



The Mechanism Design Approach to Student Assignment

Author(s): Parag A. Pathak

Source: *Annual Review of Economics*, 2011, Vol. 3 (2011), pp. 513-536

Published by: Annual Reviews

Stable URL: <https://www.jstor.org/stable/42940197>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



Annual Reviews is collaborating with JSTOR to digitize, preserve and extend access to *Annual Review of Economics*

JSTOR

The Mechanism Design Approach to Student Assignment

Parag A. Pathak

Department of Economics, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02142, and NBER, Cambridge,
Massachusetts 02138; email: ppathak@mit.edu

Annu. Rev. Econ. 2011. 3:513–36

First published online as a Review in Advance on
May 17, 2011

The *Annual Review of Economics* is online at
economics.annualreviews.org

This article's doi:
10.1146/annurev-economics-061109-080213

Copyright © 2011 by Annual Reviews.
All rights reserved

JEL codes: C78, I2

1941-1383/11/0904-0513\$20.00

Keywords

school choice, market design, matching

Abstract

The mechanism design approach to student assignment involves the theoretical, empirical, and experimental study of systems used to allocate students into schools around the world. Recent practical experience designing systems for student assignment has raised new theoretical questions for the theory of matching and assignment. This article reviews some of this recent literature, highlighting how issues from the field motivated theoretical developments and emphasizing how the dialogue may be a road map for other areas of applied mechanism design. Finally, it concludes with some open questions.

1. INTRODUCTION

In the past decade, there has been a great deal of activity and excitement among economists who study the design of systems used to assign students to schools. Theory has matured to a point at which economists have been able to advise a handful of U.S. school districts on their allocation procedures, and hundreds of thousands of students have been assigned to schools via new mechanisms. Moreover, the initial evidence suggests that these mechanisms are improvements over previous alternatives.

The potential for mechanism design and matching theory to illuminate the practical design of student assignment systems was brought to light by Abdulkadiroğlu & Sönmez (2003). This article summarizes features of some public school choice plans in the United States, describes desiderata for assignment mechanisms, and proposes two alternative mechanisms. These alternative mechanisms are adaptations of widely studied mechanisms in the literature on matching and assignment markets, dating back to seminal contributions by Gale & Shapley (1962) and Shapley & Scarf (1974).

After Abdulkadiroğlu & Sönmez (2003) appeared, a reporter for the *Boston Globe* contacted the authors. The newspaper published an article describing the flaws with Boston's student assignment system. The article also explained how alternatives might share features of the system used to assign medical students to residency programs in the United States, known as the National Residency Matching Program (NRMP) (Cook 2003). Around the same time, in May 2003, Alvin E. Roth was contacted by officials in the New York City (NYC) Department of Education for advice on their high school admissions process. As part of the Children's First Initiative, the mayor and chancellor centralized the organization and governance of the NYC public schools. One major change was the creation of over 100 new small high schools, which dramatically increased the supply of choice options. Although the district had experimented with various forms of school choice for decades and had developed procedures to assign students to schools, many aspects of the plan were problematic and generated widespread dissatisfaction (Herszenhorn 2004). Some NYC officials were aware of the NRMP, which had been reformed in the mid-1990s (Roth & Peranson 1999), and contacted Roth wondering whether similar ideas could be employed to place high school students.

The result was a collaboration involving Abdulkadiroğlu, Roth, Sönmez, and myself in various combinations to assist Boston and NYC in aspects of the design of their new mechanism. Confronting aspects of the existing theory with real-world challenges led to new theoretical problems and issues. The initial article of Abdulkadiroğlu & Sönmez (2003), together with practical developments in Boston and NYC, ushered in a new decade of research on the mechanism design approach to student assignment.

The purpose of this article is to review some of these developments focusing on the interplay between work in the field designing mechanisms and theory. At the outset, I emphasize that this article is not a comprehensive literature review. Rather, I focus on a subset of issues that have been motivated from field experiences and hence only a subset of contributions. As with most selective surveys, my own papers probably get more attention than they deserve.

The organization of this article is as follows: Section 2 provides background on school choice and describes the canonical model. The next three sections describe theoretical issues that arose from practice, related to how students are prioritized at schools

(Section 3), market size (Section 4), and heterogeneous levels of sophistication (Section 5). Section 6 concludes with some open questions.

2. BACKGROUND

2.1. Rationale for School Choice

School choice is a popular and widespread education reform in urban districts.¹ Most U.S. states have open enrollment policies, and there are estimates that the total enrollment in these plans is greater than enrollment in charter schools and voucher programs (Holme & Wells 2008). In a choice plan, families express preferences over what schools their children may want to attend. Using this information, the district assigns children to schools according to various objectives. In residence-based or neighborhood school assignment systems, families express their preferences over schools through their choice of residential location. Critics of neighborhood school assignment challenge that only wealthier families are able to purchase the rights to better schools for their children. As a result, neighborhood school assignment may lead to school segregation and has the potential to perpetuate inequalities. School choice, however, may weaken the link between the housing market and schooling options and lead to more equitable educational opportunities.

The origins of school choice in the United States can be traced back to the history of school desegregation. Despite the legal end of segregation in public schools following the Supreme Court ruling in *Brown v. Board of Education* in 1954, in the subsequent decades many urban districts continued to be de facto segregated. As a result, throughout the 1970s and 1980s, school districts implemented mandatory busing plans under court supervision. One of the most controversial busing plans was in Boston Public Schools. In 1974, Federal Judge W. Arthur Garrity ruled that the school committee “knowingly carried out a systematic program of segregation.” He required that Boston follow Massachusetts law requiring any school with a student enrollment that was more than half white be balanced by race. In 1975, Harvard Education School Professor Charles Willie served as a court-appointed master in the case. For the next 15 years, in a series of court proceedings, the Boston School Committee and Judge Garrity wrestled with the appropriate way to assign students to school. School assignment became an intense political struggle, and the city even erupted in violence at various points. Throughout the period, there was a drop in enrollment in Boston’s public schools, and this trend was common in other urban districts as wealthier families left urban districts for the suburbs (Baum-Snow & Lutz 2011, Boustan 2010).

Mortified observing Boston’s experience with desegregation, residents and public officials in nearby Cambridge envisioned a system of open enrollment to pre-empt court-ordered busing (Fiske 2002). In March 1981, the district abolished all neighborhood zones and adopted a comprehensive school choice plan in which a student could apply to any school in the city. Given his experience in Boston’s desegregation case, Willie was retained as a consultant and chief architect of Cambridge’s controlled choice plan, one of the first plans in the nation. With this test case in hand, Willie & Alves (1987) developed a

¹This review focuses primarily on U.S. choice plans, although plans for various secondary and postsecondary schooling options are widespread around the world (see, e.g., Balinski & Sönmez 1999 on Turkish college admissions, Burgess et al. 2009 on secondary school admissions in England, Lavy 2010 on middle school choice in Tel-Aviv, Israel, and Chiu & Weng 2009 on college admissions in China).

choice plan for Boston's public schools, following the last Garrity ruling in 1988. Both Cambridge's and Boston's plans used race as a factor to obtain balance at schools. The initial principles of the choice plan aspired toward having schools of choice, which were also diversified, in that there was a balanced distribution of students across racial, ethnic, and socioeconomic characteristics. Controlled choice was advertised as a reform plan that brought together choice, diversity, and school improvement.

The typical goals of choice plans start by allowing families to express their preferences over schooling options. In a comprehensive choice system, families can apply to any school in the district and do not have a default school. In more limited choice systems, families have a default school and can opt out through a choice application. District objectives in student placement are numerous. Some districts guarantee students' transportation to schools and, as a result, wish to minimize the costs of busing by ensuring that students do not travel too far from their homes. Moreover, a district may want to maintain neighborhood cohesion, allowing any children from a given neighborhood to attend the same school. Another common objective involves allowing children from the same family—siblings—to attend the same school. Finally, many districts desire balance across racial, socioeconomic, and ability dimensions across their schools. Willie argues that racially diverse schools have a positive impact on student achievement, and others argue for the achievement benefits of socioeconomic balance (Kahlenberg 2001).²

School choice advocates argue that choice is a way to inject competition from the marketplace into the regulated public school sector. Demand-side pressure from families would generate competitive pressure to improve schools. A key condition for regulated competition to realize market-based improvement is flexibility in supply-side responses. In a choice plan, administrators would have information on what schools were preferred and could make programming decisions based on this information. This in turn leads to better matches between students and schools and incentives for schools to improve to attract students.

