors are matched with *regions* rather than hospitals. In such a model, the choice pattern implied by hospital preferences and constraints is interpreted to be the result of a "region's preferences." The main advantage of this approach is that it enables us to treat a regional maximum quota as a usual capacity of an agent, i,e., region, and apply results from existing matching models *without* constraints. Then one can show that a stable allocation in the hypothetical model corresponds to a stable matching in the original model with constraints. Moreover, some other desirable properties in the hypothetical model can be translated to corresponding properties in the original model of matching with constraints.

However, this approach necessitates us to use a model of matching with contracts, because the model needs to distinguish matching of the same doctor to two different hospitals even if these hospitals are in the same region. More specifically, we treat these

[28]Fleiner (2003) considers a framework that is more general than Hatfield and Milgrom (2005), while he does not consider some results of interest for our purposes such as strategy-proofness. See also Crawford and Knoer (1981) and Kelso and Crawford (1982) who consider matching with wages, which is a form of contracts. Importantly, Echenique (2012) shows that in many-to-one matching with contracts, if preferences of all hospitals are substitutable, then the model is isomorphic to matching with wages. The models of Kamada and Kojima (2015, 2014a) do not reduce to matching with wages, however, because their models need to use *many-to-many* matching with contracts (Kominers (2012) finds a correspondence similar to Echenique's in many-to-many matching under an additional condition, but his condition is not satisfied in their setting).

two matchings as *different contracts between the same doctor-region pair*. The model of matching with contracts provides a useful language to describe such a distinction. With this technique, Kamada and Kojima (2015, 2014a) show that results such as the existence of a stable matching and strategy-proofness of FDA for doctors follow from corresponding results in matching with contracts.[29] Kojima, Tamura, and Yokoo (2014) use a similar technique to derive results in various models of constraints such as regional maximum quotas, regional minimum quotas, diversity constraints in schools, and cadet-branching problem, all by way of associating them to models of matching with contracts.

4.2. **Object allocation with constraints.** Object-allocation problems are often subject to constraints as well. We return to the model of object allocation problem which we described in Section 3.3, but now we allow for constraints to restrict feasible allocations.

As illustrated earlier, randomization is a leading approach to restore fairness when allocating objects, especially if monetary transfer is not allowed due to legal or ethical or other reasons. However, a major problem with randomization is that even if deterministic mechanisms under consideration have desirable properties such as Pareto efficiency, a randomization over them can lose those properties. The RP mechanism is a case in point, which is a uniform randomization over serial dictatorships with different serial orderings: Although serial dictatorship is Pareto efficient for any serial order, RP can produce an ordinally inefficient random allocation. To remedy such a drawback, a mechanism that directly produces a random assignment is promising (PS is such an example). More generally, consider a random assignment matrix $P = (P_{iA})$ indexed by doctors and hospitals, where entry P_{iA} specifies the probability that doctor i obtains a seat at hospital A. Working directly with such a matrix is useful because it makes it easy to define and study efficiency, fairness, and incentive properties.

It is not clear, however, that such a method is justified. That is, given a random assignment matrix, does there always exist a way to realize it as a lottery over feasible deterministic outcomes?

Budish, Che, Kojima, and Milgrom (2013) tackle this question. They provide a necessary and sufficient condition for this method to be justified. First, using well-known results in the combinatorial optimization literature, they identify a condition on the structure of the constraints that is sufficient to guarantee that every random assignment that satisfies the constraints in expectation can be implemented as a lottery over feasible pure assignments. That condition, called the "bihierarchy" condition, states that the entire structure

[29] They use results from Hatfield and Milgrom (2005), Hatfield and Kojima (2009, 2010), and Hatfield and Kominers (2009, 2014).

can be partitioned into two hierarchical sets (a.k.a. laminar families) of constraints. Then they demonstrate that the same condition is not only sufficient, but also necessary for guaranteeing the implementation in two-sided assignment and matching environments.