With so many competing objectives, it is not surprising that some of the goals of choice plans came into conflict, and existing plans reflected compromises. In the 1980s, the Cambridge plan, for instance, evolved into a system in which families register by choosing up to four schools. The district had a computer system that tried to assign student to their top choices, up to school capacity and making sure not to violate the district's goals on racial and ethnic composition. In some years, the district made slight changes to guidelines on balance, and in other years, with the opening and closing of school options, plans were modified to allow entry at earlier grades or changes in neighborhood boundaries. The current Cambridge plan now allows families to rank 3 out of 12 choices and uses sibling information, residential location, and income to guide assignments. Many plans evolved in a similar manner, tweaking initial designs.

There have been two major developments related to school choice policies in the past decade. First, by the early 2000s, many districts came out of court-ordered desegregation plans. Districts such as Chicago Public Schools and San Francisco Unified Public School District wished to keep choice options given that parts of the infrastructure had developed under desegregation, and parents had some experience expressing choices. This development led to the creation of choice plans with different features reflecting the historical legacies of desegregation in particular cities.

²There is some empirical literature studying these types of effects (see, e.g., Angrist & Lang 2004, Card & Rothstein 2007, Hanushek et al. 2009).

The second major development has been a change in the legal status of race as a factor in school assignment. Beginning with a highly publicized case involving racial preferences at Boston Latin School, *McLaughlin v. Boston Sch. Committee* (1996), the Boston School Committee dropped race as a factor in their choice plan in 1999. Cambridge followed suit in 2000 and replaced race with an income-based criteria. A few years later the U.S. Supreme Court broached the subject of racial preferences in two cases, *Parents Involved in Community Schools Inc. v. Seattle School District* and *Meredith v. Jefferson County (Ky.) Board of Education*. In 2007, the Court decided that the Seattle and Louisville plans were unconstitutional because of the way they used race-conscious criteria to achieve diversity. In a five-to-four decision, Chief Justice Roberts famously wrote that “the way to stop discrimination on the basis of race is to stop discrimination on the basis of race.” The dissenting argument claimed that the decision would strip local communities of the tools they need and have used to prevent resegregation of public schools. This ruling left a number of districts to modify their plans without using race as a factor. Districts adapted by using socioeconomic criteria for student placement, redrawing attendance zones, or selecting sites for new schools. Throughout, the design of choice plans involved a compromise between the competing objectives of giving choice, having a fair procedure, and ensuring that the demographic composition of schools is not too far out of balance.

2.2. The Canonical School Choice Model

The school choice model consists of I students and N schools. There are three main features: (a) preferences of students $P = (P_1, \dots, P_I)$, (b) a vector of school capacities $q = (q_1, \dots, q_n)$, and (c) school priorities $\pi = (\pi_1, \dots, \pi_n)$.

The student preferences express a strict rank ordering over schools; this ranking need not be complete. Denote the weak ordering for student i by R_i . The school capacities express the number of seats available at each school. The school priorities encode information on how applicants are ordered, or prioritized, at schools. The school choice problem is sometimes denoted by the pair (P, π) . I call this model the canonical model because it was the first model proposed by Abdulkadiroğlu & Sönmez (2003), and many subsequent developments have involved enriching it in various ways.

Problems of assignment are often categorized into two classes: one-sided and two-sided. In one-sided problems, there is a set of agents and objects. The agents have preferences over the objects and may also have existing priorities, or claims, over the objects. The normative properties of the allocation are evaluated from only their viewpoint. In two-sided problems, in contrast, both sides of the market express preferences over each other. As a result, evaluation of the properties of the allocation may depend on preferences from both sides of the market. The school choice model falls in between these two extremes.

In many U.S. school choice plans, as in one-sided problems, schools do not express preferences over students. Rather, district administrators prioritize applications at schools using some exogenous criteria. One such criteria is neighborhood or walk-zone priority. In Boston’s school choice plan, for instance, elementary school applicants obtain walk-zone priority if they reside within one mile of the school. In other districts, schools construct an ordering of students, as in two-sided problems. In Chicago, for instance, students applying for admissions to selective high schools take an admissions test. The nine schools then

order students by their test score. Schools also evaluate applicants using criteria other than test scores to determine their strict rank ordering over applicants. In the NYC high school admissions process, some schools use seventh grade attendance and grades together with interviews at schools to determine their ordering. In some choice plans, there are schools that use exogenous criteria and schools that actively rank applicants. High school admissions in NYC are a prominent example of this hybrid case.

The outcome of a school choice problem is a student assignment, or matching $\mu: I \rightarrow S$, where $\mu(i)$ indicates the school assignment of student i . There are two properties of assignments that feature centrally in the student assignment literature. A matching μ is Pareto efficient if there is no way to improve the allocation of a student without making another student worse off. It is important to note that this definition does not take the school's perspective into account in the welfare judgment.

A matching μ is stable if there is no student-school pair (i, s) such that (a) student i prefers school s to her assignment $\mu(i)$, and (b) there is another student j with lower priority than student i assigned to s under μ . This pair is called a blocking pair. In the canonical model, this concept is sometimes referred to as the elimination of justified envy rather than stability. The reason is that the canonical model is phrased as a one-sided problem, whereas the traditional interpretation of stability is based on the strategic interpretation related to the possibility of recontracting among matched pairs as in a two-sided problem. Under the one-sided interpretation, stability embodies a notion of fairness: A student should not envy another school over his assignment and have a higher claim to that school. For simplicity, I use the term stability keeping in mind these two potential interpretations. A matching μ is student-optimal if it is stable and no other stable matching is better for some students, and no worse for all students.

A mechanism φ is a systematic procedure to construct a matching for each school choice problem. That is, it is a function that maps each school choice problem (P, π) to a matching. Let $\varphi(P, \pi)$ denote the matching produced by mechanism φ for problem (P, π) . Let $\varphi(P, \pi)(i)$ denote the assignment of student i in this matching.

A mechanism φ is strategy-proof if truth-telling is a dominant strategy for all students. That is, regardless the report of the other students, a student can do no better than reporting her preference. More formally, for all players i , for all Q_{-i} (arbitrary reports of students other than i), for all \hat{P}_i (arbitrary report of player i),

$$\underbrace{\varphi((P_i, Q_{-i}), \pi)(i)}_{\text{assignment of } i, \text{ when report truth}} \quad \underbrace{R_i}_{\text{is at least as preferred as}} \quad \underbrace{\varphi((\hat{P}_i, Q_{-i}), \pi)(i)}_{\text{assignment of } i, \text{ for arbitrary report}} .$$

Strategy-proofness is a strong requirement because it simplifies the preference submission problem of participants to one in which their best possible response does not depend on the reports of others.

2.3. Mechanisms

Three mechanisms have been closely studied for the school choice problem. The first is a mechanism based on the student-proposing deferred acceptance algorithm of Gale & Shapley (1962). For (P, π) , the mechanism works as follows. In step one, each student proposes to his first choice. Each school tentatively assigns its seats to its

proposers one at a time following their priority order. Any remaining proposers are rejected.

In general, in step k , each student who was rejected in the previous step proposes to her next choice. Each school considers the students it has been holding together with its new proposers and tentatively assigns its seats to these students one at a time following their priority order. Any remaining proposers are rejected.

The algorithm terminates either when there are no new proposals or when all rejected students have exhausted their preference lists. Gale & Shapley show that a mechanism based on this algorithm produces the student-optimal stable matching. Dubins & Freedman (1981) and Roth (1982) show that truth-telling is a dominant strategy for students.

The next mechanism defined by Abdulkadiroğlu & Sönmez (2003) is an adaptation of Gale's top trading cycles (TTC) described in Shapley & Scarf (1974). To begin, assign a counter for each school that keeps track of the number of seats still available at the school. Initially set the counters equal to the capacities of the schools. Given the counters and (P, π) , the mechanism works as follows. In step one, each student points to his favorite school. Each school points to the student who has the highest priority. There is at least one cycle. Every student can only be part of one cycle. Assign every student in a cycle to the school she points to, and remove the student. The counter of each school in a cycle is reduced by one, and if it is zero, remove the school.

In general, in step k , each remaining student points to his favorite school among the remaining schools, and each remaining school points to the student with the highest priority. There is at least one cycle. Every student in a cycle is assigned the school she points to, and the student is removed. The counter of each school in a cycle is reduced by one, and if it is zero, remove the school.

The procedure terminates when either all students are assigned a school or unassigned students have exhausted their preference lists. In the original Shapley & Scarf version of TTC, agents are endowed with objects, but many variations of TTC are possible (see, e.g., Abdulkadiroğlu & Sönmez 1999, Papai 2000, Pycia & Unver 2009). In this adaptation with counters, the priorities of students are traded among themselves starting with the highest-priority students. The mechanism is strategy-proof as a direct mechanism (Abdulkadiroğlu & Sönmez 2003, Roth & Postlewaite 1977). It also produces an assignment that is Pareto efficient.

The third mechanism is the Boston mechanism, named after the system used in Boston until 2005 by Abdulkadiroğlu & Sönmez (2003). The mechanism works as follows. In step one, only the first choices of the students are considered. For each school, consider the students who have listed it as their first choice and assign seats of the school to these students one at a time following their priority order until either there are no seats left or there is no student left who has listed it as her first choice.

In general, in step k , consider the remaining students. Only the k -th choices of these students are considered. For each school with still-available seats, consider the students who have listed it as their k -th choice and assign the remaining seats to these students one at a time following their priority order until either there are no seats left or there is no student left who has listed it as his k -th choice.

Variations of this mechanism are common in many other school districts. This mechanism has the drawback that it is not strategy-proof for students, as we illustrate in the following example.

Example: Consider a problem with three students $i_1, i_2,$ and i_3 and three schools $s_1, s_2,$ and $s_3,$ each with one seat. Student preferences, $P,$ are

$$\begin{aligned} i_1 &: s_2 - s_1 - s_3, \\ i_2 &: s_1 - s_2 - s_3, \\ i_3 &: s_1 - s_2 - s_3, \end{aligned}$$

and priorities, $\pi,$ are

$$\begin{aligned} s_1 &: i_1 - i_3 - i_2, \\ s_2 &: i_2 - i_1 - i_3, \\ s_3 &: i_3 - i_1 - i_2. \end{aligned}$$

Under the student-proposing deferred acceptance mechanism, the matching produced is

$$\mu_{\text{DA}} = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix}.$$

In this matching, none of the students obtain their top choice. The matching is not Pareto efficient, but because there are no blocking pairs, it is stable.

Under the TTC mechanism, the matching produced is

$$\mu_{\text{TTC}} = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_3 \end{pmatrix}.$$

This matching is Pareto efficient, and both student i_1 and i_2 obtain their top choice. However, student i_3 and school s_1 form a blocking pair, so the matching is not stable.

Under the Boston mechanism, the matching produced is

$$\mu_{\text{BOS}} = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_3 & s_1 \end{pmatrix}.$$

This matching is Pareto efficient, but student i_2 and school s_2 form a blocking pair, so the matching is not stable. Moreover, had student i_2 reported that s_2 was her top choice, she would have received an assignment there, which demonstrates that the mechanism is not strategy-proof.

Although the first two mechanisms are strategy-proof, the third mechanism is not. This raises the following question: Are the student-proposing deferred acceptance mechanism and TTC mechanism the only two strategy-proof mechanisms for the school choice problem? Another important mechanism, a serial dictatorship, is also strategy-proof for this problem. This mechanism places the students into a queue and then processes students in order of the queue. If the ordering of students is drawn at random, then the mechanism is called a random serial dictatorship. The first student obtains his top choice, the second student obtains her top choice among schools with available seats, and so on. This mechanism is Pareto efficient and strategy-proof but does not consider the priorities in any natural way. It is important to note, however, that there are no rigorous criteria for which a serial dictatorship involves more instances of creating blocking pairs than the TTC mechanism, although recent work by Abdulkadiroğlu & Che (2010) provides a particular characterization of TTC that is relevant for the school choice problem.³ Their

³Earlier characterizations of TTC have focused on environments in which agents are endowed with objects (see Ma 1994, Sönmez 1999).

characterization can be interpreted as showing the particular way in which TTC respects the student who has the highest priority for a school: If she is not assigned to that school, then she is assigned somewhere she prefers at least as much. Although there are some characterizations of the class of efficient and strategy-proof mechanisms together with additional axioms (e.g., Papai 2000, Pycia & Unver 2009), a characterization of all strategy-proof mechanisms remains elusive.

Another issue highlighted by these examples is that the student-proposing deferred acceptance mechanism is stable, but not efficient, whereas the TTC mechanism is efficient, but not stable. It is natural to ask for a weaker requirement: Is there an efficient and strategy-proof mechanism that also produces a stable outcome whenever it exists? Kesten (2010) shows that this is not possible, highlighting a general tension between Pareto efficiency and stability. He advocates for an efficiency-adjusted deferred acceptance mechanism. Another question is under what conditions are stability and efficiency compatible? Ergin (2002) provides a set of necessary and sufficient conditions on priority structures for which an efficient mechanism is also stable.

2.4. Important Assumptions

The canonical school choice model makes a number of important assumptions that are worth highlighting. The student preferences, which are taken as given, are hedonic: Students only care about the school they are assigned independent of the other students who are assigned there. This rules out forms of peer effects or consumption externalities in preferences, as in the case in which groups of students all wish to attend the same school only when each member of the group attends. In practice, preferences may depend on a student's distance to the school, a student's own academic and demographic characteristics, and various aspects of school quality. Some aspects of school quality may depend on the realized assignment, such as the incoming grade's peer group. However, many aspects of school quality may be more certain at the time of application, such as expenditure per student, building facilities, course offerings, and the composition of students in higher grades.

The school priorities are expressed in terms of strict orderings, involving pairwise comparisons of individual students. In practice, many districts have coarser criteria that are not strict orderings. For instance, students with siblings at the school obtain a higher priority than students who do not, but among students with siblings, all students are given equal priority. Expressing priorities in terms of pairwise comparisons between students also rules out forms of complementarities for schools. However, the model does allow for certain types of complementarities through suitable definitions of what constitutes a school. Moreover, it is possible to enrich school preferences to a larger class than simple pairwise comparisons to include substitutable preferences (Roth & Sotomayor 1990). For example, in 2010, priorities at Chicago's Gifted and Enriched Academic Programs worked as follows: Students are required to take an admissions test and are assigned to one of four tiers based on an index of the socioeconomic of their census tract geographic location. Half the seats at a program are assigned solely based on the score. The other half of the seats are split between the four tiers. If there are not enough applicants in a given tier, the school admits students in the following order: the highest-scoring student in the lowest remaining tier who had not yet been admitted, the highest-scoring student in the second lowest remaining tier, and finally the highest-scoring student in the other remaining tier.

These preferences, although complex, can be accommodated by a suitable generalization of the canonical model. Finally, as discussed above, in the canonical model, school priorities are also taken as exogenous, while some districts have schools that actively rank students.

In the canonical model, the information submitted by students is only ordinal and does not convey information on preference intensities. This focus is sometimes defended on two grounds: (a) Mechanisms that elicit cardinal information may no longer be strategy-proof, and (b) submitting cardinal information may be difficult for participants. For instance, Bogomolnaia & Moulin (2001, p. 297) defend their focus on ordinal mechanisms by writing “it can be justified by the limited rationality of agents participating in the mechanisms. There is convincing experimental evidence that the representation of preferences over uncertain outcomes by vNM utility functions is inadequate. One interpretation of this literature is that the formulation of rational preferences over a given set of lotteries is a complex process that most agents do not engage into if they can avoid it.” Providing theoretical foundations for restricting attention to mechanisms that only elicit ordinal preferences is an open question (for interesting recent work in this direction, see Carroll 2011).

One recent development involves studying mechanisms that elicit some form of cardinal information (see, e.g., Abdulkadiroğlu et al. 2009a). Another feature of the information that participants can convey in the canonical model is that it is not constrained in any way. This assumption contrasts with current practice in some districts, in which there are constraints on the number of choices that can be submitted.