Remark 3. As stated above, there is a sense in which the bihierarchy structure is necessary for implementation. The necessity statement, however, is formulated *in terms of the structures of the constraints* for guaranteeing the implementation for all possible random assignments and constraint values. There are cases in which implementation of a given random assignment under a particular constraint values are possible even without the bihierarchy structure. A sufficient condition based on this idea is given by Che, Kim, and Mierendorff (2013).

Budish, Che, Kojima, and Milgrom (2013) apply the above implementation result to several object allocation problems with constraints. First, the implementation result enables them to generalize PS to various cases with maximum constraints, such as the school choice problem with diversity constraints. Just as the original eating algorithm, their generalized algorithm has each agent continuously "eat" her favorite available object over a unit time interval; the only modification to this algorithm is to declare an object "unavailable" to an agent whenever allowing her to eat more of that object will result in a violation of a constraint. The above implementation result implies that the outcome of this eating algorithm can be implemented as a lottery over feasible pure outcomes. In addition, efficiency and fairness properties of the eating algorithm in the original setting extend to their more general environment. In a similar spirit, the implementation result allows generalizing other random assignment mechanisms such as the pseudo-market mechanism by Hylland and Zeckhauser (1979).

4.2.1. *Approximate implementation.* As illustrated above, some types of constraints can be readily accommodated by the bihierarchy condition of Budish, Che, Kojima, and Milgrom (2013), enabling fairly straightforward generalizations of desirable mechanisms from simple settings to environments with those constraints. However, there are also natural constraints that violate their condition. For instance, two maximum constraints that partially overlap with each other cannot be part of one hierarchy, so examples such as school diversity constraints on race and gender cannot be handled by the above approach.[30]

[30] For this statement, I assume that there is another (natural) hierarchy that includes constraints on the number of schools that any one student can attend (typically just one school).

Because the bihierarchy condition is not only sufficient but also "necessary" for guaranteeing implementation of lotteries, so it seems to be elusive to find a more general sufficient condition.

Akbarpour and Nikzad (2014) tackle this question from a perspective closer to the approach of "approximate market design" described in Section 3. They presume that the constraints are divided into *hard* constraints that must always be satisfied and *soft* goals that can be slightly violated if necessary. In school choice, for instance, the physical capacity of a school may be a hard constraint that cannot be changed, while the diversity constraints for different socioeconomic classes could be violated by a small number if necessitated by other requirements. Akbarpour and Nikzad (2014) demonstrate that if the hard constraints have the bihierarchical structure of Budish, Che, Kojima, and Milgrom (2013), and the set of soft constraints satisfies certain conditions, then any given random assignment can be implemented as a lottery over pure assignments that satisfy all the hard constraints while violating soft constraints by a large amount only with small probabilities. Nguyen, Peivandi, and Vohra (2014) take a similar approach to study random assignment problems when the agents have limited preference complementarities.

5. Conclusion

This paper surveyed some of the recent advances in matching theory and its applications to market design. I emphasized the influence of theory to applications and vice versa. In particular, I described the way that recent development of theory has produced positive results in problems in which traditional theory has been negative.

As I mentioned at the outset, I restricted attention to a small subset of topics in matching theory. Inevitably, there are many exciting topics I had to omit. One such topic is dynamic matching. Akbarpour, Li, and Oveis Gharan (2014) and Anderson, Ashlagi, Gamarnik, and Kanoria (2015) study dynamic matching when preferences are dichotomous as in the case of kidney exchange.[31] They find that the benefit from waiting and making the market thick depends on the agents' discount factor and information structure in very striking manners. Baccara, Lee, and Yariv (2015) consider two-sided matching over time with vertically differentiated types and derive an optimal policy. These papers have started to formally study tradeoffs between thickening the market and matching agents quickly, a topic that the standard static matching models did not have a tool to analyze with.

[31] See also an earlier contribution of Ünver (2010), who studies the optimal design of kidney exchange in a dynamic environment.

Matching theory plays an increasingly important role in economic theory while also influencing policies in a fruitful manner. I think that the literature has been advancing with a healthy mix of abstract theory and practical considerations. I hope that this trend will continue and help us understand how to improve our institutions.