Finally, the efficiency notions introduced for the canonical model utilize only the ordinal information of students. That is, the objectives of the planner are only implicitly linked to productive dimensions of the assignment, such as whether students benefit from attending the school. For instance, depending on the nature of peer effects, it may be better to group students of the same ability together, but this may conflict with student preferences. Dufló et al. (2011) argue that tracking students based on ability can generate test score gains based on evidence from a field experiment in Kenya. The implications of peer effects on school choice are undeniably important but are outside the scope of this survey. Models incorporating these features will likely require the development of frameworks that impose more structure on preferences and the nature of education production as in Epple & Romano (1998).

3. COARSE SCHOOL PRIORITIES AND EFFICIENCY

One of the first ways the basic model has been enriched is by examining the implications of coarse school priorities. This issue obtained attention during the design of the NYC high school assignment process. In NYC, there are over 600 high school programs, and eighth and ninth grade students can apply to any program in the city. There are two main types of high school programs in NYC. The first are those who express a rank ordering over applicants as in screened or audition schools. The second are schools that have fixed criteria to order students, such as limited unscreened schools that give first priority to students who attend information fairs or live in various parts of the district. Abdulkadiroğlu et al. (2005, 2009b) present more institutional details about schooling options in NYC.

During the course of the design of the new mechanism, policy makers agreed to two phases to assign schools: the main round, which involved both types of schools and was to be based on student-proposing deferred acceptance, and the supplementary round, involving students who were unassigned in the main round and remaining school capacities, for which school orderings of students do not play a role.⁴ In both rounds, students would be allowed to rank up to 12 school choices.

One practical issue was how coarse priorities should be turned into strict priorities at schools. This issue was relevant for both rounds. During the course of the policy discussion, an official remarked,

I believe that the equitable approach is for a child to have a new chance with each . . . program. . . . The fact is that each child had a chance. If we use only one random number, and I had the bad luck to be the last student in the line this would be repeated 12 times and I would never get a chance. I do not know how we could explain this to a parent.

This policy discussion motivated a reconsideration of the assignment mechanisms in the presence of coarse priorities. To illustrate the issue, consider the earlier example, but now suppose that schools s_1 and s_2 are indifferent between all applicants. That is, the priorities, π , are

$$\begin{aligned} s_1 &: \{i_1, i_2, i_3\}, \\ s_2 &: \{i_1, i_2, i_3\}, \\ s_3 &: i_3 - i_1 - i_2. \end{aligned}$$

If the mechanism based on student-proposing deferred acceptance uses lotteries to convert these indifferences into strict orderings, and the resulting orderings are as in the earlier example, then both students i_1 and i_2 are assigned to their second choice, when they would be better off trading their placements with one another. If the priorities at s_1 had been strict, then one might justify preventing this trade because student i_3 forms a blocking pair with school s_1 after the trade. However, this is not a blocking pair when s_1 is indifferent between applicants. Hence the student-proposing deferred acceptance mechanism does not always produce a student-optimal stable matching and allows for efficiency loss.

Given that tiebreaking may have welfare consequences, one question raised by the quotation above is whether it is better for students to have school-specific lotteries or a single lottery draw. Abdulkadiroğlu et al. (2009b) show that any student-optimal stable matching can be produced by a single lottery draw so that school-specific lotteries only add matchings that are not student optimal relative to a single lottery draw. These statements are from an ex post perspective, and there is currently no known stronger ex ante argument for single versus multiple tiebreaking based on the distribution of matchings.

These facts suggest a number of additional questions. First, with coarse orderings at schools, is there a strategy-proof mechanism that produces a student-optimal matching? Erdil & Ergin (2008) show that no such mechanism exists. Second, is it possible to construct mechanisms that are student optimal? Erdil & Ergin (2008) advocate one proposal, stable improvement cycles, which finds a student-optimal stable matching in polynomial time. Next, is it possible to recover some of the efficiency loss of the single tiebreaker version

⁴Specialized high schools, such as Stuyvesant High School and Bronx High School of Science, are assigned through in an earlier round based on a special admissions test.

of the student-proposing deferred acceptance mechanism? Abdulkadiroğlu et al. (2009b) consider this question and show that this mechanism is on the efficient frontier. That is, suppose there exists a mechanism $\tilde{\varphi}$, which produces a Pareto-dominating matching:

$$\mu = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_3 \end{pmatrix}.$$

To show that this mechanism is not strategy-proof, consider the economy in which student i_1 only prefers school s_2 and denote the new preference profile by Q . In the problem (Q, π) , the student-proposing deferred acceptance mechanism produces the matching

$$v = \begin{pmatrix} i_1 & i_2 & i_3 \\ i_1 & s_2 & s_1 \end{pmatrix},$$

where student i_1 is unassigned. If mechanism $\tilde{\varphi}$ dominates the student-proposing deferred acceptance mechanism, it must also yield matching v . But in the school choice problem (Q, π) , student i_1 could manipulate $\tilde{\varphi}$ by submitting the elongated preference list $s_1 - s_1 - s_3$. This shows that $\tilde{\varphi}$ is not strategy-proof.

That no strategy-proof mechanism Pareto dominates the student-proposing deferred acceptance mechanism places it on the efficient frontier of strategy-proof mechanisms. The lesson that emerges is that incentives together with stability must necessarily entail efficiency loss.

Motivated by these theoretical results, Abdulkadiroğlu et al. (2009b) compare the extent of efficiency loss using data from the field from the new systems in Boston and NYC. Table 1 reports the average number of students obtaining a choice on their rank order list averaged over 250 draws of the random tiebreaker. DA-STB is the outcome of the student-proposing deferred acceptance mechanism with a single tiebreaker, and DA-MTB is the outcome with school-specific tiebreaking. SOSM is a student-optimal stable matching computed by applying the procedure of Erdil & Ergin (2008) with a cycle selection rule and with initial matching from DA-STB. When there are indifferences, there may be multiple stable matchings, so in the table we select a particular student-optimal stable matching that Pareto dominates the matching produced by DA-STB.

The first comparison is between DA-STB and DA-MTB. Even though both matchings are stable, DA-STB has more students obtaining their top choice, although the magnitude of the difference is larger in NYC, where approximately 2,000 more students obtain their top choice under a single lottery draw. Interestingly, fewer students are unassigned under DA-MTB than DA-STB, however. The second comparison between DA-STB and SOSM provides one measure of the efficiency loss due to the presence of indifferences. In NYC, approximately 1,500 students could receive a better high school choice in a student-optimal stable matching relative to the outcome produced by the single tiebreaking version of the student-proposing deferred acceptance mechanism. In Boston, alternatively, less than seven students on average obtain an improved assignment in the student-optimal matching. Abdulkadiroğlu et al. (2009b) conjecture that the main difference is that the pattern of preferences in Boston is different than in NYC, due in large part to different geographic and transportation situations and to the fact that, in Boston, the preferences are for younger children. But these empirical results raise the need for quantitative results in matching theory that provide guidance on what features of the student preferences and

Table 1 Impact of tiebreaking in New York City and Boston^a

Choice	New York City			Boston		
	DA-STB (1)	DA-MTB (2)	SOSM (3)	DA-STB (4)	DA-MTB (5)	SOSM (6)
1	32,105.3	29,849.9	32,701.5	2,251.8	2,157.3	2,256.6
2	14,296.0	14,562.3	14,382.6	309.8	355.5	307.4
3	9,279.4	9,859.7	9,208.6	154.9	189.3	154.0
4	6,112.8	6,653.3	5,999.8	59.7	76.1	58.7
5	3,988.2	4,386.8	3,883.4	27.4	34.1	27.0
6	2,628.8	2,910.1	2,519.5	4.9	6.0	4.9
7	1,732.7	1,919.1	1,654.6	2.6	2.8	2.5
8	1,099.1	1,212.2	1,034.8	1.9	0.9	1.9
9	761.9	817.1	716.7	1.2	0.4	1.2
10	546.4	548.4	485.6	0.3	0.1	0.3
11	348.0	353.2	316.3	—	—	—
12	236.0	229.3	211.2	—	—	—
Unassigned	5,613.4	5,426.7	5,613.4	112.4	104.6	112.4

^aThis table is based on data from the main round of the New York City high school admissions process in 2006–2007 for students requesting an assignment for grade 9 and from the Boston elementary school (grade K2) admissions process in 2006–2007. The table reports the number from 250 draws of a random tiebreaker. Abbreviations: DA-MTB, deferred-acceptance mechanism with school-specific tiebreaker; DA-STB, deferred-acceptance mechanism with a single tiebreaker; SOSM, student-optimal stable matching. Table taken from Abdulkadiroğlu et al. (2009b).

school priorities are responsible for these differences.⁵ They also raise the question of what type of behavior is expected in mechanisms that might improve on the student-proposing deferred acceptance mechanism.

The discussion on how to convert coarse priorities into strict priorities is also relevant for other mechanisms. In particular, in the supplementary round in NYC, schools are indifferent between applicants, so one approach might be to conduct school-specific lotteries and then apply the TTC mechanism. Pathak & Sethuraman (2011) show that this mechanism is equivalent to a random serial dictatorship. Another mechanism might be to randomly endow each student with a school seat and then let them trade. Abdulkadiroğlu & Sönmez (1998) show that this mechanism is equivalent to a random serial dictatorship. Hence three mechanisms are the same for the special case in which each school is indifferent between applicants.

Given that two alternative mechanisms are equivalent to a random serial dictatorship, there has been a renewed interest in understanding the efficiency properties of this mechanism. Bogomolnaia & Moulin (2001) point out that a random serial dictatorship may produce a matching that is not ordinally efficient. It may be possible to find a random

⁵Erdil & Ehlers (2010) provide conditions on priorities under which the constrained efficient rule is efficient.

assignment, which stochastically dominates the random assignment produced by a random serial dictatorship. That is, for each student i , the probability of receiving the k -th choice is at least as high under an alternative mechanism than under the random serial dictatorship for all k choices. They develop and analyze the probabilistic serial mechanism that produces an ordinally efficient assignment but is not strategy-proof. Using field data from NYC's supplementary round, Pathak (2007) compares the empirical performance of the probabilistic serial mechanism to a random serial dictatorship. The difference between the two mechanisms is relatively small; out of approximately 8,000 students, just over 15 more students received their top choice under the probabilistic serial mechanism and approximately 50 more students received a more preferred assignment.

This finding was one motivation for Che & Kojima (2010), who provide conditions under which the random serial dictatorship and probabilistic serial mechanism are asymptotically equivalent. Their result implies conditions under which the inefficiency of random serial dictatorship becomes small in large allocation problems. A nice feature of this result is that it is quantitative: Rather than illustrating the existence of ordinal inefficiency, they show what conditions ensure that it is quantitatively small. Another related paper is by Kesten (2009), who shows the equivalence of the random serial dictatorship and probabilistic serial mechanisms under a different set of assumptions. Manea (2009) considers a different asymptotic notion: In his model, preferences are randomly generated, and the object of interest is the likelihood that the assignment from a random serial dictatorship is ordinally inefficient. He shows that a random serial dictatorship is highly likely to produce an ordinally inefficient allocation. The reason for the apparently different result is it is only about the existence of ordinal inefficiency and not the extent of efficiency loss, as in Che & Kojima (2010).

4. MECHANISMS AND MARKET SIZE

The empirical study on NYC's supplementary round provided motivation for subsequent theoretical developments. In a similar vein, empirical and simulation evidence on the performance of two-sided matching models when there are a large number of participants played a key role in suggesting theoretical work on these topics.

In the main round in NYC, about half of the school districts submit rankings over applicants. In such a setting, there is no strategy-proof mechanism for both students and schools (Roth 1982). Returning to our example, suppose now that the schools order students in the following way:

$$\begin{aligned} s_1 &: i_1 - i_2 - i_3, \\ s_2 &: i_2 - i_1 - i_3, \\ s_3 &: i_3 - i_1 - i_2, \end{aligned}$$

but now school s_1 is one that ranks applicants, so its ordering is not from exogenous priorities. Under the student-proposing deferred acceptance mechanism, the resulting matching is

$$v = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_1 \end{pmatrix},$$

and school s_1 is assigned its second-ranked student. If, instead, school s_1 declared that student i_2 is not acceptable and only ranked i_1 , the resulting matching is

$$v' = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix}.$$

School s_1 's strategic rejection of student i_2 results in its obtaining its top choice; this shows that schools can manipulate the student-proposing deferred acceptance mechanism.

The presence of schools that actively rank applicants in NYC makes the canonical school choice model closer to two-sided matching market models (surveyed in Roth & Sotomayor 1990). In the labor market context, Roth & Peranson (1999) conduct a series of simulations on data from the NRMP and on randomly generated data. In their simulations, few agents could have benefitted by submitting false preference lists or by manipulating capacity in large markets given the reports of other agents. These simulations lead them to conjecture that the fraction of participants with preference lists of limited length who can manipulate tends to zero as the size of the market grows.⁶

The first theoretical attempt to understand these findings is by Immorlica & Mahdian (2005), who focus on one-to-one matching models. This paper is particularly innovative because a number of subsequent papers have built on and extended some of its analytical tools. Kojima & Pathak (2009) consider many-to-one matching markets with the student-proposing deferred acceptance mechanism, in which schools have arbitrary preferences such that every student is acceptable, and students have random preferences of fixed length drawn iteratively from an arbitrary distribution. They show that the expected proportion of schools that have incentives to manipulate the mechanism when every other school is truth-telling converges to zero as the number of schools approaches infinity. The key step in the argument involves showing that, when there are a large number of schools, the chain reaction caused by a school's strategic rejection of a student is unlikely to make a more preferred student apply to that school. Loosely speaking, this means in the example it is unlikely that school s_1 's strategic rejection leads student i_1 to apply to that school in the course of the student-proposing deferred acceptance mechanism under the large-market assumptions.

Roth & Peranson's (1999) simulations hold fixed the behavior of all other participants and consider deviations by particular agents one by one. As such, they are not necessarily about equilibrium implications. This consideration is where the theory pushes the envelope one step further. Kojima & Pathak (2009) conduct equilibrium analysis in the large market. With an additional condition, called sufficient thickness, they show that truthful reporting is an approximate equilibrium in a large market that is sufficiently thick.

The feedback from the field, for which, on the surface, the potential for manipulation did not appear to undermine systems based on the student-proposing deferred acceptance mechanism, on the empirical and simulation evidence left a challenge for theory: What conditions would ensure that manipulations are unlikely? Kojima & Pathak (2009) take a step toward understanding these results, although it is not yet known whether it is possible to tighten the rates of convergence. Nonetheless, this case also illustrates the potential for the development of quantitative aspects of matching market design, in which empirical and simulation evidence feeds into theoretical developments.

⁶Roth & Peranson (1999) also investigate the complications that couples create for two-sided matching markets. Kojima et al. (2010) build on large-market techniques to the student existence of stable matchings and incentives of matching mechanisms in the presence of couples.

5. CONFRONTING MECHANISMS WITH THE REAL WORLD

5.1. Levels of Sophistication

One important issue that emerges when designing and implementing actual mechanisms is that participants may not behave according to the theoretical assumptions of the models. Roth & Ockenfels (2002) summarize these considerations nicely in their description of online auctions: “In designing new markets, we need to consider not only the equilibrium behavior that we might expect experienced and sophisticated players to eventually exhibit, but also how the design will influence the behavior of inexperienced participants, and the interaction between sophisticated and unsophisticated players.”

One challenge with this statement is having a reasonable way to model unsophisticated players. In student assignment problems, a natural approach is to assume unsophisticated players simply report the truth even when it may not be in their best interest to do so. In studying the equilibrium properties of the Boston mechanism, this approach finds support based on data from laboratory experiments. For example, Chen & Sönmez (2006) show that approximately 20% of subjects in a laboratory experiment report the truth under the Boston mechanism. Recall that the major difficulty with the Boston mechanism is that participants may benefit by submitting a rank order list that is different from their true underlying preferences over schools. Loosely speaking, the Boston mechanism attempts to assign as many students as possible to their first-choice school, and only after all such assignments have been made does it consider assignments of students to their second choices, and so on. If a student is not admitted to his first choice school, his second choice may be filled with students who have listed it as their first choice. That is, a student may fail to get a place in his second-choice school that would have been available had he listed that school as his first choice. If a student is willing to take a risk with his first choice, then he should be careful to rank a second choice that he has a chance of obtaining.