References

Abdulkadiroğlu, A. (2005): "College Admissions with Affirmative Action," *International Journal of Game Theory*, 33(4), 535–549.

Abdulkadiroğlu, A., and T. Sönmez (2013): "Matching markets: Theory and practice," *Advances in Economics and Econometrics*, 1, 3–47.

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

Abraham, D. J., R. W. Irving, and D. F. Manlove (2007): "Two Algorithms for the Student-Project Allocation Problem," *Journal of Discrete Algorithms*, 5(1), 73–90.

Adachi, H. (2000): "On a Characterization of Stable Matchings," *Economics Letters*, 68(1), 43–49.

Akbarpour, M., S. Li, and S. Oveis Gharan (2014): "Dynamic matching market design," *mimeo.*

Akbarpour, M., and A. Nikzad (2014): "Approximate random allocation mechanisms," *mimeo.*

Anderson, R., I. Ashlagi, D. Gamarnik, and Y. Kanoria (2015): "A dynamic model of barter exchange," in *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1925–1933. SIAM.

Ashlagi, I., M. Braverman, and A. Hassidim (2014): "Stability in Large Matching Markets with Complementarities," *Operations Research*, 62(4), 713–732.

Ashlagi, I., Y. Kanoria, and J. Leshno (2014): "Unbalanced Random Matching Markets: The Stark Effect of Competition," *mimeo.*

Avery, C., and J. Levin (2010): "Early Admissions at Selective Colleges," *American Economic Review*, 100, 2125–2156.

Azevedo, E. M., and J. W. Hatfield (2014): "Complementarity and Multidimensional Heterogeneity in Matching Markets," *mimeo.*

Azevedo, E. M., and J. D. Leshno (2014): "A Supply and Demand Framework for Two-Sided Matching Markets," *mimeo.*

Baccara, M., S. Lee, and L. Yariv (2015): "Optimal Dynamic Matching," *mimeo.*

Balinski, M., and T. Sönmez (1999): "A tale of two mechanisms: student placement," *Journal of Economic Theory*, 84, 73–94.

Biró, P., T. Fleiner, and R. Irving (2011): "Matching couples with Scarf's algorithm," *Institute of Economics, Hungarian Academy of Sciences*.

Biró, P., T. Fleiner, R. Irving, and D. Manlove (2010): "The College Admissions Problem with Lower and Common Quotas," *Theoretical Computer Science*, 411(34-36), 3136–3153.

Biró, P., R. W. Irving, and I. Schlotter (2011): "Stable matching with couples: an empirical study," *Journal of Experimental Algorithmics*, 16, 1–2.

Biró, P., and F. Klijn (2013): "Matching with couples: a multidisciplinary survey," *International Game Theory Review*, 15(02).

Biró, P., D. F. Manlove, and I. McBride (2014): "The hospitals/residents problem with couples: Complexity and integer programming models," in *Experimental Algorithms*, pp. 10–21. Springer.

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

Budish, E., and E. Cantillon (2012): "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," *American Economic Review*, 102, 2237–2271.

Budish, E., Y.-K. Che, F. Kojima, and P. Milgrom (2013): "Designing random allocation mechanisms: Theory and applications," *American Economic Review*, 103(2), 585–623.

Che, Y.-K., J. Kim, and F. Kojima (2015): "Stable matching in large economies," *mimeo*.

Che, Y.-K., J. Kim, and K. Mierendorff (2013): "Generalized Reduced-Form Auctions: A Network-Flow Approach," *Econometrica*, 81(6), 2487–2520.

Che, Y.-K., and Y. Koh (2015): "Decentralized College Admissions," forthcoming, *Journal of Political Economy*.

Che, Y.-K., and F. Kojima (2010): "Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms," *Econometrica*, 78(5), 1625–1672.

Crawford, V., and E. M. Knoer (1981): "Job Matching with Heterogeneous Firms and Workers," *Econometrica*, 49, 437–450.

Dubins, L. E., and D. A. Freedman (1981): "Machiavelli and the Gale-Shapley algorithm," *American Mathematical Monthly*, 88, 485–494.

Echenique, F. (2012): "Contracts versus salaries in matching," *The American Economic Review*, pp. 594–601.

Echenique, F., and J. Oviedo (2006): "A theory of stability in many-to-many matching," *Theoretical Economics*, 1, 233–273.

Echenique, F., and J. S. Pereyra (2014): "Strategic Complementarities and Unraveling in Matching Markets," forthcoming, *Theoretical Economics*.

Echenique, F., and M. B. Yenmez (2013): "How to control controlled school choice," *forthcoming, American Economic Review*.

Ehlers, L., I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim (2014): "School Choice with Controlled Choice Constraints: Hard Bounds versus Soft Bounds," *Journal of Economic Theory*, 153, 648–683.

Fleiner, T. (2003): "A Fixed-Point Approach to Stable Matchings and Some Applications," *Mathematics of Operations Research*, 28, 103–126.

Fragiadakis, D., and P. Troyan (2013): "Market Design under Distributional Constraints: Diversity in School Choice and Other Applications," mimeo.

Gale, D., and L. S. Shapley (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–15.

Goto, M., N. Hashimoto, A. Iwasaki, Y. Kawasaki, S. Ueda, Y. Yasuda, and M. Yokoo (2014): "Strategy-proof matching with regional minimum quotas," in *Thirteenth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2014)*, pp. 1225–1232.

Hafalir, I. E., M. B. Yenmez, and M. A. Yildirim (2013): "Effective affirmative action in school choice," *Theoretical Economics*, 8(2), 325–363.

Hatfield, J. W., and F. Kojima (2008): "Matching with Contracts: Comment," *American Economic Review*, 98, 1189–1194.

——— (2009): "Group Incentive Compatibility for Matching with Contracts," *Games and Economic Behavior*, pp. 745–749.

——— (2010): "Substitutes and Stability for Matching with Contracts," *Journal of Economic Theory*, pp. 1704–1723.

Hatfield, J. W., F. Kojima, and Y. Narita (2015): "Promoting school competition through school choice: a market design approach," *Available at SSRN 1984876*.

Hatfield, J. W., and S. D. Kominers (2009): "Contract Design and Stability in Matching Markets," mimeo.

——— (2014): "Contract Design and Stability in Many-to-Many Matching," Discussion paper, Unpublished mimeo.

Hatfield, J. W., and P. Milgrom (2005): "Matching with Contracts," *American Economic Review*, 95, 913–935.

Hylland, A., and R. Zeckhauser (1979): "The Efficient Allocation of Individuals to Positions," *Journal of Political Economy*, 87, 293–314.

Immorlica, N., and M. Mahdian (2005): "Marriage, Honesty, and Stability," *SODA 2005*, pp. 53–62.

Kamada, Y., and F. Kojima (2014a): "General Theory of Matching under Distributional Constraints," mimeo.

——— (2014b): "Stability Concepts in Matching with Distributional Constraints," mimeo.

——— (2015): "Efficient matching under distributional constraints: Theory and applications," *The American Economic Review*, 105(1), 67–99.

Kelso, A., and V. Crawford (1982): "Job matching, coalition formation, and gross substitutes," *Econometrica*, 50, 1483–1504.

Kiselgof, S. (2012): "Matching Practices for Universities–Ukraine," http://www.matching-in-practice.eu/index.php/matching-in-practice/universities/ukraine.

Knuth, D. E., R. Motwani, and B. Pittel (1990): "Stable Husbands," *Random Structures and Algorithms*, 1, 1–14.

Kojima, F. (2012): "School choice: Impossibilities for affirmative action," *Games and Economic Behavior*, 75(2), 685–693.

Kojima, F., and M. Manea (2010): "Incentives in the Probabilistic Serial Mechanism," *Journal of Economic Theory*, 145(1), 106–123.

Kojima, F., and P. A. Pathak (2009): "Incentives and Stability in Large Two-Sided Matching Markets," *American Economic Review*, pp. 608–627.