If a mechanism is not strategy-proof, a natural direction is to analyze its equilibrium properties. Assuming all players are sophisticated, Ergin & Sönmez (2006) characterize the set of Nash equilibria of the preference revelation game induced by the Boston mechanism under complete information and strict priorities. Consider the previous example in which student preferences, P , are

$$\begin{aligned} i_1 &: s_2 - s_1 - s_3, \\ i_2 &: s_1 - s_2 - s_3, \\ i_3 &: s_1 - s_2 - s_3, \end{aligned}$$

and priorities, π , are

$$\begin{aligned} s_1 &: i_1 - i_3 - i_2, \\ s_2 &: i_2 - i_1 - i_3, \\ s_3 &: i_3 - i_1 - i_2. \end{aligned}$$

It is possible to construct a Nash equilibrium where student i_1 reports s_1 , student i_2 reports s_2 , and student i_3 reports s_3 as top choices. The resulting matching is

$$\mu^{\text{NE}} = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix},$$

which is the student-optimal stable matching. For this problem, it is the only Nash equilibrium outcome, but more generally Ergin & Sönmez (2006) show that Nash equilibrium

outcomes of the Boston game are equivalent to the set of stable matchings. This result implies that the best possible equilibrium outcome under the Boston mechanism is equal to the student-optimal stable matching, an outcome that can be attained via a strategy-proof mechanism. Moreover, players need to have a high degree of coordination to obtain this outcome.

It is important to recognize the strong assumptions underlying this analysis: Players have complete information about the rank order lists and priorities of each other, and all players can compute their optimal strategies in the Boston mechanism. There are some sophisticated families who understand the strategic features of the Boston mechanism and have developed rules of thumb for how to submit preferences strategically. For instance, the West Zone Parents Group, a well-informed group of approximately 180 members who meet regularly prior to admissions time to discuss Boston school choice for elementary school (grade K2), recommends two types of strategies to its members. Their introductory meeting minutes on October 28, 2003, state, "One school choice strategy is to find a school you like that is undersubscribed and put it as a top choice, OR, find a school that you like that is popular and put it as a first choice and find a school that is less popular for a 'safe' second choice." This quotation only indicates some sort of strategic sophistication. It would be interesting to understand what types of evolutionary or learning rules would support the predictions of Nash equilibrium behavior in this setting. Alternatively, one could consider alternative equilibrium notions such as self-confirming equilibrium to model this situation.

During the 2005 policy discussion about abandoning the mechanism in Boston, policy makers focused on how, under a strategy-proof mechanism, if families have access to advice on how to strategically modify their rank order lists from groups like the West Zone Parents Group or through family resource centers, they can do no better than by submitting their true preferences. Superintendent Payzant's recommendation to change the mechanism emphasized this feature, and the Boston Public Schools Strategic Planning Team, in their recommendation to implement a new Boston Public School assignment algorithm, dated May 11, 2005, states, "A strategy-proof algorithm 'levels the playing field' by diminishing the harm done to parents who do not strategize or do not strategize well."

The model in Pathak & Sönmez (2008) has both sincere families who report the truth and sophisticated families who best respond to the preference revelation game induced by the Boston mechanism. We characterize the Nash equilibria of this game and compare the equilibrium outcomes with the dominant-strategy outcome of the student-proposing deferred acceptance mechanism.

There are two main results. The first is a characterization of the equilibrium outcomes of the Boston game as the set of stable matchings of a modified problem in which sincere students lose their priorities to sophisticated students. This result implies that there exists a Nash equilibrium outcome in which each student weakly prefers her assignment to any other equilibrium assignment. Hence the Boston game is a coordination game among sophisticated students.

Returning to our main example, suppose that student i_2 is sincere and hence reports $s_1 - s_2 - s_3$ in the preference revelation game induced by the Boston mechanism. Our characterization implies that i_2 loses priority at each school other than his top choice, and

the Nash equilibrium outcome is simply the set of stable matching with the following priorities:

$$\begin{aligned}\pi_{s_1} &: i_1 - i_3 - i_2, \\ \tilde{\pi}_{s_2} &: i_1 - i_3 - i_2, \\ \tilde{\pi}_{s_3} &: i_3 - i_1 - i_2,\end{aligned}$$

where i_2 is ordered last at school s_2 and s_3 . Because the set of Nash equilibrium outcomes of the Boston game is equal to the set of stable matchings of this modified economy, the Nash equilibrium outcome for the example is

$$\mu^{\text{NE}} = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_3 & s_1 \end{pmatrix}.$$

Sincere student i_2 obtains his last choice, under the Boston mechanism, when previously he obtained her second choice as the Nash outcome, when he was sophisticated.

Next we compare the equilibria of the Boston game to the dominant-strategy outcome of the student-proposing deferred acceptance mechanism. We show that any sophisticated student weakly prefers her assignment under the Pareto-dominant Nash equilibrium outcome of the Boston game over the dominant-strategy outcome of the student-proposing deferred acceptance mechanism. In the example, student i_1 is assigned to s_1 (her second choice) and student i_3 is assigned to s_3 (his third choice), whereas under the Nash equilibrium of the Boston game, both receive their top choice. When only some of the students are sophisticated, the Boston mechanism gives a clear advantage to sophisticated students provided that they can coordinate their strategies at a favorable equilibrium.

In our example, the Boston game has a unique equilibrium, but in general, there may be many stable matchings, and hence the equilibrium may no longer be unique. There is, however, evidence in the literature that suggests that the size of the set of stable matchings may be very small in real-life applications of matching models. Using data for the years 1991–1994 and 1996 for the thoracic surgery market, Roth & Peranson (1999) show that there are two stable matchings each for 1992 and 1993, and one stable matching each for 1991, 1994, and 1996. One caveat of these computational experiments is that the thoracic surgery market used the hospital-optimal stable mechanism in these years, and truth-telling is not a dominant strategy for interns or for hospitals under this mechanism. So it is theoretically possible that the small number of stable matchings is an implication of preference manipulation.

The same computational exercise is on firmer ground for the school years 2005–2006 and 2006–2007 for Boston Public Schools student admissions when a strategy-proof mechanism is used. The results of these computational experiments are similar to those of Roth & Peranson: At grade K2 for the school years 2005–2006 and 2006–2007, there is only one stable matching for either year. At grade 6, the situation is not very different. For the school year 2005–2006, there are only two stable matchings, and among more than 3,200 students, only two are affected by the choice of a stable matching. For the school year 2006–2007, there are also two stable matchings, and among more than 2,900 students only three are affected by the choice of a stable matching. The likely reason this occurs is that for most students the factors that give a student higher priority at a Boston school (i.e., proximity and the presence of a sibling) also make it more preferable for the student.

These computational experiments suggest that, although multiple equilibria are a theoretical possibility under the Boston game, it likely affects a very small minority of students because the set of Nash equilibrium outcomes is equal to the set of stable matchings of an

augmented economy in which sincere students lose priority to sophisticated students. Using data for the school years 2005–2006 and 2006–2007 and admission to grade K2 and grade 6, Pathak & Sönmez (2008) run computational experiments by randomly setting 20% of students to be sincere and the rest to be sophisticated. They calculate the student-optimal stable matching and the school-optimal stable matching for the resulting augmented economy and repeat the same exercise 1,000 times to calculate the number of students affected on average by the multiplicity of the Nash equilibria. The same experiment is repeated for the cases in which 40%, 60%, and 80% of the students are sincere, respectively. Table 2 summarizes the results.

Most of the time, the augmented economy has a unique stable matching, and more specifically no more than 0.38 students (less than 0.013% of students) are affected on average by the multiplicity of the Nash equilibria in each case. Hence, although the main result does not theoretically extend to all equilibria, the computational experiments suggest that multiplicity may not be a significant problem in our application.

What about sincere students in the new mechanism? In the example, the sincere student is better off because she receives her second choice. However, this is not a general result, as Pathak & Sönmez (2008) show. Although no sophisticated student loses priority to any other student, some of the sincere students may gain priority at a school at the expense of other sincere students by ranking the school higher on their preference list. As a result, it is possible that a sincere student might benefit from the Boston mechanism.