Kojima, F., P. A. Pathak, and A. E. Roth (2013): "Matching with Couples: Stability and Incentives in Large Markets," *Quarterly Journal of Economics*, 128, 1585–1632.

Kojima, F., A. Tamura, and M. Yokoo (2014): "Designing Matching Mechanisms under Constraints: An Approach from Discrete Convex Analysis," .

Kominers, S. D. (2012): "On the correspondence of contracts to salaries in (many-to-many) matching," *Games and Economic Behavior*, 75(2), 984–989.

Lee, S. (2012): "Incentive Compatibility of Large Centralized Matching Markets," Unpublished mimeo.

Liu, Q., and M. Pycia (2013): "Ordinal Efficiency, Fairness, and Incentives in Large Markets," Unpublished mimeo.

Manea, M. (2009): "Asymptotic Ordinal Inefficiency of Random Serial Dictatorship," *Theoretical Economics*, 4(2), 165–197.

Milgrom, P. (2009): "Assignment messages and exchanges," *American Economic Journal: Microeconomics*, 1(2), 95–113.

Milgrom, P., and I. Segal (2014): "Deferred-Acceptance Auctions and Radio Spectrum Reallocation," mimeo.

Ministry of Health, Labour and Welfare (2009): "On opinions expressed in the public comment procedure (in Japanese)," http://www.mhlw.go.jp/shingi/2009/04/dl/s0423-8a.pdf.

Nguyen, T., A. Peivandi, and R. Vohra (2014): "One-sided matching with limited complementarities," .

Nguyen, T., and R. Vohra (2014): "Near Feasible Stable Matchings with Complementarities," .

Pathak, P. A. (2007): "Lotteries in Student Assignment," Ph.D. thesis, Harvard University.

Pathak, P. A., and T. Sönmez (2008): "Leveling the playing field: Sincere and sophisticated players in the Boston mechanism," *The American Economic Review*, 98(4), 1636–1652.

Roth, A. E. (1982): "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 7, 617–628.

——— (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.

Roth, A. E. (1991): "A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom," *American Economic Review*, pp. 415–440.

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

——— (2007): "Repugnance as a Constraint on Markets," *Journal of Economic Perspectives*, 21, 37–58.

——— (2008): "Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions," *International Journal of Game Theory*, 36, 537–569.

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

Roth, A. E., T. Sönmez, and U. Ünver (2004): "Kidney Exchange," *Quarterly Journal of Economics*, 119, 457–488.

——— (2005): "Pairwise Kidney Exchange," *Journal of Economic Theory*, 125, 151–188.

——— (2007): "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences," *American Economic Review*, 97, 828–851.

Roth, A. E., and M. A. O. Sotomayor (1990): *Two-sided matching: a study in game-theoretic modeling and analysis*. Cambridge University Press, Cambridge.

Sönmez, T. (2013): "Bidding for army career specialties: Improving the ROTC branching mechanism," *Journal of Political Economy*, 121(1), 186–219.

Sönmez, T., and T. Switzer (2012): "Matching with (branch-of-choice) contracts at United States Military Academy," forthcoming, *Econometrica*.

Sönmez, T., and M. Ünver (2006): "School Choice with Aggregate Capacity Constraints," mimeo.

Sönmez, T., and U. Ünver (2009): "Matching, Allocation, and Exchange of Discrete Resources," forthcoming, Handbook of Social Economics, eds. Jess Benhabib, Alberto Bisin, and Matthew Jackson, Elsevier.

Sönmez, T., and U. Ünver (2010): "Course Bidding at Business Schools," *International Economic Review*, 51(1), 99–123.

Ueda, S., D. Fragiadakis, A. Iwasaki, P. Troyan, and M. Yokoo (2012): "Strategy-proof mechanisms for two-sided matching with minimum and maximum quotas," in *AAMAS*, pp. 1327–1328.

Ünver, M. U. (2010): "Dynamic kidney exchange," *The Review of Economic Studies*, 77(1), 372–414.

Westkamp, A. (2013): "An Analysis of the German University Admissions System," *Economic Theory*, 53(3), 561–589.