This model and the computational experiments enrich the discussion of the rationale for changing the Boston mechanism. Even with players with heterogeneous levels of sophistication, changing the mechanism does not unambiguously benefit sincere students. Hence, under the assumptions of the model, this policy change cannot be seen as a Pareto improvement even for this subset of players. Rather, the idea of leveling the playing field only indicates that sophisticated students lose their strategic rents under the new mechanism.

Table 2 Average number of students receiving different schools in student-optimal versus school-optimal matching^a

	Fraction of sincere students			
	20%	40%	60%	80%
2005–2006				
Grade K2	0.14	0.08	0.04	0.01
Grade 6	0.38	0.20	0.07	0.01
2006–2007				
Grade K2	0.03	0.01	0.00	0.00
Grade 6	0.24	0.14	0.05	0.01

^aThis table is based on data provided by Boston Public Schools for round one of their admissions process in 2005–2006 and 2006–2007. Table taken from Pathak & Sönmez (2008).

5.2. New Approaches to Incentive Constraints

In his Nemmers Prize lecture, Milgrom (2009) argues that practical experiences implementing auction mechanisms led him to reconsider the nature of incentive constraints for applied mechanism design problems. He argues that incentive-compatible mechanisms can have bad properties, and in his view perhaps too much emphasis has been placed on incentives as constraints in mechanism design.

For instance, Day & Milgrom (2008) consider package auctions and consider an incentive metric that is to minimize the incentives to misreport in core-selecting auctions. This notion is not based on equilibrium, but it may highlight a potentially relevant consideration for the setting of a package auction. Erdil & Klemperer (2010) argue instead that in core-selecting package auctions it may be preferable to consider a bidder's marginal incentive to deviate, rather than his maximal incentive to deviate (the best possible deviation). Their argument is that the marginal incentive is not as sensitive to other bidders' behavior and hence may be easier to calculate. This criteria leads them to advocate the minimum revenue core outcome closest to some given point that does not depend on the winners' bids, unlike the proposals that consider distance to the Vickrey-Clarke-Groves payment (which depends on winners' bids).

In student assignment problems, it is also useful to consider other ways to think about incentive constraints. Pathak & Sönmez (2011) explore a formalization of how easy a mechanism may be to manipulate or game. They compare two direct mechanisms based on the following notion: Mechanism φ is more manipulable than mechanism ψ , if whenever ψ can be manipulated, φ can also be manipulated (even though the converse does not hold). This notion allows for various formulations, depending on whether the manipulating agent is the same across the problems. Like Day & Milgrom (2008), this notion is not intended to be based on equilibrium. However, it has the benefit of an equilibrium interpretation, as described in Pathak & Sönmez (2011).

To provide one illustration of this general idea in the context of student assignment, we return to a practical issue from NYC. One feature of NYC's new mechanism is that it only allows students to submit a rank order list of their top 12 choices. Based on the strategy-proofness of the student-proposing deferred acceptance mechanism, the following advice was given to students: "You must now rank your 12 choices according to your true preferences." For a student who has more than 12 acceptable schools, truth-telling is no longer a dominant strategy. In practice, between 20%–30% of students rank 12 schools. This issue was first theoretically investigated by Haeringer & Klijn (2009). In general, there are many equilibrium outcomes, and these depend on high levels of sophistication among participants.

In contrast, it is possible to show that the greater the number of choices a student can make, the less vulnerable the constrained version of student-proposing deferred acceptance mechanism is to manipulation. More formally, let GS be the student-proposing deferred acceptance mechanism and GS^k be the constrained version of the student-proposing deferred acceptance mechanism in which only the top k choices are considered. Pathak & Sönmez (2010) show that if $\ell > k > 0$, and there are at least ℓ schools, then GS^k is more manipulable than GS^ℓ . This result provides a formal criterion to encourage a district to relax constraints on rank order lists.

Policy makers seem to dislike the idea of gaming, presumably because it is costly and some participants may be able to bear its costs more easily than others. Given the prevalence of this

sentiment, it is surprising that we have few models of how gaming or manipulation is undesirable. These intuitions may rest on procedural aspects of a mechanism, rather than the properties of the outcomes of a mechanism. Understanding these dimensions of mechanisms will provide an important bridge between mechanism design in theory and in practice.

5.3. Experiments

The other way theoretical developments have confronted the real world is through experiments. During initial meetings with the Strategic Planning Team at Boston Public Schools, school officials had studied the experiment in Chen & Sönmez (2006) closely. This experiment compares the performance of students in the Boston mechanism, the student-proposing deferred acceptance mechanism, and the TTC mechanism. One nice feature of the experiment is that it is able to induce participant preferences, so it can compare how these are related to submitted preferences. The experiment finds that there is a higher degree of preference manipulation under the Boston mechanism than under the two alternatives, and this negatively impacts efficiency.

Since this initial experiment, there has been a flurry of additional experiments. Many of these experiments are intended to fill in areas in which theory is silent or gives only weak predictions. For instance, Calsamiglia et al. (2010) investigate the performance of mechanisms in the presence of constraints on the number of schooling options one can list. They are motivated by theoretical work on this topic by Haeringer & Klijn (2009). Featherstone & Niederle (2008) investigate the role of incomplete information in the Boston mechanism. They are motivated by recent discussions highlighting how the Boston mechanism, although manipulable, may be able to elicit preference intensity (see, e.g., Abdulkadiroğlu et al. 2009a, Miralles 2008).

6. CONCLUSION

In the past decade, problems related to student assignment have invigorated theoretical research on matching and assignment models. The literature reviewed here provides examples of how field experience implementing mechanisms can motivate subsequent theoretical developments. Although impossibility results indicate that there is no student-optimal stable assignment when there are coarse priorities, field evidence suggests that this may not significantly impact student welfare in Boston but impacts thousands of students in NYC's high school choice plan. The simulations of Roth & Peranson (1999) showed that, despite impossibility results on strong incentive properties of two-sided matching mechanisms, market size may ameliorate strategic issues. Finally, field evidence on heterogeneous levels of sophistication among participants in Boston motivated the examination of models in which players have varying understanding of the choice plan. This and subsequent work illustrated the importance of the consideration of procedural aspects of student assignment mechanisms in addition to the conventional focus on the outcomes produced by mechanisms.

Roth (2002) advocates for the creation of an engineering-style branch of applied mechanism design. Three cases described here—coarse priorities, market size, and heterogeneous levels of sophistication—are areas in which theoretical developments can trace their origins to particular engineering episodes. Their existence reinforces the argument for recording and creating a literature on case studies of applied mechanism design.

Of course, many interesting questions remain. In particular, some of this work is part of an emerging quantitative theory of matching market design, which moves away from impossibility and knife-edge results. In a quantitative theory, comparative statics can inform us on how the magnitude of certain issues may depend on features of the environment and can provide guidance for these situations. Another wide-open area involves building bridges between laboratory experiments and evidence on actual play in mechanisms in the field. This work, however, is challenging as measuring true preferences in the field is considerably more difficult than in the lab. Finally, much remains to be done to examine the effects of particular student mechanisms on outcomes beyond assignment, such as student achievement, and to broaden the scope of the design objectives to include the overall organization of the educational system.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I thank Atila Abdulkadiroğlu, Fuhito Kojima, Alvin E. Roth, and Tayfun Sönmez for allowing me to draw upon our joint work and Esther Duflo for helpful comments and suggestions. Financial support was provided by the National Science Foundation.

LITERATURE CITED

- Abdulkadiroğlu A, Che Y-K. 2010. *The role of priorities in assignment indivisible objects: a characterization of top trading cycles*. Work. Pap., Columbia Univ.
- Abdulkadiroğlu A, Che Y-K, Yasuda Y. 2009a. Expanding choice in school choice. *Econ. Res. Initiat. Duke Res. Pap.* 20, Duke Univ.
- Abdulkadiroğlu A, Pathak PA, Roth AE. 2005. The New York City high school match. *Amer. Econ. Rev.* 95:364–67
- Abdulkadiroğlu A, Pathak PA, Roth AE. 2009b. Strategy-proofness versus efficiency in matching with indifference: redesigning the New York City high school match. *Am. Econ. Rev.* 99(5):1954–78
- Abdulkadiroğlu A, Sönmez T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66:649–701
- Abdulkadiroğlu A, Sönmez T. 1999. House allocation with existing tenants. *J. Econ. Theory* 88:233–60
- Abdulkadiroğlu A, Sönmez T. 2003. School choice: a mechanism design approach. *Am. Econ. Rev.* 93:729–47
- Alves M, Willie C. 1987. Controlled choice assignments: a new and more effective approach to school desegregation. *Urban Rev.* 19(2):67–88
- Angrist JD, Lang K. 2004. Does school integration generate peer effects? Evidence from Boston's Metco program. *Am. Econ. Rev.* 94(5):1613–34
- Balinski M, Sönmez T. 1999. A tale of two mechanisms: student placement. *J. Econ. Theory* 84:73–94
- Baum-Snow N, Lutz B. 2011. School desegregation, school choice and changes in residential location patterns by race. *Am. Econ. Rev.* In press
- Bogomolnaia A, Moulin H. 2001. A new solution to the random assignment problem. *J. Econ. Theory* 100:295–328

- Boustan L. 2010. *School desegregation and urban change: evidence from city boundaries*. Unpublished manuscript, Univ. Calif., Los Angeles
- Burgess S, Greaves E, Vignoles A, Wilson D. 2009. What parents want: school preferences and school choice. *Cent. Market Public Org.* 09/222, Univ. Bristol
- Calsamiglia C, Haeringer G, Klijn F. 2010. Constrained school choice: an experimental study. *Am. Econ. Rev.* 100:1860–74
- Card D, Rothstein J. 2007. Racial segregation and the black-white test score gap. *J. Public Econ.* 91:2158–84
- Carroll G. 2011. *On direct mechanisms with ordinal preferences*. Unpublished manuscript, Mass. Inst. Technol.
- Che Y-K, Kojima F. 2010. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica* 78:1625–72
- Chen Y, Sönmez T. 2006. School choice: an experimental study. *J. Econ. Theory* 127:202–31
- Chiu S, Weng W. 2009. Endogenous preferential treatment in centralized admissions. *Rand J. Econ.* 40:258–82
- Cook G. 2003. School assignment flaws detailed. *Boston Globe*, Sept. 12, Metro Desk
- Day R, Milgrom P. 2008. Core-selecting package auctions. *Int. J. Game Theory* 36(3):393–407
- Dubins LE, Freedman DA. 1981. Machiavelli and the Gale-Shapley algorithm. *Am. Math. Mon.* 88:485–94
- Duflo E, Dupas P, Kremer M. 2011. Peer effects, teacher incentives, and the impact of tracking: evidence from a randomized evaluation in Kenya. *Am. Econ. Rev.* In press
- Epple D, Romano R. 1998. Competition between private and public schools, vouchers, and peer group effects. *Am. Econ. Rev.* 88(1):33–62
- Erdil A, Ehlers L. 2010. Efficient assignment respecting priorities. *J. Econ. Theory* 145:1269–82
- Erdil A, Ergin H. 2008. What's the matter with tie-breaking? Improving efficiency in school choice. *Am. Econ. Rev.* 98:669–89
- Erdil A, Klemperer P. 2010. A new payment rule for core-selecting package auctions. *J. Eur. Econ. Assoc.* 8:2–3, 537–47
- Ergin H. 2002. Efficient resource allocation on the basis of priorities. *Econometrica* 70:2489–98
- Ergin H, Sönmez T. 2006. Games of school choice under the Boston mechanism. *J. Public Econ.* 90:215–37
- Featherstone C, Niederle M. 2008. *Ex ante efficiency in school choice mechanisms: an experimental investigation*. Unpublished manuscript, Stanford Univ.
- Fiske E. 2002. Controlled choice in Cambridge, Massachusetts. In *Divided We Fail: Coming Together Through Public School Choice*, ed. R Kahlenberg, pp. 167–208. New York: Century Found.
- Gale D, Shapley LS. 1962. College admissions and the stability of marriage. *Am. Math. Mon.* 69:9–15
- Haeringer G, Klijn F. 2009. Constrained school choice. *J. Econ. Theory* 144:1921–47
- Hanushek EA, Kain JF, Rivkin SG. 2009. New evidence about *Brown v. Board of Education*: the complex effects of school racial composition on student achievement. *J. Labor Econ.* 27(3):349–83
- Herszenhorn D. 2004. Council members see flaws in school-admissions plan. *New York Times*, Nov. 19, Metrop. Desk, p. 3
- Holme JJ, Wells AS. 2008. School choice beyond district borders: lessons for the reauthorization of NCLB from interdistrict desegregation and open enrollment plans. In *Improving on No Child Left Behind*, ed. R Kahlenberg. New York: Century Found.
- Immorlica N, Mahdian M. 2005. Marriage, honesty, and stability. *SODA '05 Proc. 16th Annu. ACM-SIAM Symp. Discrete Algorithms*, pp. 53–62. Philadelphia: SIAM
- Kahlenberg RD. 2001. *All Together Now: Creating Middle-Class Schools Through Public School Choice*. Washington, DC: Brookings Inst.
- Kesten O. 2009. Why do popular mechanisms lack efficiency in random environments? *J. Econ. Theory* 144:2209–26

- Kesten O. 2010. School choice with consent. *Q. J. Econ.* 125:1297–348
- Kojima F, Pathak PA. 2009. Incentives and stability in large two-sided matching markets. *Am. Econ. Rev.* 99:608–27
- Kojima F, Pathak PA, Roth AE. 2010. *Matching with couples: incentives and stability in large markets*. NBER Work. Pap. 16028
- Lavy V. 2010. Effects of free choice among public schools. *Rev. Econ. Stud.* 77:1164–91
- Ma J. 1994. Strategy-proofness and the strict core in a market with indivisibilities. *Int. J. Game Theory* 23:75–83
- Manea M. 2009. Asymptotic ordinal inefficiency of random serial dictatorship. *Theor. Econ.* 4(2):165–97
- Milgrom P. 2009. *The promise and problems of (auction) market design*. Presented at Nemmers Prize Lect., Northwestern Univ.
- Miralles A. 2008. *School choice: the case for the Boston mechanism*. Unpublished manuscript, Boston Univ.
- Papai S. 2000. Strategyproof assignment by hierarchical exchange. *Econometrica* 68:1403–33
- Pathak PA. 2007. *Essays on real-life allocation problems*. PhD diss. Harvard Univ.
- Pathak PA, Sethuraman J. 2011. Lotteries in student assignment: an equivalence result. *Theoret. Econ.* 6:1–17
- Pathak PA, Sönmez T. 2008. Leveling the playing field: sincere and sophisticated players in the Boston mechanism. *Am. Econ. Rev.* 98:1636–52
- Pathak PA, Sönmez T. 2010. *School admissions reform in Chicago and England: comparing mechanisms by their vulnerability to manipulation*. NBER Work. Pap. 16783
- Pycia M, Unver U. 2009. *A theory of house allocation and exchange mechanisms*. Work. Pap., Univ. Calif., Los Angeles
- Roth AE. 1982. The economics of matching: stability and incentives. *Math. Oper. Res.* 7:617–28
- Roth AE. 2002. The economist as engineer: game theory, experimentation, and computation as tools for design economics. *Econometrica* 70:1341–78
- Roth AE, Ockenfels A. 2002. Last-minute bidding and the rules for ending second-price auctions: evidence from eBay and Amazon auctions on the Internet. *Am. Econ. Rev.* 92:1093–103
- Roth AE, Peranson E. 1999. The redesign of the matching market for American physicians: some engineering aspects of economic design. *Am. Econ. Rev.* 89:748–80
- Roth AE, Postlewaite A. 1977. Weak versus strong domination in a market with indivisible goods. *J. Math. Econ.* 4:131–9
- Roth AE, Sotomayor MAO. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econom. Soc. Monogr. Cambridge, UK: Cambridge Univ. Press
- Shapley L, Scarf H. 1974. On cores and indivisibility. *J. Math. Econ.* 1:23–28
- Sönmez T. 1999. Strategy-proofness and essentially single-valued cores. *Econometrica* 67:677–90